

Ubuntu in Artificial Intelligence (AI) Governance: Towards an Inclusive and Democratic Technological Future

Gabriel Kofi Akpah, MA

St. Paul's Interterritorial Major Seminary-Regent-Freetown (Sierra Leone)

Abstract

The rapid development of artificial intelligence (AI) presents new opportunities, but at the same time, it poses significant ethical challenges. In this paper, I explore the potential for Ubuntu—a Southern African philosophy that emphasises community, interconnectedness, and mutual care—to guide AI governance. Ubuntu offers a critical lens through which one can comment on the effect of AI on society, underscoring values such as inclusivity, empathy, and collective well-being. In the future, infusing principles of Ubuntu within the governance of AI will supply a more holistic approach with prime human dignity and social justice at the forefront. I argue that the inclusion of Ubuntu in AI policy and regulation can help lower biases, increase accountability, and ensure transparency in AI systems. By a normative critical approach, I unpack the philosophical underpinnings of Ubuntu, its bearings on contemporary ethical debates in AI, and the potential to transform AI governance. Comparative analyses with existing ethical frameworks underline what is peculiar about the contribution that Ubuntu can make toward democratic engagement and inclusivity in AI development and deployment. I conclude by putting forward some concrete actions for policy decision-makers, technologists, and scholars in taking Ubuntu principles into AI governance, underscoring the fact that global collaboration plays a very integral part in shaping good ethical futures for AI. I thus call for a paradigm shift in this all-inclusive AI ecosystem where technology remains only a means to better human flourishing and social cohesion.

Keywords: artificial intelligence, Ubuntu, AI governance, ethical AI, social justice, human dignity, collective responsibility.

Introduction

In the last few years, artificial intelligence (AI) has changed industries and societies all over the world. From healthcare to finance, education, and entertainment, the idea of AI being able to learn from data, make decisions based on that learning, and even outperform humans at some tasks makes it easy to understand why technologists and society alike are so excited about it. However, this unprecedented progress in AI capabilities has been paralleled by a series of urgent ethical challenges that today's society must overco-

me if it is to responsibly harvest its full potential. Some of these ethical challenges include privacy invasion, algorithmic bias, accountability and transparency, and the threat of work displacement—among others that require urgent attention [Kearns and Aaron, 2019]. These challenges drive home the importance of governance frameworks that will guide the development and deployment of AI technologies that promote human dignity and social justice above anything else, rather than exacerbating existing inequalities.

Current approaches to AI ethics, while essential, tend to reflect predominantly Western individualistic paradigms, which may overlook the relational and communal dimensions of human life. This gap calls for alternative perspectives that prioritise inclusion, empathy, and social cohesion. This paper, therefore, introduces 'Ubuntu,' an African-rooted philosophical theory that is grounded in the maxim "I am because we are," as an alternative framework for AI governance. Ubuntu's focus on community, interconnectedness, mutual

care, and group well-being, along with the acknowledgement of each individual's intrinsic value, provides a comprehensive framework for tackling the ethical dilemmas presented by AI. By applying the concepts of Ubuntu to the design, policy formulation, and regulatory control of artificial intelligence, we can develop governance frameworks that are culturally sensitive, participatory, and clearly geared toward human dignity and social justice.

This paper adopts a normative ethical approach, with references to African communitarian philosophy, to criticise and rebuild modern artificial intelligence governance. At the same time, it takes an applied philosophical approach that translates the moral principles of Ubuntu into policy suggestions. The analysis is placed at the intersection of ethics, technology, and political philosophy, aimed at enhancing a pluralistic and globally informed discourse on artificial intelligence ethics. The discussion follows an order

Photo by Gabriel Heinzer on Unsplash



Artificial Intelligence (AI) and Ethics

The growth in AI technologies has been so rapid that it has raised a number of serious philosophical debates regarding the ethical implications and impact of such technologies on society. Artificial intelligence entails technologies like machine learning, natural language processing, computer vision, and robotics that permit machines to perform tasks associated with human intelligence.

Such technologies have huge potential for many industries by providing efficiency and innovative solutions. They also raise critical ethical challenges that must be addressed for their development and use to be responsible.

