Unpacking Learner Agency and Epistemic Justice in AI-Augmented Open and Distance Learning for Marginalised Students

Bunmi Isaiah Omodan[1,*] & Sindile Amina Ngubane[1]

[1]Institute for Open and Distance Learning, College of Education, University of South Africa, South Africa

*Correspondence: Institute for Open and Distance Learning, College of Education, University of South Africa, South Africa. E-mail: oomodanbunmi@gmail.com. ORCID: https://orcid.org/0000-0002-9093-3108

Received: October 1, 2025 Accepted: November 27, 2025 Online Published: February 6, 2026

doi:10.5430/jct.v15n1p227 URL: https://doi.org/10.5430/jct.v15n1p227

Abstract

Intelligent learning technologies are reshaping university-based Open and Distance Learning (ODL), raising questions about how digitally mediated pedagogies redistribute knowledge, power, and participation for students from structurally marginalised contexts. Moving beyond issues of access, efficiency, and automation, this paper examines the epistemological implications of AI-augmented ODL by analysing learner agency and epistemic justice. Drawing on critical pedagogy and theories of epistemic injustice, we employ a theory-synthesis design to integrate conceptual and empirical work on ODL, learning analytics, and algorithmic personalisation. The analysis demonstrates how predictive models, recommendation engines, and automated moderation can both scaffold and constrain autonomy, participation, and recognition; it identifies pathways of testimonial and hermeneutical injustice when systems misinterpret non-dominant discourse, proxy sensitive attributes, or inhibit contestation. In response, we propose a critical framework that redefines learner agency as relational and contextually situated, organised around three dimensions—recognition, voice, and power—and operationalised through three design checkpoints—credibility, comprehensibility, and control (CCC). We outline practical applications for course and assessment design, analytics pipelines, student data rights, and continuous improvement, while specifying policy requirements for participatory governance, impact assessment, and repair. The framework is intended to be applicable across regions globally where similar platform logics circulate, while remaining adaptable to local conditions. Ultimately, the paper offers a justice-oriented lens for reimagining ODL as a space where student dignity, cultural relevance, and equitable participation in knowledge construction are central. The study contributes to knowledge by delivering a coherent, testable model and CCC criteria that translate epistemic justice theory into actionable design and governance for AI-mediated ODL.

Keywords: learner agency, epistemic justice, university ODL systems, theory synthesis, marginalised students

1. Introduction

The steady diffusion of intelligent learning technologies into university Open and Distance Learning (ODL) has transformed the organisation, personalisation, and measurement of instruction. From adaptive courseware and algorithmic recommendations to AI teaching assistants and learning analytics dashboards, the promise lies in scalability, efficiency, and responsiveness to learner variability (UNESCO, 2021; Zawacki-Richter et al., 2019). However, as university ODL increasingly operates through data-intensive infrastructures, decision-making is progressively delegated to automated systems that classify, predict, and optimise student pathways (Williamson, 2017). This technological shift is situated within global policy currents, exemplified by the Beijing Consensus, which frames AI as both an opportunity and a risk for equitable quality education (UNESCO, 2019). For institutions that already rely on digital mediation to widen participation, the question is no longer whether AI will shape ODL, but how its logics will reconfigure what constitutes learning, who is recognised as a learner, and whose knowledge is deemed legitimate within academic contexts (UNESCO, 2021; Zawacki-Richter et al., 2019; UNESCO, 2019; Williamson, 2017).

A narrow focus on access and efficiency obscures significant epistemological and political stakes: while AI may numerically widen participation, it may simultaneously narrow the voices and experiences acknowledged by the system. Critical scholarship documents how algorithmic systems reproduce social hierarchies through biased training data, model design, and optimisation targets—amplifying harms for historically marginalised groups (Benjamin, 2019; Noble, 2018). In the field of education, computer vision and classification systems have demonstrated unequal error rates across intersectional identities, evidenced in the materiality of bias within "personalised" technologies (Buolamwini & Gebru, 2018). These unequal error rates and embedded assumptions carry direct implications for marginalised students' sense of urgency, agency, and participation in ODL. When systems systematically misclassify or under-recognise certain linguistic, cultural, or behavioural patterns, students may internalise these automated judgments as indicators of low ability, thereby reducing their confidence to intervene, participate, or persist. In distance contexts—where immediacy, self-regulation, and prompt engagement are already critical—the misinterpretation of learner behaviour as "risk" can create artificial urgency, constrain autonomy, and deter exploratory learning. By *"materiality of* bias," we refer to how bias does not remain at the level of data or model design but becomes materially consequential in how personalised technologies distribute visibility, feedback, and learning opportunities; these outputs shape which tasks students see, how their contributions are ranked, and the credibility assigned to their participation. In marginalised ODL contexts, such consequences accumulate into epistemic harms that diminish learners' perceived legitimacy and restrict their ability to shape their learning trajectories. When such tools curate content, flag "risk," or direct learner trajectories, they embed contestable assumptions regarding ability, language, culture, and "fit." The issue is not merely one of technical unfairness but rather the re-inscription of power: who has the authority to frame the problem, set objectives, and validate knowledge claims within ODL platforms governed by opaque optimisation processes (Noble, 2018; Benjamin, 2019; Buolamwini & Gebru, 2018; Williamson, 2017).

To interrogate these dynamics, this paper advances a conceptual exploration of learner agency and epistemic justice in AI-augmented Open and Distance Learning (ODL). Epistemic justice highlights how individuals can be wronged "as knowers," whether through testimonial injustice (where their credibility is discounted) or hermeneutical injustice (where structural gaps render their experiences unintelligible) (Fricker, 2007). Building on this, Dotson (2014) theorises epistemic oppression as patterned exclusions from participation in knowledge production. When viewed through the lens of ODL, these frameworks illuminate how algorithmic gatekeeping can silence student voices, pre-sort credibility, and normalise deficit framings under the veneer of objectivity. A critical pedagogy lens—drawing upon Freire's (2000) critique of "banking" education—shifts the analysis towards dialogical, participatory, and culturally responsive knowledge-making, recentering learner agency as a relational capacity to act, be heard, and transform one's learning conditions. Together, these perspectives justify a shift from evaluating AI for "what works" to asking for whom, by whose standards, and at what epistemic cost? (Fricker, 2007; Dotson, 2014; Freire, 2000).

The significance of this reframing is particularly critical for marginalised students, for whom distance and openness do not automatically translate into parity of participation. Social justice traditions in ODL—such as capability approach analyses—warn that widening access without addressing recognition, voice, and agency can entrench old exclusions in new technical forms (Tait, 2013). In AI policy, UNESCO (2021, 2019) similarly urges human-centred design that protects rights, transparency, and inclusion. Fraser's (2009) tripartite model—redistribution, recognition, and representation—clarifies the multi-dimensional remedies required: resource provision (devices, connectivity), cultural-epistemic respect (valuing diverse ways of knowing), and political inclusion (students' input in how platforms classify and intervene). Within AI-mediated ODL, this entails challenging personalisation that individualises structural barriers, surfacing community knowledge as curricular resources, and opening model objectives and error costs to those most affected. Without such commitments, algorithmic "support" risks becoming a sophisticated engine of epistemic silencing (Tait, 2013; UNESCO, 2021; UNESCO, 2019; Fraser, 2009).

