Ubuntu and the Algorithm: Rethinking Ethical Decision-Making in the AI-Driven University

Melisa M. Chawaremera
Institute of Education, Emeris, Faculty of Law, Johannesburg, South Africa

mchawaremera@emeris.ac.za

Abstract: This conceptual paper interrogates the growing use of algorithmic decision-making systems within tertiary education through the philosophical lens of Ubuntu. As universities increasingly integrate artificial intelligence into institutional processes such as student admissions, learning analytics, academic risk profiling as well as automated assessment, algorithmic systems are frequently presented as objective and efficient mechanism for educational governance. This paper challenges these assumptions by arguing that many AI-driven systems are embedded within Western-centric epistemological traditions that privilege quantification and standardisation while overlooking the relational and communal dimensions of human development that are central to African thought. Drawing upon Afrocentric theory, Ubuntu philosophy as well as contemporary scholarship on algorithmic governance, this paper argues that the rapid adoption of algorithmic decision-making in higher education raises critical ethical concerns regarding accountability and human centred educational engagement that is rooted in context. In particular, this paper contends that algorithmic systems may unintentionally reproduce historical inequalities and institutional exclusions while masking such outcomes beneath the appearance of technological objectivity. The discussion positions Ubuntu as more than a principle. Instead, it positions it as an alternative epistemic framework that is capable of reshaping how ethical decision-making in tertiary education is conceptualised. Through its emphasis on relational personhood and contextual sensitivity, Ubuntu offers a critical lens through which limitations of data driven educational governance may be exposed and reimagined. This paper therefore proposes relational accountability as an Ubuntu-informed approach to algorithmic governance in higher education. Through its emphasis on meaningful human oversight and contextual review, relational accountability offers a practical framework for ensuring that technological innovation remains responsive to the students and communities it serves. All in all, the paper argues that the future of artificial intelligence in tertiary education should not be measured solely by efficiency and predictive capability, but by the extent to which technological systems preserve the human dignity and relational ethics that education itself is intended to cultivate.

Keywords: Algorithmic Decision-Making (ADM); Higher Education; Neutrality; Ubuntu.

1. Introduction

Universities have always been institutions of judgment. Decisions concerning admission and progression shape not only individual educational trajectories but also broader societal understandings of merit and success. Historically, these decisions have been exercised by human actors operating within institutional frameworks that, despite their imperfections, remained capable of engaging with context, discretion and the complexities of human experience. Increasingly, however, such judgments are being delegated to algorithmic systems that promise greater efficiency. The rapid integration of artificial intelligence (AI) into higher education has accelerated this transformation. Universities across the globe now employ algorithmic systems to support student admissions and academic risk profiling (Selwyn 2019). These technologies are frequently presented as neutral instruments capable of improving institutional decision-making through data-driven insights. Their appeal is understandable. As they are faced with growing student populations and persistent resource constraints, higher education institutions are under pressure to make decisions more quickly and at scale. However, the growing reliance on algorithmic decision-making (ADM) raises important questions regarding the nature of educational governance itself. While algorithms may be capable of identifying patterns and generating predictions from vast quantities of data, they simultaneously influence how institutions recognise and respond to human beings. As Amoore (2020) observes, algorithmic systems do not merely analyse reality, as they also actively participate in constructing it by shaping what becomes measurable. Decisions that once relied upon human interpretation have increasingly become matters of statistical inference as well as computational prediction. This shift reflects broader developments within what Rouvroy and Berns (2013) describe as algorithmic governmentality (a mode of governance in which behavioural data and predictive analytics increasingly inform institutional decision-making). It is therefore within such systems that students become intelligible through datasets and risk indicators rather than through the rich social and personal contexts that shape educational experience. Consequently, there is a growing tendency to privilege what can be measured over what can be understood (Chawaremera 2025). The assumption that algorithmic systems are objective and impartial has been the subject of significant scholarly criticism. Pasquale (2015), Eubanks (2018), Noble (2018) and Benjamin (2019) have demonstrated that algorithms are neither neutral nor detached from the social environments in which they are designed and deployed. More recently, Thalpage (2023) similarly notes that