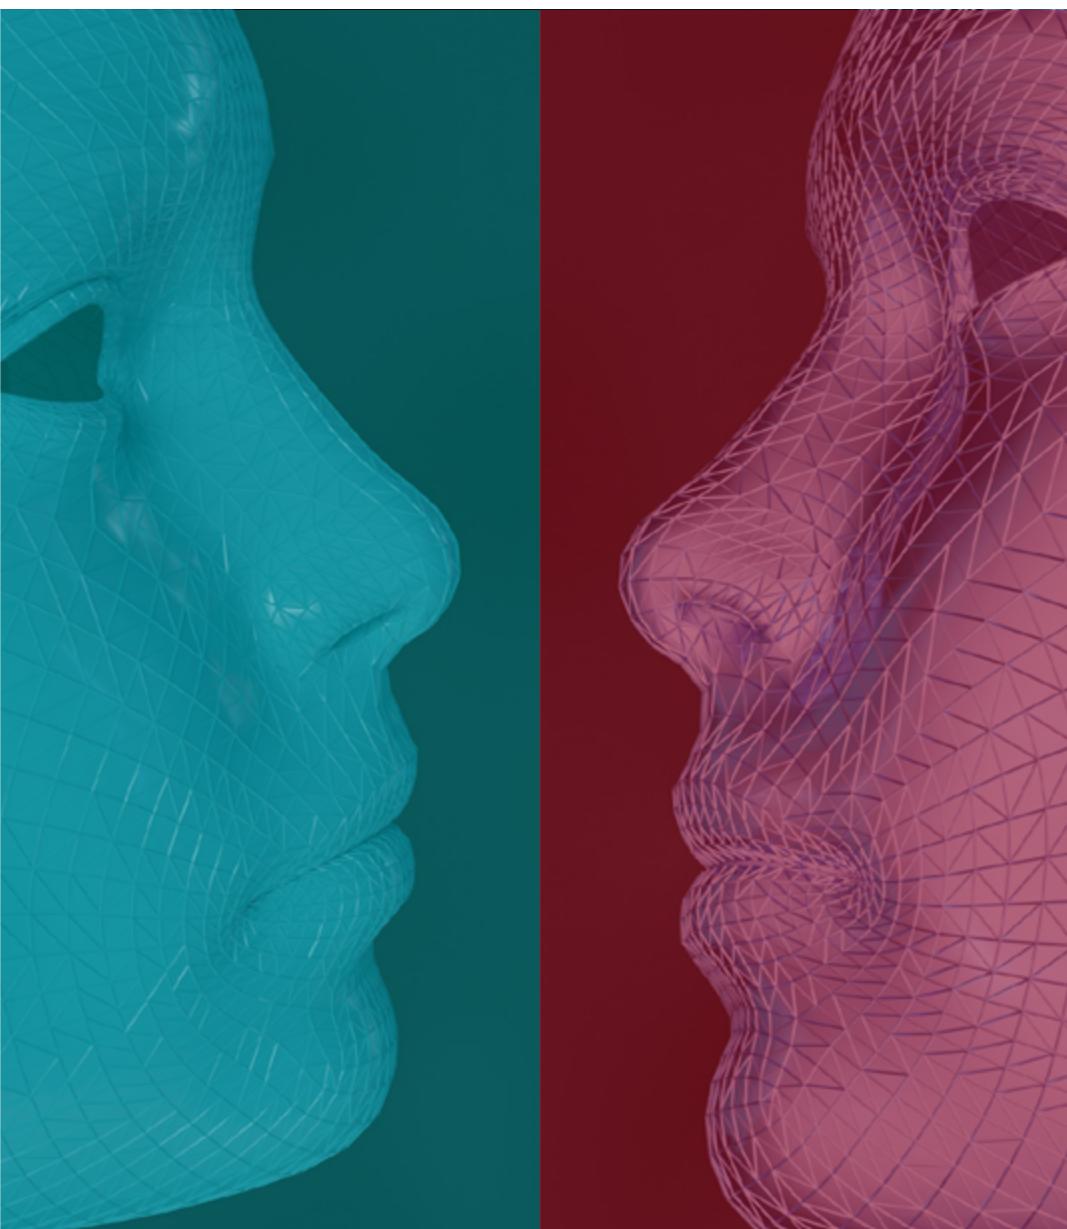
One of the most important ethical issues within AI is bias. AI is trained using data, and if the data is biased, the AI picks it up and amplifies it. This is of particular concern in applications such as hiring and law enforcement, where biased AI systems can lead to discrimination against certain groups. One of the challenges posed by bias in AI is that it brings out the need for careful consideration of training data and the implementation of strategies that can mitigate bias, ensuring AI systems are fair and equitable. As Russell and Norvig explain, “algorithms can only be as good as the data they are trained on, and if that data reflects existing biases, the AI system will, too” [Russell and Norvig, 2016: 568].

Accountability is yet another critical ethical issue. The more autonomous an AI system becomes, the more difficult it is to pinpoint accountability for its actions. Especially in applications like autonomous vehicles or AI-driven medical diagnosis, where mistakes could involve very grave consequences, clear accountabi-

lity frameworks are essential to establish the liability of outcomes on the part of individuals or organisations. Bostrom argues that the development of superintelligent AI presents special challenges of accountability since “the actions of a superintelligent AI could be unpredictable and potentially beyond human control” [Bostrom, 2024: 211].

Another important theme in the ethical discourse around AI is transparency. Many AI systems, especially those based on deep learning, are “black boxes,” making it challenging to understand why they make certain decisions. A lack of transparency might impede understanding, trust, and verification of AI decisions. Improvement in transparency translates to developing methods for interpreting and explaining AI decisions that build trust among users and stakeholders in general, better positioning them to make informed decisions. Russell and Norvig contend that the importance of transparency is underscored by the fact that “interpretable AI systems are essential for ensuring that decisions made by AI are understandable and justifiable” [2016: 603]. These ethical challenges must be addressed as AI technologies evolve in order for their benefits to be reaped with reduced potential harm. This is an interdisciplinary task, one that calls for cooperation among technologists, ethicists, policymakers, and society at large in the development of guidelines and frameworks encouraging the responsible development and use of AI.

By doing so, AI will be harnessed to improve lives without compromising ethical principles.



© Steve Johnson on Unsplash

Ubuntu Philosophy: Foundations and Principles

Ubuntu is a Nguni Bantu expression derived from Southern Africa, which carries immense philosophical depth, often translated as “I am because we are” or “humanity towards others” [Ramosé, 2002]. This philosophy highlights the nature of human beings as interdependent parts of the community, whereby one’s identity, life, and well-being are fundamentally tied to other people’s well-being. It is not only a cultural expression but one that has actively moulded social relations, government, and conflict management in different African societies for ages [Tutu, 1999]. Over the years, Ubuntu has served as an essential pillar for social unity and shared responsibility. In pre-colonial African societies, Ubuntu helped create social peace and constructive collaboration among the people. It steered social behaviour by ensuring that conduct always had a social dimension and rationale [Letseka, 2012]. Its prominence escalated globally during the South African apartheid era, when it served as part of the reconciliation framework post-apartheid. One of the strongest proponents of Ubuntu, Archbishop Desmond Tutu, emphasised its role in mending societal divisions, advocating for the choice to forgive instead of seek vengeance [Tutu, 1999].