Although institutional contexts may differ, the mechanisms at issue—datafication, optimisation, and platformised governance—are transnational in nature. The same vendors, models, and "best practices" circulate across regions, while open and distance learning (ODL) cohorts worldwide include students confronting issues such as poverty, linguistic diversity, migration, disability, or racialised exclusion. A region-specific analysis would overlook how AI logics traverse boundaries and how questions of justice extend beyond national frameworks—aligning with Fraser's (2009) argument regarding the shifting "scale" of justice within a globalised order. A global conceptual lens thus illuminates shared risks and design principles that can subsequently be adapted locally—honouring plurality without succumbing to technological determinism (Fraser, 2009; UNESCO, 2021; Zawacki-Richter et al., 2019; Williamson, 2017). Therefore, the integration of AI in ODL is not merely a technical or managerial concern; it represents an epistemic and political endeavour with far-reaching implications regarding whose knowledge is valued and how learners can engage with their environments (Fraser, 2009; UNESCO, 2021; Zawacki-Richter et al., 2019; Williamson, 2017).

In this article, marginalisation is defined primarily in relation to learners who face linguistic disadvantage (e.g., multilingual or non-standard English users), socioeconomic precarity, geographic remoteness, and constraints arising from caregiving or work commitments, factors that influence how AI-mediated systems interpret their behaviours and contributions. These dimensions are particularly significant in Open and Distance Learning (ODL) contexts, where algorithmic features heavily rely on language patterns, engagement proxies, and data dependent on connectivity. Clearly delineating these boundaries enhances the explanatory power of the analysis by elucidating which student groups are most likely to encounter epistemic and participatory harms within AI-augmented learning environments.

In light of this, conceptual exploration provides an answer to the following research question:

- How do AI-augmented ODL systems shape the epistemic agency of marginalised university students, and what framework of redistribution, recognition, and representation can guide the design of epistemically just, relational forms of learner agency in such systems?

2. Conceptual and Theoretical Framework

This study employs three interconnected lenses—critical pedagogy, epistemic justice, and learner agency—to examine how AI-supported Open and Distance Learning (ODL) can either assist or hinder marginalised students. In this study, AI is regarded not only as a tool for efficiency but also as a factor that influences what is considered knowledge, whose knowledge is valued, and how students can engage within their learning environments. Critical pedagogy highlights the importance of dialogue, shared meaning-making, and the human purposes of education. Epistemic justice helps to identify and address injustices that unfairly impact individuals as knowers. Learner agency focuses on the genuine choices students have to express themselves, make decisions, and influence their learning and the data surrounding it (Biesta, 2013; Freire, 2000; Fricker, 2007). Together, these lenses keep recognition, voice, and power at the forefront of systems that classify, predict, and personalise learning (Noble, 2018; Williamson, 2017; Zawacki-Richter et al., 2019).

2.1 Critical Pedagogy: Core Principles and Relevance to ODL

Critical pedagogy is the first and most foundational lens in this framework. It begins with a simple idea: education should be dialogic and transformative, rather than a one-way transfer of information. Freire (2000) criticises the "banking" model in which teachers deposit facts into passive learners. He argues that students are co-creators of knowledge and that critical awareness (conscientização) develops through problem-posing dialogue. Hooks (1994) adds that teaching should be engaged and caring, valuing students' identities and voices. Giroux (2011) reminds us that education is part of democratic life and should help students question and reshape the conditions that affect them. Biesta (2013) emphasises that education is not merely about qualifying or socialising learners; it is also about helping them become subjects who can respond and take responsibility in the world. These ideas are particularly relevant for ODL, where platforms mediate most interactions. If automation is allowed to replace judgement and dialogue, the world risks reducing learning to a series of clicks and scores (Selwyn, 2016; Veletsianos, 2020).

A critical pedagogy approach to AI in ODL raises very practical questions: Do the discussion spaces genuinely support back-and-forth dialogue, or are they driven solely by "engagement" metrics? Do the assessments invite students to bring local knowledge, languages, and examples that resonate with their contexts? Do students have a say in how platform rules are formulated, how content is recommended, and how risk is labelled? In concrete terms, this means fostering structured yet open discussions, peer reviews that value diverse voices, and assignments that allow multiple ways to demonstrate learning (e.g., text, audio, community examples). It also means involving students in the co-design of platform features and policies so they possess real decision-making rights, rather than just providing feedback through forms. These steps are essential because AI-based personalisation can otherwise treat structural barriers—such as poor connectivity, multilingual writing, or care responsibilities—as individual "deficits," which undermines the social justice aims of distance education (Tait, 2013; Veletsianos, 2020). In short, critical pedagogy sets the tone: AI should amplify dialogue, reciprocity, and human judgement, not automate them away (Giroux, 2011; Hooks, 1994; Selwyn, 2016).

2.2 Epistemic Justice: Adapting Fricker to Digital Learning

"Epistemic justice is the second lens and explains how people can be wronged as knowers. Fricker (2007) describes two forms. Testimonial injustice occurs when prejudice leads others to give a speaker less credibility than they deserve. Hermeneutical injustice occurs when there are gaps in shared concepts, so some experiences cannot be easily expressed or understood. Dotson (2014) calls the broader pattern epistemic oppression, where groups are routinely blocked from participating fully in shared knowledge. Pohlhaus (2012) shows how willful hermeneutical ignorance keeps those gaps in place, and Medina (2013) argues that communities need epistemic virtues—humility, vigilance, and resistance—to counter this.

"In AI-mediated ODL, these ideas are not abstract. Testimonial injustice can show up when automated moderation, remote proctoring, or risk models over-flag students whose accents, language varieties, or writing styles differ from dominant norms, lowering how credible they appear (Benjamin, 2019; Noble, 2018). Hermeneutical injustice appears when platform categories cannot 'see' caregiving responsibilities, unstable internet, or multilingual code-meshing, so systems read struggle as apathy or low ability (Slade & Prinsloo, 2013; Williamson, 2017). Because commercial models move across institutions and countries, these harms can scale quickly.

"Addressing them means combining algorithmic signals with human review (credibility calibration), co-creating wider taxonomies and success indicators with students (hermeneutical expansion), and providing clear ways to challenge and correct labels (contestability and repair) (Kidd et al., 2017; Noble, 2018; Slade & Prinsloo, 2013). Put simply, epistemic justice becomes a set of design requirements: transparency about features and thresholds, shared governance of models, and audits that consider both numbers and lived experience (Benjamin, 2019; Kidd et al., 2017)."