increasingly complex AI systems often operate through opaque processes that undermine accountability and meaningful scrutiny. The reality on the ground is that they frequently reproduce historical inequalities and normative assumptions while simultaneously obscuring them beneath the appearance of technological objectivity. The widely cited Amazon recruitment algorithm illustrates this concern. Although designed to improve efficiency in hiring, the system learned from historical recruitment data that reflected existing gender disparities and consequently penalised applications associated with women (demonstrating how algorithmic systems can inadvertently encode and perpetuate historical patterns of exclusion while appearing neutral and objective) (Dastin 2018). In educational contexts, this raises important concerns regarding fairness and the possibility that institutional exclusions may be reinforced through automated systems that are difficult to interrogate or challenge. As a result, existing frameworks often provide limited guidance when evaluating the social and communal consequences of ADM within contexts characterised by historical inequality, epistemic marginalisation and ongoing struggles for educational transformation. This concern resonates with Ndlovu-Gatsheni's (2021) critique of the "cognitive empire", wherein dominant systems of knowledge frequently universalise particular epistemological assumptions while marginalising alternative ways of knowing. Viewed in this light, the challenge presented by ADM in African higher education is not merely technical, but also epistemic. *Ubuntu* offers an alternative point of departure. It is often expressed through the notion that a person becomes a person through other persons and *Ubuntu* advances a relational understanding of human existence in which dignity, as well as moral responsibility emerge through interconnectedness and communal belonging (Ramose 2002; Eze 2010; Metz 2011). Rather than viewing human beings as isolated and autonomous actors, *Ubuntu* emphasises on relationships and contextual understanding as the foundations of ethical life. Such a perspective presents a significant challenge to algorithmic systems that rely upon abstraction and standardisation to produce decisions affecting human lives. This paper argues that prevailing approaches to ADM in higher education remain epistemologically incomplete because they are largely informed by Western-centric conceptions of rationality. While these values undoubtedly contribute to institutional effectiveness, they fail to adequately account for the relational dimensions of human development that are central both to educational practice and to African philosophical thought. Consequently, ethical debates concerning AI in higher education require a broader framework capable of engaging with both the questions of technical accuracy and also with questions of relational justice. By drawing upon *Ubuntu* philosophy and literature on algorithmic governance, this paper advances *Ubuntu* as an alternative epistemic framework through which the ethical implications of AI-driven decision-making in tertiary education may be critically examined. Rather than the extreme approach of rejecting technological innovation, this paper advocates for a more contextually responsive and relational approach to educational governance by way of proposing that the legitimacy of AI in higher education should be measured by the extent to which technological systems preserve the relational ethics and contextual sensitivity that education itself is intended to cultivate.

2. Methodology

This study adopts a qualitative conceptual desktop research methodology. It does not seek to generate empirical findings, but rather to critically synthesise existing scholarship in order to interrogate the ethical implications of ADM in higher education and to develop an Ubuntu informed account of relational accountability. Sources were selected on the basis of their direct relevance to the research question and contribution to the conceptual themes developed in the paper. Preference was given to peer reviewed journal articles and recent literature on AI and higher education, while seminal works were retained where necessary to establish the theoretical foundations of algorithmic governance and Ubuntu philosophy. Literature that dealt exclusively with technical aspects of AI without engaging questions of governance, ethics, education or human impact was generally excluded. The selected literature was analysed thematically around three intersecting concerns which include (i) the claimed neutrality and objectivity of algorithmic systems; (ii) the relational and contextual limitations of data-driven educational decision-making; and (iii) the potential contribution of Ubuntu to ethical governance in higher education. The analysis is therefore interpretive rather than systematic.