Culturally, various proverbs and sayings in Africa capture, embody, and communicate the value of Ubuntu. For example, the Nguni proverb *Umutu ngumuntu ngabantu* translates to “A person is a person through other people.” This emphasises that one’s identity and being are shaped by social links, which supports the notion of communal relationships [Ramosé, 2002]. This communal focus stands in stark contrast to the Western philosophy of individualism,

serving as yet another perspective on humanity and society. Social discourse on ethics and governance has increasingly recognised Ubuntu values as important for inclusivity, empathy, and respect. The Ubuntu approach also helps respond to contemporary issues such as social disparities, violence, and irreparable damage to the environment [Smith & Neupane, 2018]. There is a need to embrace Ubuntu today so that societies can nurture respect for individuality and enhance well-being among their members.

The philosophy of Ubuntu is also underpinned by principles that foster a balanced and just society. Some of the more distinctive ones include communalism, participative decision-making, and consensus building, which dictate social relationships and structures. People tend to achieve their maximum potential in Ubuntu through active participation and contribution to a particular community, instead of setting individualistic goals. Therefore, communalism is the principle of achieving one’s full potential through community [Ramosé, 2002]. In addition, communalism allows individuals to build a sense of belonging and responsibility towards each other, whereby everyone works towards shared goals. The philosophy of communalism can also be seen in the various cooperative practices exercised in Africa. Families and communities work together, strengthening and supporting one another. Furthermore, communalism contributes to more just governance; through its advocates, policies are made to ensure equity of resources and address social disparities [Letseka, 2012]. The proportional representation of particular groups requires strengthening socially and politically

distributive justice. The expected positive effects of enhanced communalism to a greater extent involve moderation in the misuse or overuse of authority. Thus, communalism expects leaders to act more like trustees of the community. This differs from hierarchical and authoritarian frameworks, advocating for a horizontal and participatory system of governance [Ramos, 2002].

Participatory decision-making as an integral aspect of Ubuntu articulates respect for collective opinion and inclusiveness at all levels. This pillar guarantees that every member of the community impacts decisions regarding their lives, which improves accountability and transparency [Smith & Neupane, 2018]. As it is commonly accepted, participatory decision-making means that all relevant groups are invited to discuss and deliberate. This approach improves the decision-making process and cultivates a sense of pride and commitment from the local community. It reduces marginalisation and exclusion risks, ensuring that policies or actions are developed according to the diverse needs and aspirations of the people [Tutu, 1999]. Regarding organisational and governance matters, participatory decision-making can be achieved through community forums, public hearings, and other consulting arrangements that allow direct interaction between decision-makers and the community. These approaches stimulate discussions and negotiations, allowing societies to make decisions that are

acceptable and advantageous to all [Letseka, 2012]. In addition, this model of participatory decision-making expands on democratic values by enhancing fairness, equity, and social justice within society. Ramos [2002] asserts that Ubuntu drives individuals to value others, which in turn enhances collective intelligence and collaboration toward better and more sustainable results.

Furthermore, the concept of consensus-building is directly associated with participatory decision-making under the Ubuntu framework. This approach aims to arrive at agreements that are acceptable to everyone involved, prioritising the group's welfare over individual needs and majority domination [Smith & Neupane, 2018]. It fosters dialogue among the involved parties as they debate and negotiate with each other to identify the best strategies to reach a compromise. These strategies create respect and limit rampant disagreements since decisions are made collaboratively [Letseka, 2012]. In regard to resolving disputes, consensus-building focuses on practices that seek to restore relationships and re-establish structured social orders. It emphasises building trust rather than focusing on punitive actions intended to offer punishment as a means of establishing order among community members [Tutu, 1999]. This strategy resonates well with the focus of Ubuntu, which is centred on forgiving and healing collectively, making it a humane approach instead of the adversarial setting that ju-

stices systems operate in. In governance, consensus-building improves the acceptability and support of policies and initiatives, thereby enhancing their usefulness as well as legitimacy. It fosters ongoing conversations and participation, leading to governance that is flexible and proactive in addressing new issues as they arise [Ramos, 2002]. In addition, consensus is rooted in Ubuntu as a basis for fostering cohesion and long-term stability because decisions stem from shared values and principles accepted by all. This glorifies a cohesive community that can withstand complexities and changes with collective reliance [Smith & Neupane, 2018].

Ubuntu has a very appealing array of humanity to reward us with, based on community development and nurturing through solidarity. The principles of social well-being and harmony are achieved through communalism, participatory decision-making, and consensus-building. Individuals and communities are motivated to act cooperatively as morally guided principles foster dialogue and care beyond self-interest. The interdependence and social responsibility are informed by Ubuntu as a critique of Western individualism. It provides an ethical approach to some of the world's problems, like inequality, climate change, and social fragmentation. Ubuntu is still a philosophy that fortifies Africa and the globe because it aims toward the collective good, and its inclusion in AI governance is not just good but imperative.