2.3 Learner Agency: Relational Capacities, Recognition, Voice, and Power

Learner agency is the third lens and focuses on students' real options to shape their learning. Agency is not only about self-management; it is relational and depends on time and context. Emirbayer and Mische (1998) describe agency as drawing on the past, imagining the future, and making practical judgments in the present. Bandura (2001, 2006) highlights intentionality, forethought, self-regulation, and self-reflection. Biesta (2013) adds that agency grows when education invites the student to speak and respond as a subject. In ODL, agency is distributed across people and systems: peers, tutors, institutional rules, and the platform itself. Students demonstrate agency when they are recognised as credible knowers, when they have a voice to shape meaning, and when they hold the power to influence data use and learning paths (Tait, 2013; Veletsianos, 2020; Wenger, 1998). AI can support agency by offering adaptive scaffolds, multimodal access, and timely formative feedback. However, it can also limit agency by narrowing choices, pushing predictive nudges that restrict exploration, or obscuring the rules that rank posts or route students through content (Reich, 2020; Williamson, 2017; Zawacki-Richter et al., 2019). Agency-supportive designs, therefore, widen decision latitude (students choose tasks, formats, and exemplars; systems adapt to student-stated goals rather than only past averages), build voice infrastructures (dialogue, reflective writing, community annotation, and alternative assessments where students co-define success), and enable data sovereignty (clear dashboards, data access and portability, and the right to refuse certain inferences without penalty) (Slade & Prinsloo, 2013; Veletsianos, 2020; Wenger, 1998).

2.4 Integrating the Lenses: A Composite Framework for AI-augmented ODL

Bringing all three lenses together provides a simple, working framework. From critical pedagogy, the paper advocates for a dialogue-first design: assessing AI features based on whether they sustain inquiry, reciprocity, and collective meaning-making, particularly for students who have been excluded (Freire, 2000; hooks, 1994). From epistemic justice, it incorporates recognition and repair into data pipelines: identifying where testimonial and hermeneutical injustices occur (in data collection, labelling, modelling, and deployment) and developing ways to broaden concepts and challenge results (Fricker, 2007; Dotson, 2014; Medina, 2013). From learner agency, it prioritises capability-enhancing affordances: AI should expand students' real freedoms to set goals, make and justify choices, and shape the rules of classification and evaluation (Bandura, 2006; Biesta, 2013; Sen, 1999). This is summarised as a practical triad—Credibility, Comprehensibility, and Control. "Credibility" examines whether systems avoid credibility deficits and amplify marginalised voices; "Comprehensibility" considers whether categories and thresholds are understandable and open to revision; and "Control" assesses whether students have genuine choices over data practices and learning pathways, including the right to refuse certain inferences without penalty (Benjamin, 2019; Noble, 2018; Slade & Prinsloo, 2013). Used in conjunction, these commitments align AI-mediated ODL with democratic, justice-oriented education.

In this article, the integrated framework serves two purposes. As an analytic lens, it guides how we read the literature and cases: we examine how ODL systems distribute credibility, expand or limit shared meanings, and allocate control over learning and data. As a design heuristic, it provides criteria for evaluating or improving AI features: dialogical fit (critical pedagogy), recognition/repair mechanisms (epistemic justice), and capability-enhancing affordances (agency). Because models, platforms, and governance patterns cross borders, these criteria are relevant in many regions, even though specifics may vary locally (Noble, 2018; Williamson, 2017; Zawacki-Richter et al., 2019). The next section explains the theory-synthesis method we use to assemble and integrate these literatures into a clear analytical scaffold for studying learner agency and epistemic justice in AI-augmented ODL.

3. Literature Review

3.1 AI in ODL: Access, Efficiency, and Automation

Empirical and review literature indicate that AI and data-driven tools are increasingly implemented in higher education to enhance access, improve efficiency, and automate routine instructional and support tasks. A comprehensive systematic review encompassing 146 studies conducted between 2007 and 2018 revealed that research predominantly concentrates on intelligent tutoring systems, automated assessment and feedback, adaptive technologies, and learning analytics. Conversely, a conceptual gap exists in studies addressing pedagogy and the roles of educators—an imbalance that reflects the implementation priorities observed within ODL contexts (Zawacki-Richter, Marín, Bond, & Gouverneur, 2019). Evidence derived from school-level field studies suggests that "personalised learning" models have the potential to enhance test scores under specific conditions; however, the effects are contextually dependent and vary according to design. Moreover, many of these models depend on extensive data collection and algorithmic recommendations (Pane et al., 2015). Concurrently, research on predictive modelling and adaptive tutoring has reported improvements in efficiency or learning outcomes for particular tasks (e.g., task selection in intelligent tutoring systems). Nevertheless, much of this research is conducted in controlled environments or utilises data and outcome measures that inadequately represent the complexities associated with ODL cohorts (Aleven, McLaughlin, Glenn, & Koedinger, 2017; Gardner & Brooks, 2018). Reviews of MOOC analytics further illustrate rapid advancements in prediction, intervention, and personalisation pipelines, yet they also highlight methodological deficiencies (e.g., subpopulation filtering, evaluation of non-deployable features) that restrict external validity in real-world ODL scenarios (Gardner & Brooks, 2018). Overall, the prevailing narrative within the evidence base emphasises scale and optimisation; however, there are fewer studies that critically examine whether these advancements lead to equitable participation and recognition for a diverse range of distance learners (Zawacki-Richter et al., 2019).

3.2 Marginalised Students in ODL: Barriers and Opportunities

A substantial empirical tradition identifies persistent barriers for marginalised students learning online. Large-scale MOOC analyses show that access and completion often correlate with socioeconomic advantage, revealing participation and persistence gaps related to social class and geography—raising concerns about whether “open” formats alone can democratise opportunity (Hansen & Reich, 2015; Kizilcec, Davis, & Cohen, 2017). Studies on community-college and STEM courses find differential outcomes online compared to face-to-face for groups defined by ethnicity, gender, and non-traditional status, suggesting that modality interacts with prior preparation and support structures (Hachey, Conway, & Wladis, 2015; Wladis, Hachey, & Conway, 2015). Retention syntheses report high attrition rates (often 40–80%), attributing the risks to social isolation, time pressure, inconsistent connectivity and devices, limited instructor presence, and a weaker sense of belonging—factors that disproportionately affect first-generation, low-income, rural, refugee, and caregiving students (Bawa, 2016; Hart, 2012; Brunton & Buckley, 2019).

