
Algorithmic Epistemological Incongruence: A Conceptual Framework for Rethinking Fairness in AI-Enabled Human Resource Management

Review Article | Volume 1 | Issue 2 | 2026 | Article Number: 010

Received: 19 June 2026 | Accepted: 25 June 2026 | ISSN: 2979-8582 (Online)



Michael Raphael Onenyi¹, Yunusa Ademu (PhD)¹

¹Faculty of Management Sciences, Department of Business Administration, Prince Abubakar Audu University, Anyigba, Kogi State, Nigeria

 MRO [0009-0005-7272-4483](https://orcid.org/0009-0005-7272-4483)

 YA [0009-0003-3112-1112](https://orcid.org/0009-0003-3112-1112)

Correspondence: Michael Raphael Onenyi, onenyi.mr@ksu.edu.ng

Abstract

The algorithmic fairness literature has produced valuable tools for detecting bias in AI-powered human resource management (HRM) systems. However, these tools rest on an assumption that leaves one deployment configuration analytically under-specified: that bias arises from errors within a shared evaluative framework. As AI-HRM systems designed in Western contexts are applied to the workers whose competence is constituted in a different epistemological tradition, a different form of misrecognition may emerge that cannot be addressed by data augmentation and calibration of fairness metrics. This paper introduces Algorithmic Epistemological Incongruence (AEI) to name that condition. AEI builds upon construct validity theory (Messick, 1995; Cronbach and Meehl, 1955) by shifting the unit of analysis beyond measurement adequacy to evaluative ontology alignment, posing not whether the AI system measures its construct correctly, but whether the construct that it measures is the one that is relevant to the deployment context. Drawing on a theoretically guided synthesis of six bodies of literature, this review develops the Afrocentric Algorithmic Management Framework (AAMF), using African Ubuntu-grounded organizational contexts as its primary analytical site. Three pathways of misrecognition, including Relational Performance Misattribution, Communicative Capital Devaluation, and Historical Stratification Encoding, specify how AEI can give rise to Epistemic Algorithmic Bias (EAB). Three sequential design orientations to each of the pathways are given by the Afrocentric Design Architecture, and the institutional condition that allows their continued operation is provided by Communal Algorithmic Governance. It advances three propositions, three boundary conditions, and a seven question research agenda.

Keywords: Algorithmic Epistemological Incongruence, Epistemic Algorithmic Bias, Ubuntu Philosophy, Ai In Hrm, Algorithmic Fairness, African Management, Decolonial Ai

Introduction

Hiring, performance management, and career decisions in the world now use artificial intelligence tools. There are over 250 AI tools applicable to HR functions alone (Köchling & Wehner, 2020), and by 2024 about half of all organizations worldwide had already adopted AI into at least one HR function (Tambe et al., 2019; Turner, 2025). These solutions cover the entire employment lifecycle algorithmic resume screening, video interview analysis, predictive attrition modeling, and constant performance monitoring (Fabris et al., 2024). These tools have proliferated in most African nations via multinational vendor agreements, adoption by NGOs, and subsidiary requirements, frequently with minimal consideration of local circumstances (Ayana et al., 2024). The deployed tools were largely designed, trained and tested in the United States or Western Europe settings that dominate the global AI research and commercial development (Mohamed et al., 2020; Muldoon & Wu, 2023). They come in with internal assumptions, that is, what good performance entails, how competence can be manifested, and what are the signs of professional competence.

Those assumptions are instantiated within the systems evaluative architecture and not in their data alone. The implication of such assumptions built within the systems architecture is that in situations where performance is construed in different ways, where contribution is relational and not individual, where expertise is manifested orally as well as written, where career histories are shaped by historical constraints of access, and not by ability, such architectural assumptions can generate systematic misunderstanding of worker ability.

The algorithmic fairness literature is suitably equipped to identify bias, which functions via the use of demographic proxies or unrepresentative data (Mehrabi et al., 2021; Fabris et al., 2024). It has generated significant instruments: bias taxonomies, fairness measures, and audit models. Nonetheless, it makes one configuration under-specified: when the knowledge framework of the workers being evaluated by the AI system does not align with the evaluative framework of the AI system. This gap is discussed in this paper.

Two premises, taken together, generate a gap in the existing literature. Premise A: The body of literature on algorithmic fairness presupposes the existence of a common evaluative context. This assumption has been incorporated in its best-known fairness measures; equalized odds (Hardt et al., 2016), demographic parity, and individual fairness (Dwork et al., 2011), all of which assume the existence of a ground truth definition of performance or capability, which the AI system and the deployment context agree on. Premise B: Cross-contextual AI-HRM implementation, implements evaluative frameworks in settings that might not be homogeneous with those premises regarding the content of competence and its expression. When both conditions are true at the same time, as the evidence indicates is the case, then a sub-class of algorithmic misrecognition will be possible, which is not catered to with standard fairness audits. Algorithmic Epistemological Incongruence (AEI) is introduced to name and characterize that class.

AEI is the central concept of the Afrocentric Algorithmic Management Framework (AAMF) developed in this review. The remaining factors: the three misrecognition pathways, the outcome concept (EAB), the Afrocentric Design Architecture (ADA) and the governance condition (CAG) all are based on AEI. The framework does not substitute the toolkit of the algorithmic fairness literature, but builds upon it.

2. Approach

The theoretical approach of interpretive synthesis (Snyder, 2019; Jaakkola, 2020; Gilson & Goldberg, 2015; Torraco, 2005; Webster & Watson, 2002) is applied in this paper. This is the right approach when a theoretical issue is spread into a number of independent literatures and the key contribution is to unify them, rather than to introduce new empirical evidence. Instead of searching all sources available on each

topic, the synthesis is selectively based on six bodies of literature to discover what each one adds to the main problem of analysis, and what the others fill in. The procedure is iterative: finding convergent insights across disciplines, where the convergent insights are not connected, and developing the AAMF as the organizing structure the cross-literature gap suggests.

3. Six Literatures and What Each Leaves Open

Each cluster below is read for the specific insight it contributes to the central problem and the specific gap it leaves that makes AEI necessary.

Algorithmic Fairness and HRM-AI

This body of literature has produced valuable empirical results and instruments of governance. The most systematic recent overview, the multidisciplinary survey of fairness and bias in algorithmic hiring by Fabris et al. (2024) captures the current integration of AI tools throughout the recruitment process and notes that most studies on fairness focus on bias in agreed-upon evaluative systems, and the question of cross-context validity remains unanswered. As demonstrated by Buolamwini and Gebru (2018), the error rates of dark-skinned people are higher when using facial recognition systems. Raghavan et al. (2020) record the reproduction of occupational stratification by hiring algorithms. Shin (2022) examines the reasons why algorithmic discrimination is hard to address using traditional legal remedies. Mehrabi et al. (2021) offer a twenty-three bias taxonomy. Bandara et al (2025) propose HR Algorithmic Bias Management Capability (HABMC) as a governance framework at the firm level. According to Kim et al., (2024), power imbalances, and damaging opacities are among the strategic HRM implications of AI adoption.

It is important to comprehend what this literature leaves open. Its foundational fairness criteria, including equalized odds (Hardt et al., 2016) and individual fairness (Dwork et al., 2011), all presuppose a fixed, shared definition of the outcome being predicted. Equalized odds involve having a ground truth outcome label; demographic parity involves an agreement on the definition of qualified; individual fairness involves a similarity measure that both sides agree on. These standards are set such that they identify distortions in an agreed evaluative framework. They are not made to doubt the validity of the framework itself to the context of deployment.