AI Governance: Current Challenges and Ethical Imperatives

The healthcare, financial, and educational sectors are being transformed by the recent evolutions made in artificial intelligence over the past few decades. However, the advancements in the usability of AI technologies bring their own sets of problems regarding system governance, particularly in relation to bias, transparency, and accountability. These are only some of the myriad problems that are AI system-specific and require immediate solutions for the creation and application of AI technologies that are beneficial for human society.

The instability of governance with AI systems poses one of the greatest problems to contemporary society: bias. Machine learning models are built using sophisticated algorithms that undergo 'training' using large datasets that often exhibit glaring biases, such as those based on gender, race, and even socioeconomic class. As a result, when applied in the real world, these systems are highly likely to yield biased results. A case in point is the discriminatory error rates found in facial recognition technologies, where some populations, mainly Black people, perform worse than White people [Buolamwini and Gebru, 2018]. Biased algorithms for hiring also tend to work against women and minority candidates, thus worsening existing discrimination in the workplace [O'Neil, 2016].

Bias is just one aspect of the issue that AI system producers have to deal with. Another dimension that poses a myriad of challenges to developers, users, and regulators is the so-called "black box" configuration of numerous AI systems, which makes understanding the decision-making processes and the tools used for enabling tho-

se resolutions nearly impossible. This system's lack of transparency makes it difficult to identify and resolve any biases. Users may not completely understand how the algorithm is reaching its conclusions. Thus, the problem of bias mitigation needs to be addressed in a more comprehensive manner, which includes advanced technological approaches like algorithmic fairness methodologies alongside ethical considerations [Angwin et al., 2022].

Another important concern regarding AI governance is the attention paid to transparency issues.

nical issue; it raises fundamental ethical questions regarding the ability of the systems to be assessed, controlled, and entrusted with responsibility. Clearly defined parameters for AI algorithms must be established to maintain public confidence and safeguard against harm that may be caused by suboptimal algorithms.

Responsibility within AI governance is arguably the most disputable concern. If there are errors or even damage caused by AI systems, whether through unintentional bias or failure to deliver accurately, who is deemed responsible? This inquiry has become particularly acute for autonomous vehicles, AI in healthcare, and military systems. A responsible institution such as Calo [2015] captures the impact of AI on decision-making succinctly: "which of the developers, users, or the AI itself is to bear the responsibility," as it creates a legal and moral sense of vacuum. This suggests that giving AI systems the autonomy to perform decision-making tasks generates intricate problems of responsibility and accountability—more so in legal spheres. If an autonomous vehicle causes an accident, determining liability is not straightforward. Should the driver, who retains control over the vehicle, be held subordinate to the law? Is the developer of the AI system responsible for programming the vehicle's decision-making processes? Thus, there exists a plethora of scenarios where responsibility can be evaded. Judges relying on problematic algorithms for sentencing may grant unjust sentences, yet the absence of an opposing will renders banal claims of fairness and due process irrelevant.