3. The Rise of ADM in Higher Education

The increasing integration of AI into higher education forms part of a broader global shift towards data-driven governance. Across both public and private sectors, organisations are increasingly relying on algorithmic systems to assist with complex decision-making processes that were historically exercised by human actors. Universities have not remained immune to this transformation. As institutions face growing demands for efficiency and student success, ADM systems have emerged as attractive tools for managing increasingly complex educational

environments. ADM refers broadly to the use of computational systems to assist, recommend or automate decisions through the analysis of large quantities of data (Yeung 2018). Rather than relying solely on human judgment, ADM systems identify patterns, they generate predictions and then they produce recommendations that influence institutional action (Yeung 2018). In higher education, these systems are increasingly embedded within routine academic and administrative functions. One of the most prominent applications of ADM is learning analytics. Learning analytics platforms collect and analyse data generated through students' interactions with institutional learning management systems, digital libraries, discussion forums and online assessments (Ngulube and Ncube 2025). Through these generated analyses, universities seek to identify patterns of engagement, predict academic performance and flag students who may be considered at risk of failure or attrition. Advocates argue that such systems enable earlier intervention and more targeted student support. Similarly, predictive analytics are increasingly used to support student retention initiatives. By analysing historical academic records, attendance patterns and engagement metrics, universities can generate risk profiles that estimate the likelihood of a student discontinuing their studies (Selwyn 2019). These predictions often inform institutional decisions regarding academic support and resource allocation. In some contexts, algorithmic systems are also utilised during admissions processes to assist in evaluating applicants and forecasting future academic success. The adoption of these technologies is frequently justified through claims of efficiency and evidence-based decision-making. Universities are under increasing pressure to demonstrate measurable outcomes and improve retention rates as well as optimise the use of limited resources. Within this context, algorithmic systems are often portrayed as rational and neutral mechanisms capable of reducing human error while enhancing institutional effectiveness. However, the growing presence of ADM in higher education signifies more than a technological development. If anything, it reflects a broader transformation in how educational institutions understand students and institutional success. Educational processes that were once characterised by contextual judgment are increasingly mediated through systems that prioritise quantification and data-driven classification. Students are no longer understood solely through their lived educational experiences but are increasingly represented through digital traces and predictive probabilities. This transformation does not necessarily render algorithmic systems undesirable. Undeniably, many of these technologies have the potential to support student success and improve institutional responsiveness. Nevertheless, their growing influence raises important questions regarding the assumptions embedded within these systems and the values they privilege. Before considering how alternative philosophical frameworks such as *Ubuntu* may contribute to these debates, it is firstly imperative to interrogate the (problematically) widely accepted belief that algorithmic systems operate as neutral instruments of educational governance.

4. The Myth of Algorithmic Neutrality

The widespread adoption of ADM within higher education has been accompanied by a persistent assumption that computational systems are inherently more objective than their human counterparts. Universities frequently justify the deployment of predictive analytics, automated assessment tools and risk-profiling systems on the basis that algorithms are capable of reducing human bias as well as improving institutional efficiency. While such claims possess intuitive appeal, they often obscure a more fundamental reality which is that algorithms do not operate independently of human values, assumptions or even social contexts. The notion of algorithmic neutrality has attracted significant criticism within contemporary scholarship. O'Neil (2016) famously describes algorithms as "opinions embedded in mathematics", emphasising that computational systems are shaped by a series of human decisions concerning which variables matter, which outcomes should be prioritised, and which forms of behaviour are deemed desirable. Consequently, algorithmic outputs should not be understood as purely objective reflections of reality as they are also, as a matter of fact, products of underlying assumptions and value judgments embedded within their design. The now infamous Amazon recruitment algorithm illustrates this concern. While it was developed to streamline hiring processes, the system was trained on historical recruitment data that reflected the technology sector's longstanding preference for male candidates. As a result, the algorithm learned to penalise applications associated with women and to favour characteristics commonly found in the profiles of previously successful male applicants, thereby reproducing historical patterns of exclusion while maintaining the appearance of neutral and data-driven decision-making (Dastin 2018). Every stage of algorithmic design, from data collection and model construction to implementation and evaluation, reflects normative choices that inevitably influence the outcomes produced by the system. Therefore, once again, algorithmic systems cannot be understood as neutral observers of reality. It is important to note that they are socio-technical artefacts that both reflect and reproduce the assumptions embedded within the environments from which they emerge (Benjamin 2019; Noble 2018). Historical datasets often contain traces of existing institutional preferences and historical exclusions. When such data are used to train