Epistemic Injustice and AI

This cluster provides the philosophical vocabulary for naming the harm more precisely. Fricker (2007) recognizes testimonial injustice which occurs when the credibility of a person should be unjustly undermined as a knower and hermeneutical injustice which occurs when the interpretative equipment required to comprehend an experience is structurally absent. Kumar (2025) extends them to AI systems, demonstrating how annotation practices of data can marginalize non-Western forms of knowledge in a systematic manner. Milano and Prunkl (2025) prove that algorithmic profiling generates hermeneutical injustice by robbing people of the resources to challenge the assessment of themselves. According to Turner (2025), the epistemic degradation issue: AI management tools disrupt the ability of workers to deliberate, establish values, and challenge decisions that he distinguishes as the information asymmetry constraint and the value deskilling constraint. Using Bowker and Star (1999) and Foucault (1977), Milella and Cabitza (2026) states that the ability to decide what is worthy of being considered as evidence is the most significant form of epistemic power.

This is a contextual gap. These models deal with injustices to individuals in Western liberal democratic organizations. They do not hypothesize collective epistemic harm to communities whose traditions of

knowledge are not architecturally represented in AI assessment systems. That cannot be given by African philosophical grounding that this cluster alone can offer.

Decolonial AI Theory

Mohamed, et al. (2020) demonstrates that AI creation recreates colonial forms of knowledge production. Muldoon and Wu (2023) show that the global AI supply chain is a colonial supply chain; it exploits majority-world data workers and keeps producing hegemonic Western knowledge. Birhane (2021) relates this to Ubuntu epistemology by asserting that AI systems created under the premise of individualism can cause epistemic damage in places where personhood and knowledge are relational. The MASAKHANE project is reported by Ghajiga et al., (2026) and empirically demonstrated that African-language AI cannot be constructed by adding to the Western training systems, it must be designed in different ways.

The difference here is scale. Decolonial AI theory is applied to global supply chains. It recognizes AI systems are epistemologically loaded with the assumptions of their place of origin, but never indicates which aspects of HRM-specific evaluation architecture are loaded, or what design modifications would unload them. It is the contribution of translation by the AAMF.

African management and Ubuntu philosophy

Ubuntu-grounded practices are reported in Mangaliso (2001) and Mangaliso et al. (2018) as effective management technologies, rather than cultural residuals that are pre-modern. The ethnographic research conducted by Tayali (2025) in three organizations in Zambia reveals that HR practices based on Ubuntu generate a win-win situation both for the employees and employers. According to Creighton-Randall (2024), the five Ubuntu organizational values are survival, solidarity spirit, compassion, respect, and dignity that have quantifiable community-building impacts. Grobler and Powell (2025) reveal that Ubuntu leadership is, at the same time, transformational, inclusive, and ethical. Odion and Agbonwanegbe (2025) posit that communalism of Ubuntu is a different model of governance. The traditions of African management have been disrupted by the colonialists, as recorded by Inyang (2009), indicating that the imposition of Western management systems on African countries is a continuous historical process, which began being disrupted by the colonialists.

There are two critical clarifications that forestall essentialist misreading. First, Ubuntu is not an African cultural property homogenous across Africa, but an institutionally variable feature of particular organizations, which is always evaluated at an organizational level. Second, structurally parallel relational traditions of governance can be found in West African and North African settings, which increase the range of applicability of the framework. The temporal gap created by this literature is that it is concerned with pre-digital HRM. The AAMF applies its insights to algorithmic governance.

Postcolonial Management Theory

Frenkel (2008) and Jack and Westwood (2009) establish that management theory reproduces colonial knowledge hierarchies. Jackson (2025) contributes to the discipline by stating that the dynamics of capital structure has been ignored in postcolonial theory: the cultural turn has decoupled culture and social structure, and it is therefore impossible to analyze the reproduction of the universalizing logic of capitalism by the relationships between AI-HRM vendors and colonies. Mayasandra et al. (2006) demonstrate the dynamics of othering that Western knowledge management structures created in IT organizations. Odion and Agbonwanegbe (2025) amalgamate feminist and postcolonial approaches to show that Ubuntu communalism is a different paradigm of governance that is being overshadowed by imported Western management.

The recognized gap (Frenkel, 2008; Jackson, 2025) is the lack of actionable frameworks that are construct-precise. The AAMF is created to bridge that gap.

Algorithmic Management and Workplace Democracy

Turner (2025) gives the most lucid explanation of the means of undermining epistemic agency by AI management tools with information asymmetry and value deskilling barriers. The two constraints are both intensified in Ubuntu-grounded contexts, where individual democratic voice is not the norm but collective deliberation. When AMTs undermine collective deliberation, it is not only that individual workers suffer but also that the community itself suffers in its ability to govern itself.

Miceli et al. (2025) propose Workers Inquiry as Research Methodology (WIRM) - the focus of worker epistemic authority by participatory action approaches. Their operation in Global South labor environments informs the governance design at CAG. The distance is the liberal democratic working assumption: the assumption in the Turner model is that individuals will act democratically. The communal model of epistemic governance adopted by Ubuntu needs an alternative model of accountability.

4. How AEI Relates to Existing Bias Constructs

The critical reader will inquire on whether AEI is actually different to existing bias constructs or whether it is merely a re-description of something already present in the literature. This section addresses that question directly.

AEI builds upon the construct validity theory (Messick, 1995; Cronbach & Meehl, 1955) and puts the unit of analysis as measurement adequacy into evaluative ontology alignment. Construct validity theory asks: does the measure capture the construct it claims to measure? AEI asks one level deeper: is the construct being measured the relevant one for the deployment context? It is not the rejection of construct validity but an upstream extension. AEI inquires about the applicability of the construct before inquiring about the sufficiency of the measure. It is that extra step of analysis that makes AEI stand apart, not reducible to, extant bias taxonomies. Table 1 brings this to reality by contrasting AEI with the four most similar existing constructs.

TABLE 1: How AEI Relates to — and Extends Beyond — Existing Bias Constructs

Bias Type	Where the Error Lives	How to Fix it	Why AEI Extends Beyond it
Measurement Bias (Mehrabi et al., 2021)	Wrong proxy used within an agreed evaluative model. Both parties agree on what is being measured; the proxy just fails to capture it accurately	Replace the proxy with a better measure of the same construct.	Both parties share the same evaluative framework, the same definition of what competence means. AEI begins one level deeper: the framework itself may not be valid for the context, not just the proxy within it.
Sampling / Representation Bias (Mehrabi et al., 2021)	Training data under-represents certain demographic groups, producing	Add more representative training data from underrepresented groups.	Adding more African workers to training data does not resolve AEI if the evaluative criteria still do not recognize how those workers express competence.

	skewed statistical distributions.		AEI is in the architecture, not the data composition.
Historical / Label Bias (Raghavan et al., 2020; Bertrand & Mullainathan, 2004)	Outcome labels in historical data reflect past structural inequalities, who were hired or promoted because of access, not ability.	Reweight or adjust outcome labels; apply counterfactual fairness procedures.	HSE is a pathway through which AEI produces harm, but AEI is broader. AEI can produce EAB through RPM and CCD even in a system with clean, balanced historical labels. The pathways share the same root (AEI) but operate through different mechanisms.
Construct Validity Failure (Messick, 1995; Cronbach & Meehl, 1955)	The measure does not capture the construct it claims to measure, for example, a test labeled 'leadership ability' that only captures verbal fluency.	Redesign the measure to better operationalize the agreed construct.	Construct validity failure assumes a shared understanding of the construct and asks whether the measure captures it. AEI extends this by shifting the unit of analysis: it asks whether the construct being measured is the relevant one in the deployment context. If 'performance' is defined as individual output and the worker's performance is relationally constituted, the problem is not measurement inadequacy, it is evaluative ontology misalignment
Algorithmic Epistemological Incongruence (AEI)	The evaluative ontology of the AI system does not correspond to the knowledge framework through which workers in the deployment context constitute competence and performance.	Redesign the evaluative architecture, what is being evaluated and how competence is defined, not just the proxies or data.	AEI extends construct validity theory by shifting the unit of analysis from measurement adequacy to evaluative ontology alignment. It is not a replacement for existing bias constructs, it addresses the architectural layer that sits upstream of all of them.