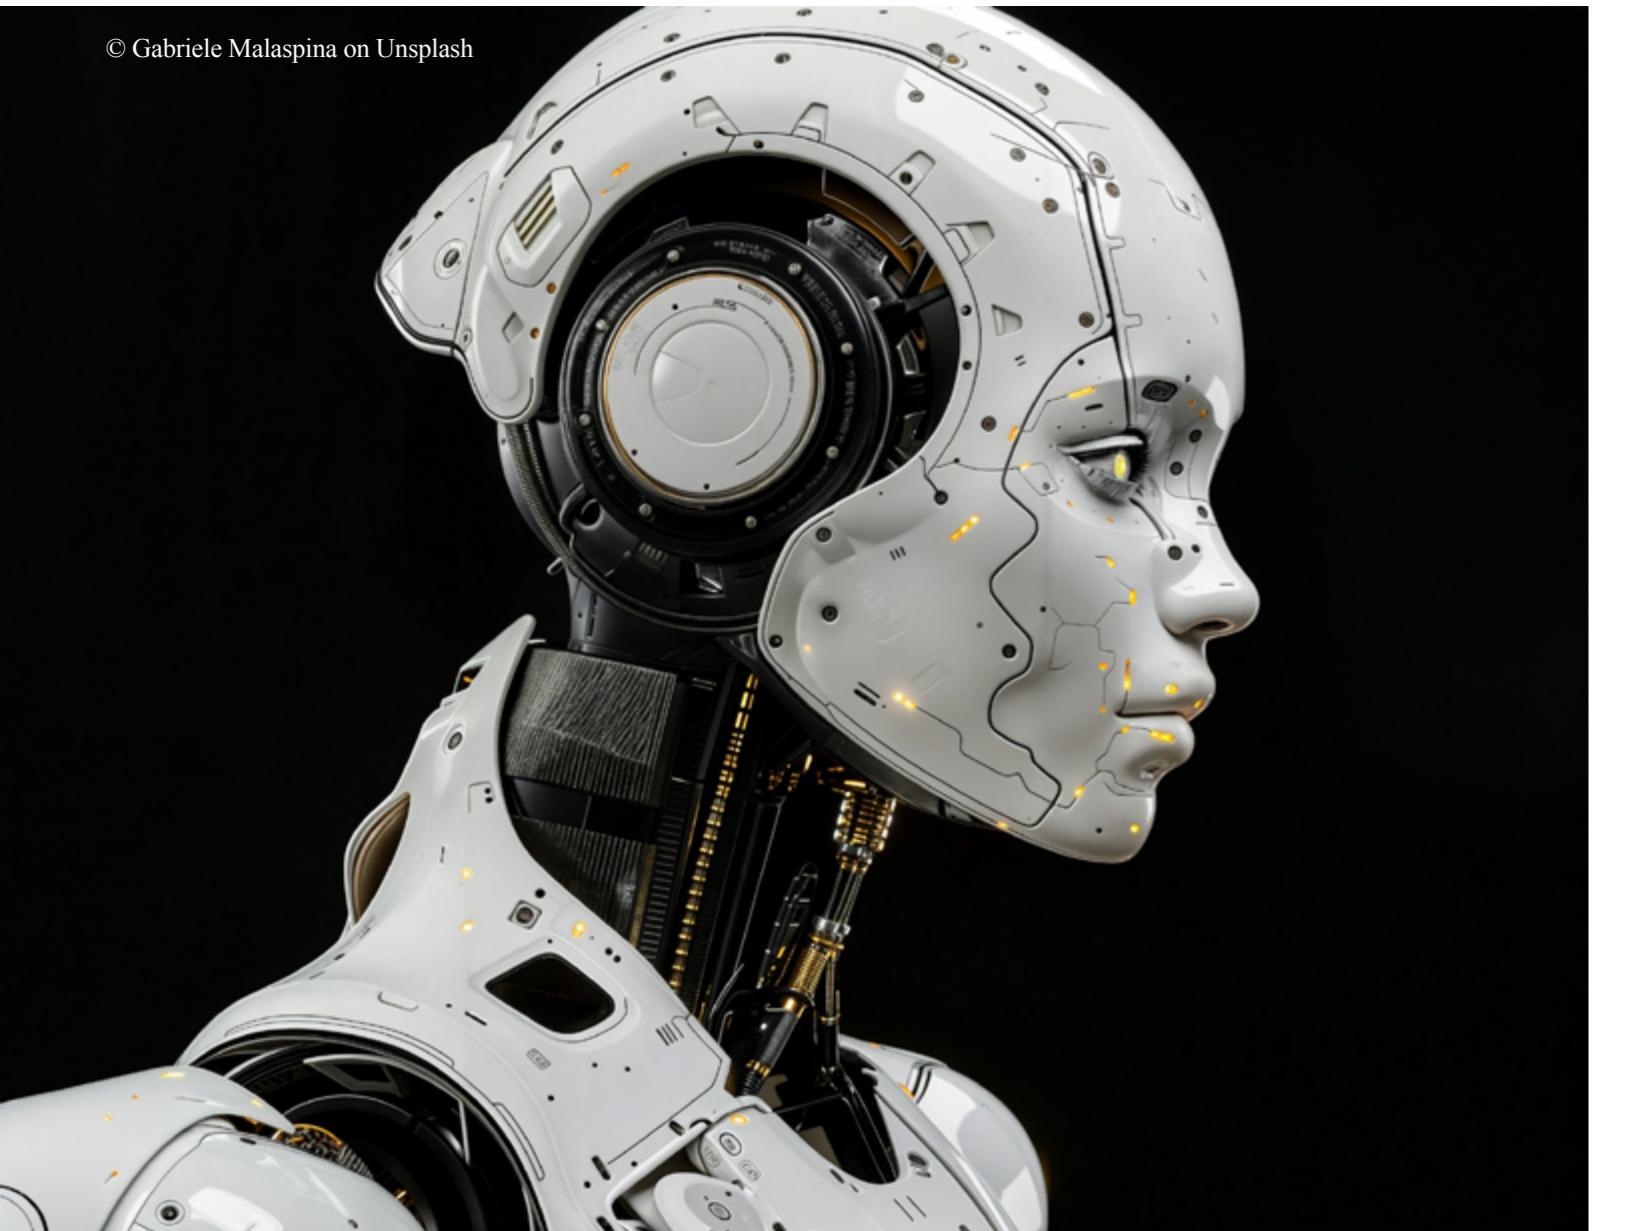
Integrating Ubuntu Philosophy into AI Governance

With the integration of AI systems into society, legal frameworks must be adjusted accordingly in order to set clear delineations of responsibilities to mitigate danger to society. Forward-thinking scholars devote their time and intellect to examining harm that indisputably exists: Goodmans and Flaxman [2017] base their research on liability depending on the extent of foreseeability, human interaction, and evidential clarity of the system. While computer science policy allows for the perfecting of the integration of AI into society, moral and ethical boundaries must always exist to maintain a healthy balance.

Ubuntu is an African philosophy that contests the idea of individualism in Western philosophy and promotes the essence of being human in community with others. From the Southern African perspective, individualism is not in tune with humanity, and this is where Ubuntu comes in. In contrast to individualism, which promotes self-interest, Ubuntu promotes the interests of the community. The intention of this paper is to introduce the philosophy of Ubuntu as a potential paradigm for AI ethics and governance. In particular, we are concerned with the dual problem of fair treatment of individuals and groups and ensuring that

AI technology serves the interests and well-being of humanity as a whole. The philosophical tenets of Ubuntu resonate profoundly with the individual, the community, and society; they explain how one becomes or lives and grows through the community. When applied to AI governance, the concept of Ubuntu offers a fresh perspective on how AI systems should be developed, implemented, and governed. An Ubuntu-inspired perspective is neither seriously naive nor too pessimistic about human nature. It focuses not only on technical efficiency but also on the ethical and social responsibilities of AI developers and users.