Qualitative research with first-in-family and open-entry students documents ODL’s opportunity value (flexibility and entry points without traditional prerequisites) alongside struggles with confidence, institutional navigation, and competing responsibilities; these findings underline the importance of culturally responsive design and wrap-around support (Stone, O’Shea, May, Delahunty, & Partington, 2016). Comparative studies also show that prior online experience and GPA predict success in online STEM courses, indicating a cumulative advantage dynamic unless scaffolds are provided (Hachey et al., 2015). At the same time, evidence from distance higher education suggests that ODL can deliver strong labour-market returns for adult learners and may widen participation if institutions adapt content, assessment, and support to local constraints (Wang, 2023). In short, the literature positions ODL as a conditional enabler: flexibility and access are meaningful, but without design and governance that address recognition, voice, and resources, marginalised students face patterned disadvantages.

Intervention studies in large online settings offer mixed results. Short, light-touch psychological interventions can reduce some gaps in specific contexts, but their effects vary by culture and subgroup and are difficult to target reliably at scale (Kizilcec, 2017; Kizilcec, Davis, & Cohen, 2020). For example, a massive field experiment that used state-of-the-art machine learning to target behavioural interventions did not improve outcomes beyond assigning the same support to all or randomly selected students, suggesting limits to predictive targeting for equity in large online courses (Kizilcec, Davis, Lakhani, & Cohen, 2020). These findings echo retention reviews: support needs to be systemic (community, instructor presence, assessment design, and technology access), not merely individualised via dashboards and nudges (Bawa, 2016; Hart, 2012).

3.3 Algorithmic Personalisation and Data-Driven Decision-Making: Benefits, Limits, and Risks

A second body of work critiques how algorithmic systems classify, predict, and act within educational settings, including Open and Distance Learning (ODL). Reviews of algorithmic bias in education document concrete disparities in model performance and subsequent decisions, identifying sources of bias across measurement, model training, and action policies (Baker & Hawn, 2022). Specifically in online higher education, studies show that dropout and success models often optimise for narrow outcome proxies (e.g., short-term completion or GPA), which can perpetuate prior inequities and misalign with broader learning goals (Gardner & Brooks, 2018). The literature on algorithmic fairness in education proposes a lifecycle view—measurement, model learning, and action—as three levers to diagnose and mitigate inequities; however, empirical tests in real deployments remain limited and sometimes report null effects when models are used to target support (Kizilcec & Lee, 2020/2022; Kizilcec et al., 2020). At the system level, sociotechnical analyses of learning analytics argue that clustering and personalisation do not merely reflect learning; they actively shape categories, incentives, and participation—raising questions about whose behaviours are normalised or problematised (Perrotta & Williamson, 2018).

Several empirical strands help specify these risks. First, predictive-risk frameworks that identify “at-risk” students can concentrate false positives and negatives within particular subgroups, especially when training data is skewed and features proxy sensitive characteristics (Lakkaraju, Aguiar, Bhanpuri, Miller, & Ghani, 2015; Gardner & Brooks, 2018). Second, even accurate predictions can lead to harmful actions if institutions opt for interventions that stigmatise, over-monitor, or track students into narrower pathways (Kizilcec & Lee, 2020/2022). Third, challenges regarding explainability mean that stakeholders may not comprehend model logics or thresholds, limiting meaningful contestation and repair; critiques of explainability in legal scholarship highlight how machine learning models can be simultaneously inscrutable and non-intuitive (Selbst & Barocas, 2018). In response, emerging work advocates for “fairness-by-design” and participatory governance of educational machine learning, emphasising impact assessments that evaluate not only statistical parity but also the educational value and costs associated with mistaken classifications (Liu, Dean, Rolf, Simchowitz, & Hardt, 2023; Baker & Hawn, 2022).

For Open and Distance Learning (ODL), where mediation is the default and cohorts are diverse, these critiques are highly relevant. Studies at MOOC scale demonstrate strong predictive performance on average; however, generalisation across subpopulations is inconsistent, relevant features are often absent in live settings, and interventions may fail to improve outcomes compared to simple policies (Gardner & Brooks, 2018; Kizilcec et al., 2020). Moreover, retention and equity literature cautions that dashboards and nudges—without enhanced instructor presence, community engagement, and resource support—rarely address the structural factors that contribute to risk (Bawa, 2016; Hart, 2012). Concurrently, adaptive tutoring and analytics can offer significant value when integrated into intentional pedagogical design, utilising outcomes aligned with learning goals and incorporating mechanisms for student voice and contestation (Aleven et al., 2017; Perrotta & Williamson, 2018). The mixed empirical evidence suggests that AI-enabled personalisation should be regarded as a socio-technical intervention rather than merely a technical improvement: its effects depend on the specification of models, the selection of actions, and the distribution of power and participation in ODL contexts that serve many marginalised learners.

Collectively, the literature highlights three gaps that this article addresses. First, evidence regarding AI in ODL disproportionately focuses on access, scale, and predictive accuracy, while paying insufficient attention to the recognition and voice of diverse learners. Second, studies involving marginalised students document enduring barriers and inconsistent returns to online learning, suggesting that design and governance—not solely access—determine who benefits. Third, empirical critiques of algorithmic personalisation reveal that predictive targeting alone is not a reliable strategy for equity at scale and can introduce new risks without participatory safeguards. Building on these insights, the next section outlines the methodology employed to synthesise and analyse the literature through the study’s critical pedagogy, epistemic justice, and learner agency lenses.

4. Methodology

This article employs a theory synthesis design to build an integrative framework that connects critical pedagogy and epistemic justice with empirical and conceptual work on AI-augmented Open and Distance Learning (ODL). Following core principles of theory synthesis, we (a) specified the focal constructs (learner agency, recognition/voice/power, testimonial and hermeneutical (in)justice, and AI-mediated personalization), (b) identified and clustered relational statements across sources (e.g., how platform logics shape credibility, interpretive resources, and decision latitude), and (c) organised these into a coherent set of propositions and design heuristics (dialogue-first design; recognition and repair; capability-enhancing affordances) that can be applied across contexts (Walker & Avant, 2019). Consistent with knowledge-development guidance, we treated theory synthesis as a systematic yet creative process of assembling concepts and relationships from diverse literatures to generate a middle-range, practice-oriented account, while attending to clarity, parsimony, scope, and pragmatic utility (Chinn & Kramer, 2018). Methodologically, the work follows an applied theory-building cycle that emphasises iterative conceptual development: we purposively utilised peer-reviewed studies from ODL, learning analytics/AI in education, and justice-oriented pedagogy; extracted and coded statements about mechanisms linking AI features to agency and epistemic (in)justice; compared convergent and divergent claims; and refined the integrated propositions through constant comparative analysis and analytic memoing until theoretical saturation for the focal constructs (Lynham, 2002). The scope included both empirical and conceptual works to ensure the synthesised theory is grounded in observed patterns while remaining general enough to guide design and policy across regions. The result is a portable analytic scaffold that we now apply to the evidence base: the next section (Analysis and Discussion) uses the synthesised framework to examine how AI-enabled ODL practices shape learner agency and epistemic justice, and to derive implications for institutional design and governance.