predictive systems, past patterns may become embedded within future decision-making processes. In this sense, algorithms participate in shaping institutional realities by determining which characteristics become visible and actionable. Within higher education, this concern is particularly significant. Learning analytics platforms, predictive risk models and automated assessment systems rely upon the conversion of complex educational experiences into quantifiable indicators. Student engagement becomes measurable through clicks, logins and participation metrics while academic success becomes associated with behavioural patterns identified within historical datasets. Also sadly, risk becomes transformed into a statistical probability. While such quantification may offer useful institutional insights, it simultaneously narrows the lens through which students are understood. This tendency reflects what Pasquale (2015) describes as the problem of the "black box", whereby increasingly influential algorithmic systems operate in ways that remain inaccessible to those affected by their decisions. Students may be identified as academically at risk without understanding how such classifications were reached over and above a superficial, surface level generic reason. Also, educators may receive predictive insights without meaningful opportunities to interrogate the assumptions underpinning them. As a result, decisions that carry significant educational consequences can acquire an aura of objectivity simply because they emerge from computational processes. The risks associated with such systems are not merely theoretical. Eubanks (2018) demonstrates how automated decision-making in public services has frequently reinforced existing forms of social disadvantage under the guise of efficiency and evidence-based governance. Similarly, Noble (2018) illustrates how seemingly neutral technologies can reproduce and amplify broader societal inequalities through the data and assumptions upon which they rely. These insights are particularly relevant to higher education institutions that increasingly rely upon algorithmic systems to enrol, identify, categorise and evaluate students.

It is pertinent for this paper to note that the challenge is not simply that algorithms may produce biased outcomes. Human decision-makers are themselves susceptible to bias, inconsistency and error. The deeper concern is that algorithmic systems often conceal their normative foundations behind the language of technical objectivity. As Amoore (2020) argues, algorithmic systems actively construct the conditions through which particular futures become more likely than others. The thresholds and risk indicators embedded within these systems shape institutional responses and influence the opportunities available to individuals (mostly students). For higher education institutions, this raises an important question which is to ask what conception of the student is embedded within ADM systems. Predominantly data-driven approaches tend to conceptualise students through predictive probabilities and individual performance indicators. Such representations may be useful for administrative purposes, yet they risk overlooking the relational, contextual and communal dimensions of educational experience that cannot easily be captured through data analytics alone. The danger is not merely theoretical. If an admissions algorithm is trained on historical institutional data, it may learn to recognise privilege as potential and disadvantage as risk. In the same way that Amazon's recruitment algorithm internalised historical gender biases embedded within past hiring decisions (Dastin 2018), an admissions system operating within a historically unequal educational landscape could inadvertently reproduce racial and socio-economic disparities while presenting its outcomes as neutral and evidence-based assessments of academic suitability. It is at this point that the limitations of prevailing approaches to algorithmic governance become most apparent. Existing debates often focus on technical solutions such as transparency and fairness. While these interventions remain important, they generally operate within the same epistemological framework that produced the problem in the first place. They seek to improve algorithmic systems without fundamentally questioning the assumptions about knowledge, personhood and human development that underpin them. The challenge therefore extends beyond technical accountability and requires a deeper interrogation of the philosophical foundations upon which ADM rests. To undertake such an inquiry, this paper turns to Ubuntu as an alternative epistemic framework.

5. *Ubuntu* as an Alternative Epistemic Framework

5.1 *Ubuntu* and Relational Personhood (Beyond Individualism)