© Gabriele Malaspina on Unsplash



AI ethics, as informed by Ubuntu, requires a fundamental rethinking of not just how AI systems are designed, but also how they are deployed and overseen in society. The individualism, profit motive, and absence of community typically associated with technological development are directly at odds with the Ubuntu ethos. Advancing AI governance in a way that's even partially Ubuntu-informed means embracing many of the key principles associated with that African philosophy. These include, but are by no means limited to three core aspects. Firstly, communalism, as opposed to individualism which is a model that's common in many parts of the world, but which is also directly at odds with what happens inside a typical AI system, with its individual instead of collective models of understanding and generating human language. Secondly, respect for human dignity, which is classically associated with Kant and also found in Ubuntu. Finally, making decisions in an inclusive and participatory way, as opposed to top-down decision-making.

To understand how to harness Ubuntu for AI governance, we must first understand its core tenets. For a start, the governance of AI by Ubuntu would require a monumental shift in our thinking. Most modern societies view artificial intelligence predominantly as a means to achieve greater efficiency and profitability. Those societies are, in turn, governed by frameworks that somewhat pay lip service to the notion of these technologies having “positive social impact” - whatever that means. Fairness, transparency, and accountability are terms that pop up all too often in these ostensibly progressive frameworks.

In addition, the emphasis placed by Ubuntu on human dignity and interconnectedness demands that AI systems respect the inherent worth of all people, fostering inclusivity and eschewing anything that would dehumanise or marginalise any population. If AI development is infused with the ethical imperatives of Ubuntu, it will enhance the social responsibility and governance of AI and thereby improve the capacity of AI to serve the people. A key element that distinguishes Ubuntu is its emphasis on inclusive decision-making. In typical African societies, decisions are made in a way that ensures all members have a say. This is not only a moral imperative but also a recipe for creating governance structures that are fair, transparent, and accountable. Why not apply these same principles to AI governance? Issues like bias, transparency, and accountability in AI could use a dose of the good governance principles that Ubuntu advocates.

Implementing Ubuntu in AI governance could lead to the establishment of inclusive governance frameworks that would actively involve all stakeholders in the decision-making processes surrounding AI. With these frameworks in place, it isn't just the developers and policymakers who would have a say; the frameworks would also welcome the wider public into the conversation, including those often marginalized communities who are the first to feel the impact of AI technologies. Creating societies of people who know better is one approach to participatory governance in AI. This could take the form of panels or councils that are semi-deliberative or fully deliberative. A council of this sort, if populated with a broad cross-section of so-

ciety, could serve as an advisory panel or even a regulatory panel, providing a level of oversight to the development and deployment of AI technologies. Whatever the governance structure, the assurance that ethical principles are guiding AI technologies requires a level of dialogue with diverse groups that is far beyond what AI and its societal implications could command even a few years ago. Engaging in this dialogue is itself a societal implication of AI.

Dialoguing ensures that principles rooted in Ubuntu, like those mentioned above, inform the development and deployment of AI systems. In addition, it is possible to structure participatory decision-making in AI governance through open public consultations and feedback mechanisms. These would enable individuals and communities to express concerns and provide perspectives on the social implications of AI systems to be deployed at scale. Public engagement like this not only bolsters trust in AI technologies but also guarantees that their design reflects the sorts of values and needs individuals and communities expect of them.

At the heart of Ubuntu lies the principle of consensus-building. African traditional communities often have lengthy discussions and negotiations to make a decision that involves mutual understanding—that is, an understanding that serves the whole group in a way that benefits them as a community. This reaching of a decision ensures that all perspectives have been considered; it guarantees that the decision is a group decision, not one made by some individual with authority (such as a chief). By using this principle in a computer science context, we

are building a system that behaves more like a community than like a dictatorship. AI governance might find a path to addressing the ethical problems of AI by building consensus. That path would not be straight. It could take a long time. But if it went anywhere, it would go to the kinds of decisions that many people find acceptable and that many different types of stakeholders have had the opportunity to weigh and consider.

Building consensus among stakeholders may be crucial to achieving the design and operation of advanced AI systems in a way that produces good outcomes and avoids harmful ones. Through processes of multi-stakeholder engagement, it is possible to build an institutionalised consensus within AI governance structures. These processes involve working with diverse sets of stakeholders that together form the kinds of dialogue needed for consensus-building and also help identify a more socially inclusive set of governance mechanisms for AI. The inclusion of Ubuntu's principles of governance could help ensure a balance between technological innovation and ethical considerations. Ubuntu stresses not just the

importance of local communities, but the principle that underpins local empowerment: governance. And that's an area where AI lags. Ensuring that communities have a real say in how local, potentially life-altering AI systems are designed and deployed is critical. If we don't, then what will likely happen is that some powerful interests will impose an external technology on a community. And it might be a really powerful technology—like a powerful AI system.

But if the AI system is designed without input from the community, then what's to stop designers from programming in all kinds of biases, just as has happened with some (not all) powerful technologies that came before AI? Equipping local communities to govern AI technologies also means furnishing them with the tools and know-how to understand and engage with AI. This might involve training programs and other educational initiatives that help make the technology and the decision-making around it transparent and understandable to the average community member and local elected official. It's hard to see how a community can participate meaningfully in the deci-

sion-making processes governing the use of powerful technologies like AI if it does not comprehend how the technology works at some basic level.