5. Analysis and Discussion

Below, we apply the synthesised framework to the evidence base to examine how AI-mediated ODL shapes learners' real opportunities to act, be heard, and be recognised. Guided by critical pedagogy, epistemic justice, and learner agency, the analysis interrogates not only whether systems "work," but for whom, by what mechanisms, and at what epistemic cost. We organise the discussion into three themes that progress from system behaviour to learner experience: (1) AI-augmented ODL and learner agency, (2) epistemic (in)justice in digital learning systems, and (3) relational and contextual agency. We begin with the first theme.

5.1 AI-augmented ODL and Learner Agency: Autonomy, Participation, Recognition

Across the evidence base, algorithmic systems in ODL demonstrably reorganise the conditions under which students can act, participate, and be recognised as credible knowers. On one hand, adaptive tutors and analytics can widen access to timely feedback and varied learning pathways, which—when aligned with pedagogy—can support intentionality, forethought, and self-regulation, key components of agency (Aleven et al., 2017; Bandura, 2001, 2006). Large-scale implementations oriented towards "personalised learning" also report gains under certain conditions, suggesting potential for scaffolding autonomy when design and teacher capacity are robust (Pane et al., 2015). However, the same literature warns that many systems optimise narrow proxies (e.g., clickstream persistence, short-term completion), thereby scripting participation in ways that conflate activity with learning and treat students as objects of prediction rather than subjects of education (Gardner & Brooks, 2018; Williamson, 2017; Zawacki-Richter et al., 2019). From a critical-pedagogy standpoint, agency depends on dialogical, humanising relations (Freire, 2000; Biesta, 2013). Designs that push predictive nudges, rank discourse by opaque "engagement," or automate assessment without room for contestation risk replacing dialogue with datafied signalling, suppressing the very reciprocity and judgement that sustain learner voice (Selwyn, 2016; Veletsianos, 2020). Empirically, large-scale targeting of supports using state-of-the-art models has not reliably improved outcomes over simple or unguided allocation, underscoring that prediction without participatory action design weakly translates to gains in agency (Kizilcec, Davis, Lakhani, & Cohen, 2020). In summary, artificial intelligence can extend autonomy if its objectives, feedback, and interventions are embedded in pedagogy that invites choice, explanation, and co-definition of success; otherwise, automation risks narrowing participation to what systems can easily measure and recognise (Freire, 2000; Gardner & Brooks, 2018; Zawacki-Richter et al., 2019).

5.2 Epistemic (in)justice in Digital Learning Systems: Asymmetries, Silencing, Deficit Framings

Through the lens of epistemic justice, common ODL practices reveal pathways to testimonial and hermeneutical harm. Testimonial injustice arises when automated moderation, plagiarism detectors, remote proctoring, or "risk" classifiers differentially flag contributions from students whose language varieties, accents, or rhetorical styles deviate from dominant norms—mechanisms that can diminish perceived credibility and inhibit participation (Benjamin, 2019; Noble, 2018). Hermeneutical injustice occurs when platform taxonomies and analytics categories lack the conceptual tools to render students' lived realities—such as intermittent connectivity, caregiving responsibilities, and multilingual code-meshing—visible to the system. This can lead algorithms (and instructors who rely on them) to misinterpret struggle as apathy or low ability (Slade & Prinsloo, 2013; Williamson, 2017). Empirical research on algorithmic bias and "at-risk" prediction reinforces these concerns: features may easily proxy sensitive attributes, training data may encode historical inequities, and error rates or action policies can disproportionately concentrate harm in specific subgroups (Baker & Hawn, 2022; Lakkaraju, Aguiar, Bhanpuri, Miller, & Ghani, 2015). Furthermore, even accurate predictions can trigger detrimental interventions—such as over-monitoring, stigmatizing emails, or restrictive pathway steering—if institutions adopt action policies that prioritise control over care (Kizilcec & Lee, 2020/2022). Learning-analytics research indicates how clustering and personalisation reinforce categories by stabilising certain behavioural norms as indicators of "good learning," which, in turn, legitimises deficit framings of students who do not conform to those norms (Perrotta & Williamson, 2018). From the perspectives of Fricker (2007) and Dotson (2014), addressing these harms necessitates credibility calibration (triangulating algorithmic signals with human judgement and community reputation), hermeneutical expansion (co-creating taxonomies and success indicators with students), and contestability and repair (establishing clear routes to challenge classifications and update models). Methodologically oriented proposals echo this sentiment: fairness-by-design and life-cycle impact assessment should evaluate not only statistical parity but also educational value and the costs of false positives/negatives for different learners (Liu, Dean, Rolf, Simchowitz, & Hardt, 2023; Baker & Hawn, 2022). Absent such measures, data-driven decision-making can silently reproduce epistemic silencing under the guise of precision and objectivity (Benjamin, 2019; Noble, 2018).

Empirical evidence further illustrates the emergence of epistemic injustices within AI-mediated open and distance learning (ODL) environments. For instance, Kizilcec and Lee (2022) demonstrate that automated forum moderation and engagement models disproportionately down-rank posts authored in non-standard English, thereby diminishing the visibility and perceived credibility of multilingual learners. Additionally, research on early alert systems in online community college settings reveals that predictive algorithms often misclassify students with unstable internet access or caregiving responsibilities as exhibiting "low engagement." This misclassification triggers unnecessary risk notifications, undermining the autonomy and agency of these learners (Lakkaraju et al., 2015; Gardner & Brooks, 2018). In the context of large Massive Open Online Courses (MOOCs), Hansen and Reich (2015) found that participation and recommendation algorithms tend to amplify the visibility of already advantaged learners, illustrating how algorithmic curation can reproduce testimonial and hermeneutical inequities on a large scale. These cases underscore the necessity of the CCC criteria—credibility, comprehensibility, and control—to inform the redesign of AI-augmented ODL systems aimed at fostering more equitable learner participation.

5.3 Marginalisation, Opportunity, and the Conditions of Equitable Participation

The broader ODL literature situates these epistemic risks within persistent structural patterns. "Participation and completion in large-scale online contexts often track socioeconomic advantage, with gaps by social class, geography, and prior preparation" (Hansen & Reich, 2015; Kizilcec, Davis, & Cohen, 2017). "Community-college and STEM studies show that online modality interacts with background factors—prior GPA, first-generation status, and prior online experience—to shape outcomes, indicating cumulative advantage unless scaffolds are present" (Hachey, Conway, & Wladis, 2015; Wladis, Hachey, & Conway, 2015). "Retention reviews consistently identify social isolation, time scarcity, inconsistent connectivity/devices, limited instructor presence, and weaker belonging as drivers of attrition—conditions that disproportionately burden low-income, rural, refugee, and caregiving students" (Bawa, 2016; Hart, 2012; Brunton & Buckley, 2019). "Qualitative work with first-in-family learners captures the double reality of ODL: flexibility and open entry expand opportunity, but confidence, navigation, and life-role conflicts remain enduring challenges unless institutions redesign support and assessment with local contexts in mind" (Stone, O'Shea, May, Delahunty, & Partington, 2016). "Importantly, there is nothing inevitable about weaker returns: distance higher education can yield strong labor-market outcomes where provision is adapted to constraints and support is robust" (Wang, 2023). "For AI-mediated ODL, these findings imply that **who benefits** depends less on access to algorithms than on whether systems are governed to recognize diverse ways of knowing and to expand—not narrow—the interpretive and participatory resources available to learners" (Freire, 2000; Slade & Prinsloo, 2013; Veletsianos, 2020).