The practical implication is significant. Since the source of AEI is architectural, being situated in what the system judges, rather than how it judges, the interventions that are necessary are also architectural. RPM is

not fixed by recalibration of a proxy variable. Increasing the training data of the African candidates does not solve CCD. HSE needs to reweight historical outcome labels but does not need to do so with the other pathways. The three analytically distinct design orientations are needed as each pathway needs a different kind of fix.

5. Algorithmic Epistemological Incongruence

Definition

AEI is characterized as the extent to which the epistemological presuppositions in the evaluation architecture of an AI-HRM system are not aligned with the knowledge structures in which workers targeted constitute, articulate and practice competence, performance and professional identity.

The three boundaries explain the non-AEI. First, AEI is an architectural, not a data property. It is in the design logic of the system, the model of what performance is, a definition of what valid evidence looks like, not the demographic composition of training data. Second, AEI is not synonymous with cultural distance (Hofstede, 1980): cultural distance is a measure of the difference in values among people; AEI is a measure of the epistemological validity of a particular evaluative architecture in a particular organizational environment. Third, AEI is not partisan. It is the state of architecture in which a particular type of bias can occur.

AEI's Three Levels

At three levels, AEI can act, and the determination of the level where it is active is important to understand what type of intervention is effective.

On the most fundamental level — ontological — AEI deals with the model of the worker. Western AI-HRM presupposes the presence of an individual autonomous agent whose performance can be disaggregated as compared to others. Ubuntu epistemology makes the worker relationally constituted: competence and performance cannot be usefully dissociated with relationships (Metz, 2011; Tayali, 2025). This is no cultural choice, but an ontological one.

At the middle level — epistemological — AEI concerns what counts as valid evidence of competence. Western artificial intelligence systems view educational qualification and proficiency in written English, as well as the output measures, as legitimate cues. Epistemological traditions in Africa can consider expertise that has been validated by community, multilingualism, and oral knowledge as equally valid (Santos, 2014; Asante, 1987).

At the surface level — evaluative — AEI concerns the specific proxies and metrics used. It is possible that interventions at this level alone, modifying metrics without altering the underlying ontological structure can mitigate AEI, but not solve it. That is why the three orientations of ADA are logically sequenced and embrace one level at a time.

Operationalizing AEI: Three Measurable Dimensions

Above AEI is conceptually defined but to be of practical use in empirical research and in organizational practice, it must be measurable. In this subsection, three dimensions are presented that are '*operable*' and each of which relates to the three levels of AEI. Table 2 shows these dimensions and gives an example of indicators and recommended instruments.

TABLE 2: Operationalizing AEI; Three Dimensions, Example Indicators, and Suggested Instruments

AEI Dimension	What it Measures	Example Indicator	Suggested Instrument
Dimension 1: Ontological Misalignment (OM)	The degree to which the AI system assumes an individual-attribution model of performance in a context where performance is relationally constituted	'The AI system generates individual performance scores that cannot be attributed to team or network contributions' agree/disagree scale administered to HR managers and workers	5-point Likert organizational audit survey administered to HR managers, line managers, and employee representatives; validated against UCE scale (Mangaliso, 2001)
Dimension 2: Epistemological Misalignment (EM)	The degree to which the AI system's valid evidence criteria do not recognize forms of competence expression that are legitimate in the deployment context	'The AI system assigns lower scores to candidates whose professional expertise is primarily community-validated rather than credential-certified' assessed through proxy variable audit against local competence norms	Expert coding protocol: two raters assess each AI evaluation feature against locally elicited competence norms; inter-rater agreement; supplemented by algorithmic audit (Sandvig et al., 2014)
Dimension 3: Evaluative Proxy Misalignment (PM)	The degree to which the specific metrics and proxy variables used by the AI system diverge from the indicators that the deployment context treats as meaningful measures of professional contribution	'The performance indicators used by the AI system (e.g., output counts, response time) do not reflect the dimensions of contribution most valued by our organization' scored by HR managers against community-generated contribution dimensions	AI system feature audit: structured analysis of system documentation and vendor specifications coding each feature's evaluative assumption against context-validated contribution indicators; reported as misalignment score per feature

The dimensions are meant to be measured separately. An organization may be characterized by high Ontological Misalignment (the AI system solely relies on individual attribution) but lower Epistemological Misalignment (the AI system does accept community-approved credentials in its scoring). This autonomy implies that the dimensions can be utilized in the context of differential diagnosis, determining which of the levels of AEI is the most serious in a particular deployment situation and focusing interventions on it.

AEI is best understood as a composite score: the sum or weighted average of OM, EM, and PM scores for a given AI-HRM system in a given deployment context. The large composite AEI score will mean that the system has validity issues at various levels. A more narrow-specific intervention need is represented by a specific AEI score, high on one dimension only.

Scope: Africa-Grounded, Not Africa-Limited.

The AAMF uses African Ubuntu-grounded contexts as its primary analytical site. This is conceptually driven: the epistemological gap between the relational ontology of Ubuntu and the individualist architecture of Western AI-HRM is so great as to render AEI an obvious phenomenon, and the philosophical expertise to theorize it most advanced within African management literature.

The same condition, however, may arise in other contexts where AI systems cross significant epistemological boundaries. The misalignments are similar with indigenous community organizations where knowledge is land-based and communal; when AI performance systems use individual output metrics. East Asian collectivist organizations, in which group harmony and face-preservation inform the expression of performance, can be subjected to CCD-type misrecognition by NLP tools that are trained on individualist communication norms in the West. Innovative and artisan industries, where tacit embodied expertise is not easily tracked in terms of quantifiable output, can be subject to RPM-type misrecognition even in Western contexts. AEI is the overall state; it is in the African Ubuntu setting that the AAMF develops it due to the fact that the resources to do so are richest there.

6. From AEI to EAB: Three Misrecognition Pathways

The Outcome: Epistemic Algorithmic Bias (EAB)

EAB is characterized as systemic under-recognition of competence, performance and potential of workers due to the evaluative frameworks mismatch in AI-HRM implementation.

Systematic refers to the consistency among workers whose professional qualities are manifested by the underrepresented knowledge framework; it is not a chance error. Misrecognition builds on the recognition theory of Honneth (1995): structural failure to recognize contributions that exist, are valued and have consequences within the community in which the worker lives.