Applying Ubuntu to AI governance creates a profoundly different kind of framework, one that prioritises community well-being, participatory decision-making, and collective responsibility. It is an opportunity to engage with principles that Ubuntu embodies - fairness, transparency, and respect for human dignity - and to consider how these might be integrated into the AI systems being developed today. The 'ubuntification' of AI governance, then, is as much about kindling a discourse on the Earth that could inclusively involve all as it is about any specific recommendations one might make (for instance, to build governance structures around participatory decision-making, to ensure local communities are empowered, etc.). Though we can't be together with our brothers and sisters in various kinds of communities that AI might affect (or so we hope), we can collectively and communally use their actions and voices to help us make good decisions for all.

Controversies of Ubuntu Philosophy in AI Governance

Ubuntu has become an influential way of re-imagining artificial intelligence (AI) governance. Yet several objections keep arising that question its global applicability, conceptual precision, and practical enforceability. Critics argue that Ubuntu cannot serve as the basis for a transnational AI regime because it is embedded in the communitarian cultures of sub-Saharan Africa and cannot rightfully impose on cultures that value individual autonomy a mo-

rality derived from what might be seen as a tribal ethic [Appiah, 1998; Sen, 1999]. However, comparative moral philosophers reject a strict dichotomy between "collectivist Africa" and "individualist West." Instead, they uncover overlapping relational values across global traditions—Confucian ren, Indigenous North American minobimaatisiiwin, and Catholic social teaching's principle of solidarity [Metz, 2011; Harding, 2020]. Empirical studies of global AI ethics

consultations show broad support for principles such as relational accountability and community benefit, even in liberal democracies [Floridi & Cowls, 2019]. Thus, Ubuntu need not supplant local ethics; it can supply a complementary relational vocabulary that enriches pluralist governance frameworks [Ramosse, 2002].

In addition, the qualitative aspirations of Ubuntu—togetherness and humaneness—seem far too indeterminate to yield enforceable guidelines for the algorithmic tri-fecta of fairness, transparency, and accountability in human-computer interaction [Gyekye, 1997; Gordon, 2013]. But the accusation of vagueness overlook recent jurisprudence and policy instruments that already operationalise Ubuntu-style principles. South Africa's Constitutional Court has used Ubuntu to mould doctrines of restorative justice, data privacy damages, and administrative fairness [Mokgoro, 2015]. On that basis, the African Union's 2022 "Data Policy Framework" translates the kinds of community-centred relational duties emphasised by Ubuntu into concrete safeguards: community-centred impact assessments, collective redress, and algorithmic auditing. Legal scholars thus argue that Ubuntu offers not just principles but resources that can be rendered into statutory language.

Furthermore, others argue that rhetorics of "human dignity" may be co-opted by corporations or states to justify the extraction of data from people, all the while giving the appearance that they respect individuals and are not exercising undue control over them—an appearance that masks the power asymmetries involved. Any normative framework can be captured; the way to prevent that is to

have strong procedures and clear accountability. Ubuntu's insistence on participatory deliberation provides a measure of protection. Aspects of its vision have been tested in two very different settings: multi-stakeholder forums in Kenya's biometric ID review process and the Ghana Agricultural Consortium, and two public-interest data trusts in the USA. These are ways Ubuntu has been tried out in practice. Kenya and Ghana, however, are not the USA or Europe, and even if the level of technology proved sufficient for the trials in these settings, the context in which those trials took place was a very different one. Translating Ubuntu into contexts where privacy, consent, and the public good are understood very differently poses a real risk of creating normative conflicts [Beetham, 2018]. Polycentric governance theory [Ostrom, 2010] counsels that global baseline standards should be layered with protocols that are specific to local contexts. This governance structure is reflected in UNESCO's 2021 "Recommendation on the Ethics of AI." Thus, layering global baseline rights with local context—by what is called "subsidiarity" in governance—makes it possible for rights to influence local contexts. At this level, Ecodharma can guide local impact assessments in the use of AI, while coexisting with global rights instruments such as the ICP-CPR. The plug-and-play simplicity of commercial generative AI systems may seem remarkably incompatible with Ubuntu's widely praised consensus-driven procedures, which favour slow but sure decision-making [Sullivan, 2022].

Digital governance is being tested in Latin America and Europe, and what they show is that the best way to achieve both inclusiveness and speed is to use nested delibe-

ration. This means using small, carefully tuned citizen deliberations to feed recommendations into regulatory processes that are set up to work quickly—what some are now calling regulatory sandboxes. (In these sandboxes, regulatory staff work with businesses and other stakeholders to figure out how best to govern new types of digital services.). Neither narrow-minded nor un-clear, Ubuntu offers a worldwide relational framework that is increasingly reflected in comparative ethical discourses in artificial intelligence. Criticisms of particular cultures (or lack thereof), insufficiently clear concepts, and apparent ease of capture are significant but do not seem to be fatal to the framework. Accountable artificial intelligence increasingly seems to be something that can be enforced both operationally and in a way that is internationally resonant.

Conclusions

The rapid and transformative rise of artificial intelligence (AI) presents both tremendous opportunities and complex ethical dilemmas. As AI becomes an ever-increasing part of the fabric of modern society, its governance demands an approach that not only prioritises technological efficiency but also nurtures human dignity, fairness, and collective well-being. The integration of Ubuntu - a Southern African philosophy that emphasises interconnectedness, community, and mutual care - into AI governance offers a crucial new avenue for addressing these challenges.