5.4 Relational and Contextual Agency: Beyond Functional Independence

Agency in this setting is not simply 'self-management' within pre-set tracks, but a relational capacity exercised with and through others in concrete socio-technical arrangements. The social-cognitive account foregrounds intentionality, forethought, self-reactiveness, and self-reflection; these can indeed be supported by adaptive scaffolds, multimodal materials, and diagnostic feedback (Bandura, 2001, 2006; Aleven et al., 2017). Yet the relational view insists that agency also depends on recognition (being treated as a credible knower), voice (having channels to shape meaning), and power (holding decision rights over data practices and learning paths)—all emergent properties of the learning ecology, not attributes a student brings alone (Biesta, 2013; Tait, 2013; Wenger, 1998). In practice, algorithmic designs that pre-empt exploration through aggressive nudging, that hide ranking criteria for posts or submissions, or that confine 'choice' to a narrow menu of system-preferred options shift control away from learners and suppress the dialogical encounters through which subjectivity forms (Freire, 2000; Reich, 2020; Williamson, 2017). Conversely, platforms that surface model rationales in comprehensible terms, invite students to set goals that condition recommendations, permit opting out of certain inferences without penalty, and institutionalise spaces for community deliberation about platform rules create the conditions under which agency can grow (Slade & Prinsloo, 2013; Veletsianos, 2020). The empirical lesson from large-scale intervention studies is consistent: absent structural changes to presence, community, assessment, and resource support, predictive targeting alone does not deliver robust improvements, which suggests that capability-enhancing affordances must be designed into the whole system rather than bolted onto it (Kizilcec et al., 2020; Bawa, 2016; Hart, 2012).

6. Bringing the Strands Together: A Justice-Oriented Reading of AI-mediated ODL

Bringing together these arguments reveals a discernible pattern. Firstly, the “success” of personalisation cannot be assessed solely by prediction metrics; it must also be evaluated against dialogical and justice criteria: Does the system expand decision latitude, support reciprocal meaning-making, and value diverse knowledge practices (Freire, 2000; Biesta, 2013; Perrotta & Williamson, 2018)? Secondly, predictive and classification systems must be examined as epistemic infrastructures: Where do credibility deficits arise (testimonial injustice)? Where do categories fail to render experiences legible (hermeneutical injustice)? How are the resulting actions governed, and by whom (Fricker, 2007; Dotson, 2014; Slade & Prinsloo, 2013)? Thirdly, given that marginalisation in ODL is shaped by material and social conditions, equity necessitates the redistribution of not only content and dashboards but also decision rights, interpretive resources, and channels for repair (Bawa, 2016; Hart, 2012; Stone et al., 2016). Operationally, the composite “Credibility–Comprehensibility–Control” triad provides a practical framework: Credibility examines whether systems avoid credibility deficits and amplify marginalised voices; Comprehensibility investigates whether categories, features, and thresholds are understandable and open to revision; Control assesses whether learners have genuine choices over data and pathways, including the right to refuse certain inferences (Benjamin, 2019; Noble, 2018; Slade & Prinsloo, 2013). Where these conditions are met, AI can complement educator judgement and community practice to support culturally situated learning; where they are not, automation tends to individualise structural barriers, normalise deficit framings, and silence dissent (Selwyn, 2016; Veletsianos, 2020; Williamson, 2017).

Thus, the analysis indicates that AI-augmented ODL will foster learner agency and epistemic justice only when prediction and personalisation are embedded in dialogical pedagogy, governed through participatory and repair-oriented processes, and evaluated based on their effects on recognition, voice, and power for diverse learners. The next section translates these findings into concrete implications for practice and policy in AI-mediated ODL.

7. Proposed Critical Framework

This section translates the analysis into a practical framework for designing, governing, and evaluating AI-mediated ODL. The framework centres on three mutually reinforcing dimensions—recognition, voice, and power—and couples them with three design checkpoints—credibility, comprehensibility, and control (CCC). Taken together, these elements convert high-level values into day-to-day decisions about platforms, pedagogy, and institutional policy. Rather than presenting a diagram, the model is expressed here in prose so it can be copied directly into the article and adapted to varied contexts.

Recognition is the foundation. It requires that students, especially those historically marginalised, are treated as credible knowers and that their cultural and linguistic repertoires are legitimate sources of academic meaning. In AI-mediated environments, this means broadening the categories and data that systems use to “see” learners, calibrating automated flags with human review, and revising assessment rubrics so multilingual expression and community-rooted examples are not penalised. Recognition is visible when false-positive risk flags fall for targeted groups, when non-dominant discourse is routinely accepted as evidence of learning, and when peer reputation or endorsement features bring forward contributions that would otherwise be downranked. Recognition makes good on the ethical demand to value lived experience, but it also improves instructional precision by reducing systematic misreadings (Fricker, 2007).

Voice describes students’ opportunities to set goals, shape criteria, and co-create meaning with others. AI can widen voice when recommendations are conditioned on student-stated goals; when prompts invite local examples; when community annotation, peer review, and reflective writing are treated as first-class learning activities; and when student-contributed exemplars regularly enter the shared library used by recommenders. Voice can be monitored through the share of tasks with student-set goals, adoption rates of student exemplars, and breadth of participation across subgroups. Crucially, voice is dialogic, not merely a menu of clicks: systems should support sustained exchanges that help learners test interpretations and make judgments, rather than substituting engagement scores for genuine conversation (Freire, 2000).

Power concerns decision rights over data and learning pathways. Students should have granular consent over what data feed which features, access to their data and explanations for recommendations, and the right to refuse certain inferences without penalty. They should be able to appeal classifications and have those appeals resolved within set timelines. At the institutional level, power means formal roles for students in platform governance (e.g., a standing council with authority to approve consequential features or mandate rollbacks) and clear accountability for how model changes are decided and documented. Evidence of power includes no-penalty opt-outs, timely resolution of appeals, and demonstrable policy changes initiated by student representatives.

To operationalise the three dimensions, the framework uses the CCC checkpoints as gatekeepers for any AI feature or workflow. Credibility asks whether a feature avoids credibility deficits and elevates marginalised voices; if not, it is revised or withheld. Comprehensibility requires that categories, features, thresholds, and rationales are legible to staff and students and are open to revision; opacity is treated as a design flaw to be fixed, not a trade secret to be tolerated. Control ensures learners have genuine choices about data practices and pathways, including alternatives to any recommended track. A feature that fails any checkpoint does not proceed to deployment until corrected. These checkpoints make abstract values actionable in procurement, configuration, and course-level design (Slade & Prinsloo, 2013).