The limitations of contemporary approaches to algorithmic governance invite a deeper question on what conception of the human person underpins decision-making in higher education. Dominant AI ethics frameworks are largely grounded in liberal traditions that prioritise individual autonomy and procedural fairness. While these principles remain important, African philosophical traditions offer an alternative point of departure. While Ubuntu is not a singular or uncontested philosophical tradition, scholars such as Ramose (2002), Metz (2011) and Eze (2010) broadly converge in emphasising relationality, interconnectedness and communal belonging as

central features of African conceptions of personhood. Ubuntu, often expressed through the maxim *umuntu ngumuntu ngabantu*, rejects the notion of the individual as an isolated and self-sufficient entity. Instead, personhood emerges through relationships with others and participation within a broader community (Ramose 2002; Mokgoro 2017). Human beings are not merely autonomous actors but relational beings whose dignity and identity are constituted through mutual recognition and interconnectedness. As Eze (2010) argues, *Ubuntu* is not simply a moral principle but a worldview that foregrounds relationality as the foundation of human existence. This perspective presents a direct challenge to algorithmic systems that rely upon abstraction and standardisation. Where ADM frequently seeks to reduce individuals to measurable characteristics and behavioural indicators, *Ubuntu* insists that human beings cannot be understood apart from the relationships and contexts that shape their lives.

5.2 Educational Decision-Making Through an *Ubuntu* Lens

The relevance of *Ubuntu* becomes particularly apparent within higher education. Educational achievement is rarely the product of individual effort alone. It is shaped by family circumstances, community support, cultural identity, language, historical disadvantage and numerous other social realities that resist quantification. From an *Ubuntu* perspective, educational institutions are not merely sites of knowledge transmission but communities of human development. Waghid (2020) argues that higher education should cultivate dialogue, mutual engagement and democratic participation rather than simply produce economically productive graduates. Similarly, Metz (2011) contends that ethical action is ultimately concerned with fostering harmonious relationships characterised by solidarity and mutual respect. Viewed through this lens, the increasing reliance on ADM raises important concerns. Predictive analytics may identify behavioural patterns associated with academic performance, but they cannot fully account for the social and historical realities that shape educational experience. Risk scores and engagement metrics may provide useful information, yet, they offer only a partial representation of the student. Consequently, there is the very sad danger that students become known through data rather than understood through context.

5.3 Relational Accountability

Contemporary debates concerning AI governance have largely focused on technical mechanisms such as transparency, fairness, explainability and auditability. These interventions are indeed both valuable and necessary. However, they remain primarily concerned with improving how algorithmic systems function rather than interrogating the assumptions about human development that underpin them. *Ubuntu* broadens this discussion by introducing the concept of relational accountability. From this perspective, the legitimacy of a decision cannot be assessed solely by its procedural correctness or predictive accuracy. Rather, decisions must also be evaluated according to whether they preserve dignity, recognise context and contribute to communal flourishing. This shift is particularly significant in higher education, where decisions frequently shape opportunities, aspirations and life trajectories. An algorithm may be able to explain how a decision was reached, but explainability alone does not guarantee justice. A decision may be technically accurate while remaining socially or contextually harmful. *Ubuntu* therefore directs attention beyond technical accountability and more towards the relational consequences of institutional decision-making. The significance of *Ubuntu* lies in providing an alternative framework through which AI's role within higher education can be critically evaluated. This research submits that *Ubuntu* exposes dimensions of educational governance that prevailing approaches to algorithmic accountability frequently overlook. In essence, it is therefore from this foundation that a more human-centred approach to ADM may be developed.

6. Relational Accountability in ADM

The preceding discussion has demonstrated that prevailing approaches to algorithmic governance in higher education often emphasise technical accountability through mechanisms such as transparency and fairness. While these interventions remain necessary, this paper submits that they are insufficient on their own. They focus primarily on how algorithmic systems function rather than on how decisions affect the individuals and communities subject to them. If Ubuntu is taken seriously as an alternative epistemic framework, then accountability must extend beyond technical performance and incorporate relational considerations. This paper therefore proposes the concept of relational accountability as a complementary approach to ADM in higher education. Relational accountability recognises that educational decisions do not occur in a social vacuum. Decisions concerning admissions, progression, academic support and assessment affect human beings whose lives are shaped by complex social, cultural and historical realities. Consequently, the legitimacy of algorithmic

systems should not be measured solely by predictive accuracy or procedural consistency, but also by their capacity to recognise context and remain responsive to those affected by their outputs.