EAB appears in three forms which are empirical signatures: competence misrecognition (scores lower than expert validators determine them), performance invisibility (valid performance dimensions not reflected in the evaluative architecture), and trajectory distortion (development and promotion allocated unreasonably due to epistemologically biased evaluations). These are diagnostic categories, not independent constructs; they indicate the direction of the pathway that is functioning most strongly.

Pathway 1 - Relational Performance Misattribution (RPM)

RPM explains how AEI on the ontological level can create performance invisibility-form EAB. The channel works based on the individual-attribution design hypothesis incorporated in AI performance management systems.

The process consists of three steps. First, the AI system monitors distinctively visible behavior, such as keyboard usage, task accomplishment, the number of output, email rate, and so forth, since they are

measurable by individual-attribution models. Second, in relationally constituted performance based organizations, the most appreciated contributions of mentorship, maintenance of social harmony, transfer of knowledge, and building of trust do not produce any trace in these systems that would be individually attributable. Third, workers whose excellence is primarily relational receive lower algorithmic scores than their community-validated performance justifies.

This avenue has been based on a number of intersecting bodies of evidence, not just in the field of African management. The theory of team production (Alchian & Demsetz, 1975) has determined that individual production cannot be easily separated in a manner that does not distort production when output is produced jointly, a structural issue rather than measurement error. The development of social capital research (Nahapiet & Ghoshal, 1998) records the means of establishing organizational value by networks of relations which are real but cannot be ascribed to any individual. The results of the study of organizational citizenship behavior (OCB) (Organ, 1988; Bolino et al., 2002) indicate the validity of the performance dimensions, such as assisting colleagues, sustaining harmony in the group and aiding institutional functioning, which are shown to be missed in the standard metrics of individual performance in a systematic way. The empirical evidence of Ubuntu specific to Tayali (2025) offers the African organizational grounding. These anchors collectively determine that RPM is not a hypothetical possibility; it is a foregone conclusion of the application of individual-attribution logic to any relational performance environment.

Practical application: an RPA-conformist system would be based on graph-based contribution scoring, with the edge weights between workers determined by peer validation of facilitative contribution, and the individual scores determined as functions of the network centrality of a worker in the contribution graph, not as aggregates of the individually attributed outputs.

Pathway 2 - Communicative Capital Devaluation (CCD)

CCD explains how AEI at the epistemological level can result in competence misrecognition-form EAB. The pathway works by the linguistic epistemology of NLP-based AI evaluation tools.

The evaluation tools of NLP are mainly trained on standard American and British English corpora. This inscribes a certain linguistic epistemology: the competence and professionalism is given as a validly expressed mode of expression using particular phonological patterns, syntactic structures, and rhetorical conventions. Other candidates whose communicative repertoires do not fit this training distribution, such as African '*Englishes*', multilingual code-switching, and oral rhetorical forms, are placed as outliers. Their communicative patterns are assigned lower probability scores in the system and they generate competence penalties, with no relationship to real cognitive or professional capability.

CCD is not exactly like the regular NLP linguistic bias. Standard NLP bias has to do with the erroneous classification of language, such as a sentiment model that mistakes the African American Vernacular English to be negative (Blodgett et al., 2020) or a speech recognition model with increased rates of errors in Black speakers (Koenecke et al., 2020). These are bugs in a known system model of what the system is attempting to accomplish. CCD is another type of issue: it is the failure of the system to acknowledge other forms of communicative epistemology as legitimate forms of competence in the first place. It is not that African English is being misread by the system; it is that there is no model of African English as valid professional communication in the system.

The existence of the MASAKHANE project directly empowers this distinction (Ghajiga et al., 2025). MASAKHANE was developed as a result of the impossibility of training African-language AI on more

African data specifically because the epistemology of the training structure is based on the Western linguistic epistemology, and not only the quantity of data.

Practical implementation example: an NLP system that is EKC-compliant would be multilingual embedding models that are trained exclusively on African English corpora and code-switching data, following the MASAKHANE-type methodology, that way the system could recognize African communicative repertoires as a legitimate professional expression.

Pathway 3: Historical Stratification Encoding (HSE)

HSE explains that the evaluative level of AEI, namely, in the treatment of historical outcome data, can generate trajectory distortion-form EAB.

AI-HRM systems that have learned through past African labor market results data acquire patterns based on the distributions of historical colonial access control to occupations. The system is trained to forecast success based on features that are correlated with historical access: institutional credentials, employer brand membership, regularities in occupational trajectories. These attributes mirror the access under colonial arrangements, rather than ability. According to Muldoon and Wu (2023), it is the Quijano (2000) colonial matrix of power at the level of data.

HSE as a standalone observation is not new; the persistence of historical inequality in algorithmic systems is well-documented in labor economics (Nunn, 2008) and hiring discrimination research (Bertrand & Mullainathan, 2004). The AAMF contribution is the incorporation of HSE into the AEI. HSE is no longer a distinct type of bias; it is the mechanism by which historical inequality combines with epistemological misalignment to cause the distortion of the trajectory. The interpretation of HSE in AEI explains why a standard historical bias correction (demographic rebalancing of training data) is inadequate: the contamination is in the labels of outcome, which capture patterns of colonial access, and not in the distributions of features.

Concrete implementation example: An ATC-compliant system would use counterfactual fairness reweighting (Kusner et al., 2017) on historical outcome labels, rescaling label weights to reflect what would have been the outcomes in the absence of colonial access restriction, and then use the labels as supervisory signals during model training.

Pathway Interactions

These are the three paths that are analytically different, but not mutually exclusive. In reality, they will tend to co-exist.

A high-UCOE organization worker who also is multilingual and whose training data context is high COSD can also face performance invisibility (RPM), communicative competence penalties (CCD), and a misallocation of career trajectory (HSE). These influences will accumulate rather than merely add: a relationally invisible worker with communicative penalties is more deeply misrecognized than either of these two directions would generate alone.

The three ADA orientations which cover these pathways can also cause implementation tensions. RPA needs community validation procedures, which need time and institutional confidence to develop. EKC needs multilingual corpus development that does not run on an organizational procurement cycle. ATC needs access to historical data which might be unwelcome by vendors. The AAMF is thus provided as a framework of prioritized implementation that is contextually sensitive, rather than a blueprint to be taken in all its essentials at once.

7. Propositions

Three propositions are put forward and can be tested. Each of them is based on a certain aspect of the AAMF and connected to an empirically feasible research design.

Proposition 1 (P1) - AEI Separability:

Systems of AI-HRM with stronger Algorithms Epistemological Incongruence (AEI) (based on composite OM + EM + PM scores) will explain the variance in the gap between outcomes of evaluation through algorithmic assessment and community-validated performance assessment over and above the variance in existing statistical measures of bias such as imbalanced demographic representation and proxy discrimination. This hypothesis directly probes the question of whether the three-dimensional construct of AEI contributes explanatory value over and above that of the bias taxonomy of Mehrabi et al. (2021). It is regulated by Algorithmic Opacity Level (AOL): when the level of opaque is high, the effects of AEI add up without being noticed and its variance added to the explained variance will be bigger.

Proposition 2 (P2) - UCE Moderation:

Ubuntu Cultural Embeddedness (UCE) will positively mediate the relationship between AEI and Epistemic Algorithmic Bias (EAB) - the difference between algorithmic performance scores and community-validated performance assessments; in that higher UCE organizations will have a larger evaluation gap compared to lower UCE organizations using the same AI-HRM system. This hypothesis is a test of whether the institutional context at the theoretically anticipated direction increases the role of the Ontological Misalignment dimension of AEI.