The core philosophy of Ubuntu, which emphasises the interconnectedness of all human beings and the importance of community in shaping individual identity, offers a lens through which to critique AI's impact on society. Ubuntu focuses on inclusivity, empathy, and collective responsibility. Consequently, it challenges the individualistic tendencies that often seem to characterise the development and deployment of technologies, including AI. In counterbalancing those individualistic tendencies, Ubuntu asks us to consider, first and foremost, societal values. As increasingly influential decision-making tools, AI systems must either align with those societal values or be seen as a threat to them. Yet AI is inherently value-neutral. Thus, while Ubuntu's influence may counterbalance AI's individualistic biases—which could perpetuate societal and in-group biases—unlike ethical frameworks that focus on individualism, the Ubuntu framework focuses on the community. Its emphasis on the local community encourages

a shift in AI governance toward a much more inclusive model. In the Ubuntu framework, decisions are made with the participation of all affected parties, and there is a strong push toward consensus—even with the many difficult decisions that involve the creation and regulation of AI technologies.

With all these voices in the mix, especially those from marginalised communities, it seems likely that the kinds of insensitivity that have led to the creation of many biased AI systems could be reduced significantly. Adopting the principles of Ubuntu in AI policy and regulation is not without its difficulties. The critiques mentioned throughout this paper - such as the philosophy's cultural specificity, its North-South divide, and the apparent contradiction between its prescribed practices of decentralisation and the centralisation required for coherent global AI governance - must be taken into account. Yet these challenges are not insurmountable. They provide us an opportunity to rethink and fortify, from different cultural standpoints, the principles and practices that are necessary in the AI local global beta. Ubuntu not only offers a framework for ensuring ethics are built into global AI governance but also encourages different stakeholders to engage collaboratively across cultural divides.

In the face of these challenges, I put forward specific actions I would like to see taken by policymakers, technologists, and scholars. First, there should be a serious move afoot to develop AI governance frameworks that incorporate the core principles of Ubuntu, which is, after all,

the African equivalent of a carbon-based life form. And what do those core principles emphasise? Why, transparency, accountability, and fairness, for starters. Second, and this may strike some as a bit too cute, the AI developers and regulators of the future should be encouraged to engage in regular dialogue with a broad range of stakeholders, especially those communities most impacted by AI decisions. Ubuntu, remember, mandates not only consultation but also active involvement in the decision-making process. Third, educational and research initiatives that promote the values of Ubuntu in the design of our technologies and the development of ethical AI should be expanded.

This paper advocates for a paradigm shift in the governance of artificial intelligence, not only in its technical aspects, but broader still in the role technology ought to play.

To conclude, Ubuntu's integration into AI governance presents an excellent opportunity to reshape the discourse concerning technology and society. With its focus on community, fairness, and human dignity, Ubuntu furnishes an ethical foundation that can steer AI systems toward serving the collective good. Implementing this vision is not without its challenges, but with help from around the world and a commitment to inclusivity, we can surely construct an AI ecosystem that reflects, in all its parts and as a whole, the just and equitable society we aspire to. This is a shift whose benefits promise not only a more ethical future for AI but also a more compassionate, socially responsible technological landscape.

References

ANGWIN, J., LARSON J., MATTU S., and KIRCHNER L., 2022. Machine Bias. In *Ethics of Data and Analytics*, 1st ed. Auerbach Publications: 254-264.

APPIAH, K. A., 1998. *The ethics of identity*. Harvard University Press.

BOSTROM, N., 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.

BUOLAMWINI, J. and GEBRU T., 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. *Proceedings of the 1st Conference on Fairness, Accountability, and Transparency*: 77-91.

CALO, R., 2015. Robotics and the Lessons of Cyberlaw. *Calif. L. Rev.* 103: 513.

CHILISA, B., 2012. *Research Methodologies from an Indigenous Perspective*. SAGE Publications.

COWLS, J., et al. 2019. Designing AI for social good: Seven essential factors.

FERGUSON, T. S., 2017. *A course in large sample theory*. Routledge.

GOODMAN, B., and FLAXMAN S., 2017. European Union regulations on algorithmic decision-making and a 'right to explanation'. *AI magazine* 38.3: 50-57.

GYEKYE, K., 1996. *Values of African Culture: An Introduction*. Sankofa Publishing.

GYEKYE, K., 1997. *Tradition and modernity: Philosophical reflections on the African experience*. Oxford University Press.

KEARNS, M., and Roth A., 2019. *The ethical algorithm: The science of socially aware algorithm design*. Oxford University Press.

LETSEKA, M., 2012. In defence of Ubuntu. *Studies in Philosophy and Education* 31: 47-60.

METZ, T., 2011. Ubuntu as a moral theory and human rights in South Africa. *African Human Rights Law Journal* 11.2: 532-559.

MOKGORO, Y., 2015. Constitutional jurisprudence: Ubuntu and the law. *South African Journal on Human Rights* 31(3).

NUSSBAUM, M. C., 2011. *Creating Capabilities: The Human Development Approach*. Harvard University Press.

O'NEIL, C., 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing Group.

RAMOSE, M., 2002. *Ubuntu as African Philosophy*. Mond Books.

RUSSELL, S. J., and NORVIG P., 2020. *Artificial Intelligence: A Modern Approach*. Pearson.

SEN, A., 1999. *Freedom as development*. Oxford University Press.

SHAH, P., et al., 2019. Artificial intelligence and machine learning in clinical development: a translational perspective. *NPJ Digital Medicine* 2.1: 69.

SMITH, M. and NEUPANE S., 2018. Artificial intelligence and human development: toward a research agenda.

TAYLOR, L., Floridi L., and VAN DER SLOOT B., 2021. Group privacy. *Journal of Political Philosophy* 29(1): 68–90.

TUTU, D., 1999. *No future without forgiveness*.