6.1 Human Oversight Beyond Administrative Approval

The first dimension of relational accountability requires meaningful human oversight. While many universities continue to place humans "in the loop", such oversight frequently becomes little more than administrative endorsement of algorithmic recommendations. There is a significant difference between reviewing a decision and critically interrogating it. An *Ubuntu*-informed approach requires decision-makers to remain active participants in the evaluative process. Educators and administrators should therefore retain the authority to question and, where necessary, override algorithmic outputs. This research submits that human judgment should not function as a procedural formality but as an essential safeguard against the limitations of data-driven classification. The role of the human decision-maker is therefore not merely to approve what the system has produced, but rather, to also ensure that individual circumstances and contextual realities remain visible within the decision-making process.

6.2 Contextual Review as a Safeguard Against Reductionism

The second dimension concerns contextual review. Algorithmic systems necessarily simplify reality in order to generate predictions and classifications. However, educational experiences are rarely reducible to measurable indicators alone (Chawaremera 2025). Family responsibilities, financial hardship, language barriers, community obligations and historical disadvantage frequently shape student outcomes in ways that cannot be adequately captured through datasets. Relational accountability therefore requires institutions to establish mechanisms through which algorithmic outputs may be evaluated against contextual realities. Rather than treating predictive models as definitive representations of student potential, universities should regard them as decision-support tools whose outputs remain open to scrutiny and interpretation. Context should not be treated as an exception to ADM but as an essential component of educational governance.

6.3 Participatory Governance and Shared Responsibility

The third dimension concerns participation. Decisions regarding the design, deployment as well as the evaluation of algorithmic systems are often concentrated within institutional leadership structures or outsourced to commercial technology providers. Students and educators frequently have limited opportunities to influence how these systems operate despite being directly affected by their consequences. *Ubuntu*'s emphasis on communal responsibility suggests that algorithmic governance should be approached as a shared institutional undertaking rather than a purely technical exercise. Universities should therefore create opportunities for meaningful participation by students, academic staff and other stakeholders when determining how algorithmic systems are implemented and evaluated. Such participation enhances legitimacy while also ensuring that technological innovation remains responsive to the needs and values of the communities it serves. Taken together, these principles shift the focus of algorithmic governance from technical optimisation towards relational responsibility. They actually do not reject AI, nor do they deny the potential benefits of data-driven decision-making. But instead, they seek to ensure that technological systems remain subordinate to the educational purposes they are intended to serve. In this regard, relational accountability offers a practical framework through which universities may pursue innovation without sacrificing the dignity and context of those whose lives are shaped by institutional decisions.

7. Conclusion

The growing use of ADM in higher education represents a significant shift in how universities govern admissions, assessment, student support and academic progression. While these technologies offer clear benefits in terms of efficiency and scalability, they also raise important questions regarding the values that underpin educational decision-making. This paper has argued that prevailing approaches to algorithmic governance remain epistemologically incomplete because they are largely grounded in frameworks that prioritise quantification and procedural correctness. Although transparency, explainability and fairness remain important, they do not fully account for the relational and contextual dimensions of human development that are central to educational practice. Through drawing upon *Ubuntu* as an alternative epistemic framework, this paper demonstrated that educational decisions cannot be reduced to data points and predictive probabilities alone. Students exist within social, cultural and historical contexts that shape their educational experiences in ways that frequently resist

quantification. Consequently, the legitimacy of ADM should be evaluated not only according to technical performance but also according to its capacity to preserve human dignity and recognise context of those affected by its outcomes. To address this challenge, the paper proposed relational accountability as an Ubuntu-informed approach to algorithmic governance in higher education. At the end of the day, the future of AI in higher education should not be determined solely by what algorithms can do, but by what universities ought to become. The true measure of innovation will not lie in the sophistication of the algorithm, but in the extent to which technology remains accountable to the humanity that is actually intended to serve.

Ethics Declaration

The author declares that ethical clearance was not required for this conceptual paper, as it relies solely on the analysis of existing, publicly available academic literature and does not involve primary data collection from human or animal subjects.