Proposition 3 (P3) - EKC Independence:

Epistemic Knowledge Calibration (EKC) interventions—specifically, African-corpus NLP model training will reduce Communicative Capital Devaluation (CCD) effects on candidate evaluation scores independent of demographic parity adjustments. This proposition addresses whether CCD is a separate track to overcoming representational bias by comparing the score penalties arising from both CCD and demographic balancing.

The three propositions are chosen as they are the most consequential propositions of the framework and can be realized in practical studies such as algorithmic audit study, matched-pair experiment and organizational field study.

8. Afrocentric Algorithmic Management Framework (AAMF)

The AAMF has five components in a rational order: Problem (AEI and its three dimensions) → Pathways (RPM, CCD, HSE) → Outcome (EAB) → Design Response (ADA) → Governance Condition (CAG). Each element is summarized in table 3.

TABLE 3: AAMF — Five Elements, Roles, Grounding, and Contributions

Element	Role	Grounding	What it Adds
Algorithmic Epistemological Incongruence (AEI)	Central concept	Asante (1987); Metz (2011); Bowker & Star (1999); Messick (1995)	Extends construct validity theory by shifting the unit of analysis from measurement adequacy to evaluative ontology alignment. AEI asks not whether the system measures its construct correctly,

			<p>but whether the construct it measures is the relevant one for the deployment context. Operationalized through three dimensions: Ontological Misalignment (OM), Epistemological Misalignment (EM), and Evaluative Proxy Misalignment (PM).</p>
Three Misrecognition Pathways (RPM · CCD · HSE)	Explanatory layer	Tayali (2025); Nahapiet & Ghoshal (1998); Organ (1988); Ghajiga et al. (2025); Muldoon & Wu (2023); Nunn (2008)	<p>Three analytically distinct pathways through which AEI may produce EAB at the attribution, evidential, and historical-temporal stages respectively. Each pathway is grounded in a different body of literature, making the pathways independently anchorable and not mutually dependent. The pathways can co-occur and interact, producing compound misrecognition effects.</p>
Epistemic Algorithmic Bias (EAB)	Outcome concept	Fricker (2007); Honneth (1995); Turner (2025); Mehrabi et al. (2021)	<p>The observable consequence of AEI: systematic misrecognition of worker competence, performance, and potential arising from evaluative framework mismatch. Three manifestation forms, competence misrecognition, performance invisibility, and trajectory distortion, serve as empirical signatures for identifying which pathway is operating most strongly in a given deployment context.</p>
Afrocentric Design Architecture (ADA)	Design response	Santos (2014); Ghajiga et al. (2025); Kusner et al. (2017)	<p>Three ordered design orientations addressing AEI at each of its three levels. RPA addresses ontological misalignment through graph-based relational performance attribution. EKC addresses epistemological misalignment through multilingual</p>

corpus training and proxy recalibration. ATC addresses evaluative proxy misalignment through counterfactual fairness reweighting of historical outcome labels. Orientations are logically ordered, not interchangeable.

Communal Algorithmic Governance (CAG)	Governance condition	Metz (2011); Sloane et al. (2022); Lee et al. (2019); Birhane et al. (2022); Ayana et al. (2024)	The institutional architecture enables ADA to be implemented and sustained against vendor re-standardization pressures. Grounded in both Ubuntu relational justice and the participatory AI governance literature. In practice, CAG involves a worker governance council participating in model audit review, validation of performance metric definitions, and dispute resolution providing the binding accountability mechanism that keeps ADA orientations in place over time.
--	-------------------------	--	---

9. The Design Response: ADA and CAG

Afrocentric Design Architecture (ADA)

The three ordered design orientations contained in ADA target AEI on one of its three levels and counter one of the three pathways. These orientations logically follow each other: RPA deals with ontological misalignment, EKC deals with epistemological misalignment, and ATC deals with evaluative proxy misalignment at the historical-temporal level. The introduction of EKC in absence of RPA lowers penalties of communicative competence but does not eliminate performance invisibility. The application of both in the absence of ATC does not resolve the issue of trajectory distortion by use of HSE. The orientations are complementary to each other and when they are implemented together they minimize AEI in totality.

Relational Performance Architecture (RPA) redefines the performance analysis unit between the individual and the relational network. It is based on the relational constitution of performance in Ubuntu (Tayali, 2025; Mangaliso et al., 2018; Grobler & Powell, 2025) and the documentation of relational value creation done in the social capital literature (Nahapiet & Ghoshal, 1998). RPA needs performance attribution on network-based performance that needs to be peer-validated, longitudinal cohort tracking of performance that needs to attribute facilitative performance to facilitators, and community-validated contribution measurement as primary data. RPA is not team performance measurement; RPA does not modify the way the system evaluates performance, but rather the definition of what performance is.

Epistemic Knowledge Calibration (EKC) restructures the knowledge validity hierarchy in AI evaluation criteria. Grounded in Santos' (2014) epistemologies of the South and the MASAKHANE methodology (Ghajiga et al., 2025), EKC orients system design toward treating African communicative forms multilingual repertoires, oral knowledge traditions, indigenous credentialing as epistemologically valid evidence. It includes proxy variable retraining, African-corpus NLP model training and knowledge source diversification. EKC is a validity architecture intervention, it does not modify the presentation of evidence at the interface, and only what is considered legitimate.

The idea of Algorithmic Temporal Contextualization (ATC) is that training data is not a neutral record, but evidence of the past, which must be contextualized. Based on the Akan Sankofa ethos, which states that it is not bad to revisit and reclaim what you forgot, and research on historical inequality (Nunn, 2008), ATC orientates the training data governance to follow the colonial lineage of outcome names and implement the de-stratification processes. ATC's obligation is proportional to COSD: most urgent in South Africa, Zimbabwe, and Namibia; calibrated differently in West Africa; least urgent in non-colonized contexts.

Communal Algorithmic Governance (CAG)

CAG is the governance condition that enables ADA to be implemented and sustained. In its absence, ADA orientations are constantly under re-standardization pressure by global AI vendors and in MNC subsidiary set-ups coercive parent-company mandate (DiMaggio & Powell, 1983). CAG is not a design dimension- it is the institution architecture that maintains the other dimensions in place across time.

CAG brings the ethics of Ubuntu, as defined by Metz (2011), to the accountability of individual workers, and moves it to that of the collective. It is based on the literature of participatory AI governance, which has elaborated viable frameworks of community engagement in AI system management. WeBuildAI is an empirical example of participatory AI governance as presented in Lee et al. (2019). Sloane et al. (2022) record the circumstances in which participatory AI practices create meaningful, as opposed to superficial, community impact on system design. Birhane et al. (2022) show how participatory methods can be implemented in regard to ML fairness in particular.

Practically, CAG entails a worker governance council that has the binding mandate on three processes. First, model audit review: the council can review the output of algorithmic audits, may demand explanation of flagged decisions, and may require the council to recalibrate where EAB evidence is discovered. Second, validation of performance metrics: prior to any performance evaluation criteria being put into practice in the AI system, they are checked and approved by the council on contribution standards that are validated by the community. Third, dispute resolution: employees that believe their algorithmic rating is a poor reflection of their input have the option to take their case to the council, which may override algorithmic scores on the basis of community-verified evidence. The theory of governance based on the Arnstein (1969) ladder of citizen participation is the benchmark: the real CAG is at the levels of partnership and delegated authority, not at the lower levels of consultation and provision of information.