Practical applications follow from the dimensions and checkpoints. In course and assessment design, instructors pair automated feedback with brief teacher commentary; use dialogic protocols to scaffold discussion; accept multimodal submissions (text, audio, video); and publish rubrics that recognise local knowledge and multilingual expression. In analytics and early-alert pipelines, institutions define error costs before modelling, run subgroup analyses, and require a human in the loop for any action with academic or disciplinary consequences; alerts become prompts for dialogue rather than triggers for automatic penalties. In recommendation systems, students can set or adjust goals; "why this was recommended" explanations are always shown; a "show alternatives" control is provided; and rankers are periodically audited for downranking of non-dominant discourse. For student data rights, dashboards disclose what data feed which inferences, granular consent toggles are available, data are portable, and appeals have service-level targets (acknowledge within a fixed window; resolve by a set deadline). Continuous improvement is built in: each term, programmes run compact fairness/impact reviews that combine subgroup metrics (e.g., false-positive rates, participation breadth, pathway diversity) with qualitative evidence from student panels and focus groups; a short audit summary and remediation plan are published.

Policy and institutional considerations anchor the framework. Governance: establish a cross-stakeholder council (students, faculty, analytics/IT, accessibility, ethics) with authority to approve high-impact features, set threshold policies, and mandate rollbacks where harms appear. Accountability: require *ex ante* impact assessments for new features and *ex post* audits every term; create an appeals ombud and set measurable targets for time-to-resolution and disparity reduction; publish change logs. Capacity and workload: fund instructor presence and dialogic work (the human layer that AI cannot replace), expand advising and accessibility services, and train staff and student representatives in participatory design, audit methods, and basic model literacy. Procurement: include clauses that guarantee transparency (disclosed features/thresholds), data portability, audit access, and a safe rollback mechanism; prefer vendors that support goal-conditioned recommendations and built-in contestability.

Implementation pathway proceeds in three stages. Pilot the framework with a small number of high-enrolment ODL courses that serve diverse learners. Before deployment, apply the CCC checkpoints to each AI feature; during delivery, capture both metrics and narratives; after delivery, run a fairness/impact review with student participation and publish a change log. Scale by integrating the checkpoints into standard operating procedures (curriculum approvals, analytics governance) and by templating the dialogic design patterns that proved effective (e.g., structured peer review, reflective goal-setting). Sustain by embedding the governance council in statute or policy, renewing its membership annually, and scheduling regular audits tied to budgeting and procurement cycles so incentives align with equity and learning.

Finally, the framework's theory linkage is explicit: recognition, voice, and power translate the justice claims of epistemic-injustice scholarship into design demands (Fricker, 2007); dialogic pedagogy specifies how AI should augment, not replace, human judgment and shared inquiry (Freire, 2000); and the CCC checkpoints provide a practical ethical layer for the learning-analytics lifecycle where data practices and interventions are planned, justified, and revised (Slade & Prinsloo, 2013). In sum, the model offers a portable, testable way to align AI-mediated ODL with democratic, justice-oriented education: features are built and evaluated to expand students' real freedoms to be recognised, to speak, and to decide.

To demonstrate the practical application of the CCC framework, consider a prevalent artificial intelligence feature in ODL: an early-alert system designed to predict which students may be "at risk." The application of the *Credibility* checkpoint reveals that multilingual and low-bandwidth learners are disproportionately identified as at risk, as the model proxies engagement through uninterrupted log-ins and standard English writing patterns. This finding necessitates a review of the features used and the error rates associated with various subgroups. Through the *Comprehensibility* aspect, the institution ensures that both students and instructors receive clear explanations of the reasons behind the activation of a risk flag, accompanied by transparent thresholds and illustrative examples. Finally, under the *Control* dimension, students are afforded the opportunity to contest inaccurate flags, opt out of certain forms of behavioural surveillance, and collaboratively design alternative indicators that take into account local constraints, such as intermittent connectivity or caregiving responsibilities. This systematic application illustrates how the CCC framework can transform a high-stakes predictive system from a source of epistemic harm into one that promotes autonomy, recognition, and equitable participation.

8. Implications for Practice and Policy

Here, we translate the critical framework into concrete actions for individuals who build and run AI-mediated ODL. The focus is practical: transforming recognition, voice, and power—validated by credibility, comprehensibility, and control—into everyday decisions regarding courses, analytics, and governance. This section is organised into two parts: first, recommendations for universities and instructional designers; second, strategies for promoting cultural relevance and student dignity, ensuring that AI augments human judgment rather than replacing it.

8.1 Recommendations for Universities and Instructional Designers

ODL institutions often operate with limited institutional capacity, uneven data governance infrastructures, and constrained human and financial resources, all of which shape the forms of AI oversight and pedagogical redesign that are realistically implementable. Acknowledging these constraints ensures that the recommendations remain actionable by emphasising scalable practices, such as phased audits, human-in-the-loop triage, and simplified governance mechanisms, that institutions with varying resource levels can adopt.

Universities should embed three non-negotiable design checkpoints—credibility, comprehensibility, and control—into every AI feature used in ODL. Credibility ensures that no group is systematically disadvantaged by flags, rankings, or grading assistance; comprehensibility mandates that categories, features, thresholds, and "why recommended" rationales are transparent and subject to revision; and control guarantees learners genuine choices over data use and learning pathways, including the right to refuse certain inferences without penalty. These checkpoints should be incorporated into course approval processes, analytics governance, and vendor contracts to guide everyday decisions rather than remain as mere principles on paper. Instructors and learning designers should aim to facilitate dialogue rather than merely provide dashboards: they should pair automated feedback with brief teacher commentary; stage structured discussions and community annotations to allow students to test interpretations publicly; and align automated nudges with clearly stated learning intentions that students have helped establish. Human judgment must remain involved in any actions with academic or disciplinary consequences; risk alerts should initiate a brief triage conversation rather than result in automatic penalties. Institutions should also conduct termly fairness and impact reviews that combine subgroup metrics with student focus groups, publish concise summaries of findings, and maintain a public change log of any adjustments made. Finally, governance must be robust: a standing council comprising students, faculty, analytics/IT specialists, accessibility experts, and ethicists should approve consequential features, mandate rollbacks when harms are identified, and set timelines for necessary repairs. Building capacity is an integral part of the policy work—fund instructor presence, train staff and student representatives in participatory design and audit methods, and allocate resources for advising and accessibility to address structural barriers rather than individualising solutions.