AI Declaration

Zotero was used for the purposes of managing sources and generating the in-text references and reference list. Generative AI was used to support language editing during finalisation of the manuscript, particularly to trim text to the required word limit without losing meaning of core arguments. The author takes full responsibility for the final manuscript

Acknowledgements

The author wishes to acknowledge the Emeris (Ruimsig Campus) Academic Team and fellow Heads of Faculty for the many thoughtful discussions that continue to deepen our collective understanding of teaching, learning and the institutional systems that support student success. Particular appreciation is extended to the Office of the Dean of Academic Support at Emeris for consistently encouraging academic leaders to engage critically with student results and institutional data rather than relying solely on automated outputs. Their continued emphasis on interrogating data, understanding context and, as Margi Boosey often reminds us, to "dig deeper", has informed the thinking that underpins this paper.

References

Amoore, L. (2020) *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham, NC: Duke University Press. doi:10.1515/9781478009202.

Benjamin, R. (2019) Race After Technology: Abolitionist Tools for the New Jim Code. Cambridge: Polity Press.

Chawaremera, M.M. (2025) 'Algorithmic Teaching, Fading Thought? Rethinking Engagement in the AI Era', in *Proceedings of the 2nd International Conference on Education Research*. doi:10.34190/icer.2.1.4008.

Dastin, J. (2018) 'Amazon scraps secret AI recruiting tool that showed bias against women', Reuters, 10 October. Available at: https://www.reuters.com/article/world/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G / (Accessed: 15 June 2026).

Eubanks, V. (2018) Automating Inequality: How High-Tech Tools Profile, Police and Punish the Poor. New York: St Martin's Press.

Eze, M.O. (2010) Intellectual History in Contemporary South Africa. New York: Palgrave Macmillan.

Metz, T. (2011) 'Ubuntu as a moral theory and human rights in South Africa', African Human Rights Law Journal, 11(2), pp. 532–559. Available at: https://www.ahrlj.up.ac.za/metz-t.

Metz, T. and Gaie, J.B.R. (2010) 'The African ethic of Ubuntu/Botho: Implications for research on morality', Journal of Moral Education, 39(3), pp. 273–290. doi:10.1080/03057240.2010.497609.

Mokgoro, Y. (2017) "Ubuntu and the law in South Africa", *Potchefstroom Electronic Law Journal*, 1(1), pp. 16–32. doi:10.17159/1727-3781/1998/v1i1a2897.

Ndlovu-Gatsheni, S. J. (2021). The cognitive empire, politics of knowledge and African intellectual productions: Reflections on struggles for epistemic freedom and resurgence of decolonisation in the twenty-first century. *Third World Quarterly*, *42*(5), 882-901.

Ngulube, P. and Ncube, M.M. (2025) 'Leveraging learning analytics to improve the user experience of learning management systems in higher education institutions', *Information*, 16(5), p. 419. doi:10.3390/info16050419.

Noble, S.U. (2018) Algorithms of Oppression: How Search Engines Reinforce Racism. New York: New York University Press. doi:10.2307/j.ctt1pwt9w5.

O'Neil, C. (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishing Group.

Pasquale, F. (2015). *The Black Box Society, the Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press. https://doi.org/10.4159/harvard.9780674736061

Ramose, M.B. (2002) African Philosophy Through Ubuntu. Harare: Mond Books.

Rouvroy, A., & Berns, T. (2013). Algorithmic Governmentality and Prospects of Emancipation: Disparateness as a Precondition for Individuality. *Réseaux, 31,* 163-196. https://doi.org/10.3917/res.177.0163

Selwyn, N. (2019). *Should robots replace teachers? AI and the Future of Education*. (1st ed.) Polity Press.

Thalpage, N.S. (2023) 'Unlocking the black box: Explainable artificial intelligence (XAI) for trust and transparency in AI systems', Journal of Digital Art & Humanities, 4(1), pp. 31–36. doi:10.33847/2712-8148.4.1_4.

Waghid, Y. (2020) *Towards a Philosophy of Caring in Higher Education: Pedagogy and Nuances of Care*. Cham: Springer.

Yeung, K. (2018) 'Algorithmic regulation: A critical interrogation', Regulation & Governance, 12(4), pp. 505–523. doi:10.1111/rego.12158