10. Scope and Boundary Conditions

Core Application Conditions

Three conditions are necessary for AEI to be analytically applicable. First, AI-HRM systems must be active in decisions affecting workers. Second, the deployment context must exhibit meaningful epistemological distinctiveness from the source context regarding how competence is constituted, how performance is expressed, and what counts as valid evidence. Third, this distinctiveness must be institutionally significant, embedded in actual organizational practice.

Three Boundary Conditions

Within the applicable scope, three boundary conditions specify where AEI risk is highest and where the framework's design orientations are most urgent. Table 4 presents each with moderation direction.

TABLE 4: Three Boundary Conditions — Moderation Logic and Scope Specification

Condition	AEI risk is higher when....	AEI risk is lower when....
Ubuntu Cultural Embeddedness (UCE) Links to: <i>P2</i>	Relational performance norms are institutionally embedded as the primary standard for how contribution is assessed, decisions are made, and relationships are governed. In high-UCE contexts, the evaluative logic of an individual-attribution AI system diverges most sharply from how workers and managers understand good work. This is where performance invisibility (RPM) is most likely to occur.	The organization operates with Western-style individual performance norms, for example, a professional services firm with Western management culture. AEI is still analytically applicable but its severity is lower, and the urgency of ADA's relational orientation is reduced accordingly.
Colonial Occupational Stratification Depth (COSD) Links to: <i>P3</i>	Training data comes from sectors or organizations where colonial-era access restrictions shaped who got which jobs. In these contexts, historical outcome labels record who had access, not who had ability. AI systems trained on these labels reproduce colonial occupational stratification as a success predictor (Nunn, 2008; Bertrand & Mullainathan, 2004).	Training data comes from post-independence organizations where occupational restructuring has substantially changed the outcome distribution, or from systems that use direct capability testing rather than historical outcome data as supervisory signals. ATC remains relevant but less urgent.
Algorithmic Opacity Level (AOL) Links to: <i>P1, CAG</i> feasibility	The AI system does not explain its reasoning, the modal case in many African deployment contexts (Ayana et al., 2024). Errors accumulate across model	XAI-compliant design makes decision logic auditable. CAG's community governance function becomes technically feasible. Transparency is a necessary precondition for the

generations because nothing flags them for correction. High opacity also disables CAG's audit function: community governance mechanisms cannot operate where the system's decision logic is inaccessible.

framework's governance dimension to function without it; even a well-resourced CAG cannot exercise meaningful oversight.

Within-Africa Heterogeneity

The framework's use of Ubuntu philosophy carries an essentialism risk that must be managed explicitly. The core position: UCE is an organizational-level variable, always assessed empirically, never assumed from national or ethnic identity. Table 5 specifies four dimensions of within-Africa variation that any application of the AAMF must account for.

TABLE 5: Within-Africa Variation — Four Dimensions and Framework Implications

Dimension	What varies	What this means for the framework
Ubuntu Embeddedness (UCE)	Ubuntu is most institutionalized in Nguni and Bantu Southern and East African contexts. West Africa has analogous relational governance traditions, Yoruba <i>ẹgbẹ</i> networks; Igbo <i>ohanaeze</i> structures, that are relational but philosophically distinct. North African Islamic traditions share structural relational features but differ in content.	UCE is measured at the organizational level, not assumed from national identity. A rural community cooperative in Zambia and a multinational bank's office in the same country may exhibit very different UCE levels. The framework begins with an organizational assessment every time.
Colonial History Depth (COSD)	COSD varies significantly across Africa. South Africa and Namibia have the deepest recent colonial occupational stratification histories. West African nations have different patterns. Ethiopia was not formally colonized and has a distinct COSD profile.	ATC's de-stratification obligation is proportional to COSD. It is most urgent where colonial occupational restriction was most recent and most systematic. It is still relevant elsewhere, but requires calibration to the specific historical labor market context of the training data source, not a generic 'African colonial history'.
Urban–Rural and Sector Context	Urban professional service firms tend toward Western HR norms (low UCE). Rural cooperatives and	Framework application is sector- and organization-specific. The same AI-HRM system may exhibit high

community organizations tend toward higher UCE. Sector shapes context substantially: extractive industry, NGO, public sector, and fintech each create different institutional environments.

AEI in one organization and lower AEI in another within the same city. UCE assessment always precedes application.

Gender
Intersectionality

Ubuntu relational norms are not gender-neutral. Oyewumi (1997) shows gender operates differently in African ontologies than in Western feminist frameworks. Amadiume (1987) documents matrilineal structures that complicate simple communalism assumptions.

EAB does not operate identically for all African workers. Gender intersects with epistemological harm in ways that need specific investigation. The AAMF provides the conceptual architecture for this but does not resolve it. Research Question 6 identifies it as a priority.

11. Research Agenda

Seven research questions are advanced, sequenced from foundational construct validation through mechanism testing to scope extension. Table 6 presents each with methodological guidance.

TABLE 6: Research Agenda — Seven Questions with Methodological Guidance

Ref.	Research Question	Method & Grounding
RQ1	Can AEI's three dimensions (OM, EM, PM) be reliably measured through an organizational audit instrument, and does aggregate AEI score predict evaluation outcome gaps beyond what statistical bias measures already explain?	Scale development using expert panel + cognitive interviewing; structural equation modeling comparing AEI dimensions and statistical bias as predictors of scoring gaps; builds on Bandara et al. (2025) HABMC framework
RQ2	Do NLP-based AI-HRM tools assign systematically lower scores to candidates communicating through African ' <i>Englishes</i> ' and multilingual code-switching, independent of assessor-validated cognitive and professional capability?	Matched-pair audit experiment: communicative repertoire as manipulated variable; NLP scoring vs. expert panel; Blodgett et al. (2020) and Koenecke et al. (2020) as empirical comparators; Shin (2022) as legal framework
RQ3	Does Colonial Occupational Stratification Depth (COSD) in training data predict trajectory distortion outcomes, hiring gaps, promotion rate disparities in AI-HRM	Longitudinal organizational study (3–5 years); historical labor market genealogy analysis of training outcome labels; Nunn (2008) on

	deployments in Sub-Saharan African organizations over a multi-year period?	historical inequality persistence; Muldoon & Wu (2023) as structural framework
RQ4	Does Communal Algorithmic Governance (CAG) specifically worker governance councils with binding audit authority maintain the effectiveness of ADA orientations over time against institutional re-standardization toward Western vendor defaults?	Comparative institutional case studies: high-IIP MNC subsidiaries vs. low-IIP domestic firms; DiMaggio & Powell (1983); Sloane et al. (2022) participatory AI governance as comparator
RQ5	How does UCE vary across sector, ownership type, and urban-rural context within a single African national context, and does this organizational-level variation predict differential AEI levels for the same AI-HRM system?	Multi-site cross-sectional study within one country (following Tayali's 2025 Zambian methodology); UCE measured at organizational level; AEI assessed using audit instrument from RQ1
RQ6	How do gender and intersectional dynamics (Oyewumi, 1997; Odion & Agbonwanegbe, 2025) shape differentiated EAB patterns for African women workers, and does the AAMF require gendered adaptation?	African feminist participatory methodology aligned with WIRM (Miceli et al., 2025); African women workers as community researchers; Crenshaw (1989) intersectionality as analytical frame
RQ7	Do AI-HRM tools developed by indigenous African AI organizations exhibit measurably lower AEI across all three dimensions than functionally equivalent Western-vendor tools in high-UCE deployment contexts?	Comparative algorithmic audit: matched-function tools compared using the AEI audit instrument; Ghajiga et al. (2025) MASAKHANE as epistemic benchmark; Birhane et al. (2022) as participatory governance comparator

12. Contributions

The AAMF makes three contributions, each targeted at a specific gap.

The main contribution is to the algorithmic fairness literature. AEI extends that literature's toolkit by adding an epistemological validity dimension that the statistical discrimination paradigm does not address. More specifically, AEI builds on the construct validity theory (Messick, 1995) by shifting the unit of analysis of the measurement adequacy to an evaluation ontology alignment, an upstream analytical layer, sitting above construct validity questions are accessed. This has a direct practical implication: The HABMC governance framework, the most comprehensive existing tool to manage AI-HRM bias, needs an AEI assessment dimension to address epistemological misalignment in addition to statistical bias. Kim et al. (2024) strategic HRM research agenda should add epistemological validity as a system quality criterion alongside accuracy and fairness metric compliance.

The secondary contribution is to postcolonial management theory. The AAMF has responded to the acknowledged gap by Frenkel (2008) - critiquing without constructing, and to the call of Jackson (2025) to provide concrete structural analysis, by offering operationalisable design orientations and a measurable central construct.

The third contribution is to the Ubuntu-in-management tradition. Tayali (2025), Mangaliso et al. (2018), Grobler and Powell (2025), and Creighton-Randall (2024) have established an empirical record of HRM based on Ubuntu, in pre-digital contexts. The AAMF moves that record into the domain of algorithmic governance.

13. Practical and Policy Implications.

For AI-HRM system designers and vendors, the AEI operationalization framework, three dimensions assessed through organizational audit, expert coding, and AI system feature analysis provides a pre-deployment assessment tool. Vendors must evaluate: OM level (estimating RPM risk), EM level (estimating CCD risk through corpus analysis), and PM level before deploying in a new African context, assess: OM level (estimating RPM risk), EM level (estimating CCD risk through corpus analysis), and PM level. These three tests bring the HABMC framework by Bandara et al. (2025) to the scope of the epistemological validity domain.

For African organizations, the AEI construct shifts the evaluation question from 'does this system produce disparate impact?' to 'is this system's evaluative architecture valid for how our workers constitute competence?' This is consistent with the AU Digital Transformation Strategy (2020–2030) and the AU Data Policy Framework (2022).

For policymakers, individual-rights governance frameworks, GDPR's right to explanation, Title VII's disparate impact doctrine are insufficient for EAB, which is a collective epistemological harm. An Ubuntu-relational justice-based framework of algorithmic labor governance based on community-level algorithmic impact assessment and worker governance councils with binding audit authority should become standard features.

14. Conclusion

This paper started from a narrow but consequential observation. The literature on algorithmic fairness is founded on the premise that bias plays off within a common evaluative context. This is a logical assumption that has produced valuable tools. But it leaves one deployment configuration analytically under-specified: when AI-HRM systems are applied across contexts that do not share those assumptions.

That arrangement is called AEI. It builds upon construct validity theory, shifting the unit of analysis from the measurement adequacy norm to the evaluative ontology alignment norm, asking not whether the AI system is measuring its construct in a manner consistent with adequate measurement, but whether the construct it is measuring is the right one in the deployment situation. The AAMF operationalizes AEI in three quantifiable dimensions, determines three pathways through which it can produce EAB, suggests three design orientations to address each pathway, and defines a governance condition of sustainable implementation.

Four things make the AAMF different from a normative critique. First, the argument of discriminant validity demonstrates that AEI extends but not replaces the existing bias constructs, it works upstream of the analytical layer that the existing tools are dealing with. Second, AEI can be operationalizable in three measurable dimensions and two proposed types of instruments; it is not just a philosophical concept. Third, three testable propositions make the framework's claims falsifiable. Fourth, the real micro-examples in Section 6 and 9 relate analytical orientations to particular design processes.

The fact that the framework is based on the philosophy of Ubuntu and the African management scholarship is not incidental. These traditions are most advanced in the philosophical and empirical resources to theorize epistemological validity as a dimension of AI-HRM quality. The paper consists of an overextension, which involves bringing them into dialogue with the literature on algorithmic fairness.

References

1. African Union. (2020). The digital transformation strategy for Africa (2020–2030). African Union Commission
2. African Union. (2022). Data policy framework for Africa. African Union Commission
3. Alchian, A. A., & Demsetz, H. (1975). Production, information costs, and economic organization. *IEEE Engineering Management Review*, 3(2), 21–41. <https://doi.org/10.1109/emr.1975.4306431>
4. Amadiume, I. (1987). *Male daughters, female husbands: Gender and sex in an African society*. Zed Books.
5. Arnstein, S. R. (1969). A ladder of citizen participation. *Journal of the American Institute of Planners*, 35(4), 216–224. <https://doi.org/10.1080/01944366908977225>
6. Asante, M. K. (1987). *The Afrocentric idea*. Temple University Press.
7. Ayana, G., Dese, K., Nemoassa, H. D., Habtamu, B., Mellado, B., Badu, K., Yamba, E., Faye, S. L., Ondua, M., Nsagha, D., Nkweteyim, D., & Kong, J. D. (2024). Decolonizing global AI governance: Assessment of the state of decolonized AI governance in Sub-Saharan Africa. *Royal Society Open Science*, 11(8), 231994. <https://doi.org/10.1098/rsos.231994>
8. Bandara, R. J., Biswas, K., Akter, S., Shafique, S., & Rahman, M. (2025). Addressing algorithmic bias in AI-driven HRM systems: Implications for strategic HRM effectiveness. *Human Resource Management Journal*, 35(4), 1047–1063. <https://doi.org/10.1111/1748-8583.12609>
9. Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4), 991–1013. <https://doi.org/10.1257/0002828042002561>
10. Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), 100205. <https://doi.org/10.1016/j.patter.2021.100205>
11. Birhane, A., Isaac, W., Prabhakaran, V., Díaz, M., Elish, M. C., Gabriel, I., & Mohamed, S. (2022). Power to the People? Opportunities and Challenges for Participatory AI. *arXiv* (Cornell University). <https://doi.org/10.48550/arxiv.2209.07572>
12. Blodgett, S. L., Barocas, S., Daumé, H., III, & Wallach, H. (2020). Language (technology) is power: A critical survey of “bias” in NLP. <https://doi.org/10.48550/arXiv.2005.14050>
13. Bolino, M. C., Turnley, W. H., & Bloodgood, J. M. (2002). Citizenship behavior and the creation of social capital in organizations. *Academy of Management Review*, 27(4), 505–522. <https://doi.org/10.2307/4134400>
14. Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press. <https://doi.org/10.7551/mitpress/6352.001.0001>

15. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 77–91.
16. Crenshaw, K. (1989). Demarginalizing the intersection of race and sex. *University of Chicago Legal Forum*, 139–167.
17. Creighton-Randall O., (2024). “The Applicability of the Ubuntu Philosophy to Building Organizational Communities and its impact on Organizational Performance”. *International Journal of Business Research and Management* 1(1); DOI: <http://doi.org/112.2024/07.1010>.
18. Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52(4), 281–302. <https://doi.org/10.1037/h0040957>
19. Santos, B. de S. (2014). *Epistemologies of the south*. Paradigm Publishers.
20. DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited. *American Sociological Review*, 48(2), 147–160. <https://doi.org/10.2307/2095101>
21. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2011). Fairness through awareness. <https://doi.org/10.48550/arxiv.1104.3913>
22. Fabris, A., Baranowska, N., Dennis, M. J., Graus, D., Hacker, P., Saldivar, J., Borgesius, F. Z., & Biega, A. J. (2024). Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Transactions on Intelligent Systems and Technology*, 16(1), 1–54. <https://doi.org/10.1145/3696457>
23. Foucault, M. (1977). *Discipline and punish: The birth of the prison*. Pantheon.
24. Frenkel, M. (2008). The multinational corporation as a third space. *Academy of Management Review*, 33(4), 924–942. <https://doi.org/10.5465/amr.2008.34422002>
25. Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>
26. Ghajiga, G. S., Warlimont, D. J., & Sebakwane, L. (2025). Epistemic decolonization in the 21st century. *Ovidius University Annals, Economic Sciences Series*, 25(2), 93–101.
27. Gilson, L. L., & Goldberg, C. B. (2015). Editors’ comment. *Group & Organization Management*, 40(2), 127–130. <https://doi.org/10.1177/1059601115576425>
28. Grobler, A., & Powell, K. (2025). Ubuntu leadership. *Frontiers in Psychology*, 16, 1686493. <https://doi.org/10.3389/fpsyg.2025.1686493>
29. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Proceedings of NIPS*, 3323–3331.
30. Hofstede, G. (1980). *Culture’s consequences*. Sage.
31. Honneth, A. (1995). *The struggle for recognition*. MIT Press.
32. Inyang, B. J. (2009). Indigenous management theories in Africa. *International Journal of Business and Management*, 3(12). <https://doi.org/10.5539/ijbm.v3n12p122>
33. Jaakkola, E. (2020). Designing conceptual articles. *AMS Review*, 10(1–2), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
34. Jack, G., & Westwood, R. I. (2009). *International and cross-cultural management studies*.
35. Jackson, T. (2025). Beyond postcolonial theory. *International Journal of Cross Cultural Management*, 25(3), 581–585. <https://doi.org/10.1177/14705958251394240>

36. Kim, S., Khoreva, V., & Vaiman, V. (2024). Strategic HRM in the algorithmic era. *Human Resource Management*, 64(2), 447–464. <https://doi.org/10.1002/hrm.22268>
37. Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Toups, C., Rickford, J. R., Jurafsky, D., & Goel, S. (2020). Racial disparities in speech recognition. *PNAS*, 117(14), 7684–7689. <https://doi.org/10.1073/pnas.1915768117>
38. Köchling, A., & Wehner, M. C. (2020). Discrimination by algorithm. *Business Research*, 13(3), 795–848. <https://doi.org/10.1007/s40685-020-00134-w>
39. Kumar, D. (2025). Epistemic justice in AI. *Omniscient*, 3(3), 83–95.
40. Kusner, M. J., Loftus, J. R., Russell, C., & Silva, R. (2017). Counterfactual fairness. <https://doi.org/10.48550/arxiv.1703.06856>
41. Lee, M. K., et al. (2019). WeBuildAI. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–35. <https://doi.org/10.1145/3359283>
42. Mangaliso, M. P. (2001). Ubuntu and competitive advantage. <https://doi.org/10.2139/ssrn.1762783>
43. Mangaliso, M. P., Mangaliso, Z., Knipes, B. J., Jean-Denis, H., & Ndanga, L. (2018). Invoking Ubuntu philosophy as a source of harmonious organizational management. *Academy of Management Proceedings*, 2018(1), 15007. <https://doi.org/10.5465/ambpp.2018.15007abstract>
44. Mayasandra, R., Pan, S. L., & Myers, M. D. (2006). Viewing IT outsourcing organizations through a postcolonial lens. *IFIP*, 208, 381–396.
45. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). Bias in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
46. Messick, S. (1995). Validity of psychological assessment. *American Psychologist*, 50(9), 741–749.
47. Metz, T. (2011). Ubuntu as moral theory. *African Human Rights Law Journal*, 11(2), 532–559.
48. Miceli, M., et al. (2025). Workers' epistemic authority in AI. *AIES Conference Proceedings*, 1698–1710. <https://doi.org/10.1609/aies.v8i2.36667>
49. Milano, S., & Prunkl, C. (2025). Algorithmic profiling and injustice. *Philosophical Studies*, 182, 185–203. <https://doi.org/10.1007/s11098-023-02095-2>
50. Milella, F., & Cabitza, F. (2026). Perceiving AI as an epistemic authority or algorithm: A user study on the human attribution of authority to AI. *Machine Learning and Knowledge Extraction*, 8(2), 36. <https://doi.org/10.3390/make8020036>
51. Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI. *Philosophy & Technology*, 33, 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
52. Muldoon, J., & Wu, B. A. (2023). AI in the colonial matrix of power. *Philosophy & Technology*, 36, 80. <https://doi.org/10.1007/s13347-023-00687-8>
53. Nahapiet, J., & Ghoshal, S. (1998). Social capital and advantage. *Academy of Management Review*, 23(2), 242–266.
54. Nunn, N. (2008). Effects of Africa's slave trades. *Quarterly Journal of Economics*, 123(1), 139–176.
55. Odion, A. S., & Agbonwanegbe, O. E. (2025). Feminist and post-colonial perspectives in management. *FUDMA International Journal of Social Sciences*, 5(4), 76–91.
56. Organ, D. W. (1988). *Organizational citizenship behavior*. Lexington Books.

57. Oyěwùmí, O. (1997). *The invention of women*. University of Minnesota Press.
58. Quijano, A. (2000). Coloniality of power. *International Sociology*, 15(2), 215–232.
59. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Bias in algorithmic hiring. *FAT Conference Proceedings*, 469–481. <https://doi.org/10.1145/3351095.3372828>
60. Sandvig, C., et al. (2014). *Auditing algorithms*. ICA Conference.
61. Shin, C. (2022). *Can hiring algorithms be impartial decision-makers?* Seton Hall Law School Student Scholarship.
62. Sloane, M., Moss, E., Awomolo, O., & Forlano, L. (2020). Participation is not a Design Fix for Machine Learning. *ResearchGate*. <https://doi.org/10.48550/arXiv.2007.02423>
63. Snyder, H. (2019). Literature review methodology. *Journal of Business Research*, 104, 333–339.
64. Tayali, E. M. (2025). Unlocking mutual gains through Ubuntu : a reimagined approach to human resource management practices and employee well-being : a case of Zambia. In *STAX*. <https://doi.org/10.48730/b2hz-ta94>
65. Tambe, P., Cappelli, P., & Yakubovich, V. (2019). AI in HRM. *California Management Review*, 61, 15–42.
66. Torraco, R. J. (2005). Integrative literature reviews. *Human Resource Development Review*, 4(3), 356–367.
67. Turner, B. (2025). Workplace democracy and algorithmic management. *Critical Review of International Social and Political Philosophy*, 1–25.
68. Webster, J., & Watson, R. T. (2002). Writing literature reviews. *MIS Quarterly*, 26(2), xiii–xxiii.



©2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by-nc-sa/4.0/>).