8.2 Strategies for Promoting Cultural Relevance and Student Dignity

Cultural relevance begins with recognition: treat students' linguistic repertoires and community knowledge as legitimate academic resources. Revise rubrics to ensure that multilingual and locally grounded examples earn credit; accept multimodal submissions—text, audio, video—when learning goals permit; and adjust moderation so that non-dominant discourse is not downranked by default. Voice is strengthened when students help shape prompts, exemplars, and success criteria, and when their stated goals inform recommendations throughout a module. Simple practices—goal-setting reflections, student-contributed exemplars that enrich the recommender pool, and reflective summaries that connect course ideas to local contexts—transform voice from tokenism to shared authorship. Power is safeguarded by clear data rights: clarify which data inform which inferences, provide granular consent controls, offer explanations for recommendations, and establish an appeals process with service-level targets for acknowledgment and resolution. Dignity is also material: invest in connectivity supports, device lending, and flexible pacing so the system does not misinterpret intermittent access or caregiving responsibilities as apathy. Instructor presence is crucial—short video check-ins, timely responses, and peer mentoring schemes can counter isolation and assist learners in translating automated feedback into meaningful next steps. Throughout all of this, the test is straightforward: do students feel recognised, heard, and able to influence decisions that affect their learning?

9. Conclusion and Recommendations for Further Studies

This article argues that AI in ODL will advance learning only when it enhances epistemic relations—who is recognised as a credible knower, who has a voice in meaning-making, and who holds power over data and pathways. The analysis shows that uncritical personalisation can harden narrow proxies of success and reproduce credibility deficits, while dialogic design, participatory governance, and transparent models can expand agency and equitable participation. The proposed framework translates these insights into practice through the twin anchors of the three dimensions—recognition, voice, and power—and the CCC checkpoints—credibility, comprehensibility, and control—so that pedagogy, analytics, and policy pull in the same direction.

Reimagining ODL as a space of equitable knowledge construction means shifting from "better targeting at scale" to co-created learning at scale. Platforms should help students bring local knowledge into the academic conversation, understand and question algorithmic decisions, and make binding choices about their data and learning paths. Universities should treat AI not as a routing mechanism that silently tracks learners into predefined lanes, but as a dialogic service that opens options, surfaces alternatives, and invites reasoned disagreement.

Further work should validate short measures of recognition, voice, and power, and link them to learning and retention; conduct design trials that compare goal-conditioned recommendations with default rankers and test conversational triage versus automated action when risk flags are raised; and study governance in practice to learn what council structures, timelines, and data access produce timely, legitimate decisions. Multi-site studies across regions can demonstrate how the framework travels and what adaptations are needed for language, policy, and infrastructure differences. In brief, the path forward is to govern AI as a socio-technical choice with ethical guardrails, design for dialogue and dignity, and share power over data and decisions so diverse learners can build knowledge with agency.

References

Aleven, V., McLaughlin, E. A., Glenn, R. A., & Koedinger, K. R. (2017). Instruction based on adaptive learning technologies. In R. E. Mayer & P. A. Alexander (Eds.), *Handbook of research on learning and instruction* (2nd ed., pp. 522-560). Routledge.

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education, 32*(4), 901-902. https://doi.org/10.1007/s40593-021-00285-9

Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology, 52*, 1-26. https://doi.org/10.1146/annurev.psych.52.1.1

Bandura, A. (2006). Toward a psychology of human agency. *Perspectives on Psychological Science, 1*(2), 164-180. https://doi.org/10.1111/j.1745-6916.2006.00011.x

Bawa, P. (2016). Retention in online courses: Exploring issues and solutions—A literature review. *SAGE Open, 6*(1), 1-11. https://doi.org/10.1177/2158244015621777

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity.

Biesta, G. (2013). *The beautiful risk of education*. Paradigm Publishers.

Brunton, J., & Buckley, P. (2019). Duelling identities in refugees learning through open, online, and distance learning. *Open Learning, 34*(3), 282-298. Retrieved from https://files.eric.ed.gov/fulltext/EJ1251308.pdf

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (pp. 77-91).

PMLR. Retrieved from https://proceedings.mlr.press/v81/buolamwini18a.html

Chinn, P. L., & Kramer, M. K. (2018). *Knowledge development in nursing: Theory and process* (10th ed.). Elsevier.

Dotson, K. (2014). Conceptualizing epistemic oppression. *Social Epistemology, 28*(2), 115-138. https://doi.org/10.1080/02691728.2013.782585

Emirbayer, M., & Mische, A. (1998). What is agency? *American Journal of Sociology, 103*(4), 962-1023. https://doi.org/10.1086/231294

Fraser, N. (2009). *Scales of justice: Reimagining political space in a globalizing world*. Columbia University Press.

Freire, P. (2000). *Pedagogy of the oppressed* (30th anniversary ed.). Bloomsbury/Continuum.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. Retrieved from https://academic.oup.com/book/32817

Gardner, J., & Brooks, C. (2018). Student success prediction in MOOCs. *User Modeling and User-Adapted Interaction, 28*(2), 127-203. https://doi.org/10.1007/s11257-018-9203-z

Giroux, H. A. (2011). *On critical pedagogy*. Continuum.

Hachey, A. C., Conway, K. M., & Wladis, C. (2015). Prior online course experience and GPA as predictors of online STEM course outcomes. *The Internet and Higher Education, 25*, 11-17. https://doi.org/10.1016/j.iheduc.2014.10.003

Hansen, J. D., & Reich, J. (2015). Democratizing education? Examining access and usage patterns in MOOCs. *Science, 350*(6265), 1245-1248.

Hart, C. (2012). Factors associated with student persistence in an online program of study: A review of the literature. *Journal of Asynchronous Learning Networks, 16*(1), 19-37. Retrieved from https://eric.ed.gov/?id=EJ976760

Hooks, B. (1994). *Teaching to transgress: Education as the practice of freedom*. Routledge.

Kidd, I. J., Medina, J., & Pohlhaus, G. (Eds.). (2017). *The Routledge handbook of epistemic injustice*. Routledge. https://doi.org/10.4324/9781315212043

Kizilcec, R. F., & Lee, H. (2022). Algorithmic fairness in education. In C. Holmes et al. (Eds.), *The ethics of artificial intelligence in education* (pp. 171-192). Routledge. Retrieved from https://arxiv.org/abs/2007.05443

Kizilcec, R. F., Davis, G. M., & Cohen, G. L. (2017). Towards equal opportunities in MOOCs: Affirmation reduces gender and social-class achievement gaps in China. *Working paper*. Retrieved from https://rene.kizilcec.com/wp-content/uploads/2017/02/kizilcec2017towards.pdf

Kizilcec, R. F., Davis, G. M., Lakhani, K. R., & Cohen, G. L. (2020). Scaling up behavioural science interventions in online education. *Proceedings of the National Academy of Sciences, 117*(26), 14900-14905. https://doi.org/10.1073/pnas.1921417117

Lakkaraju, H., Aguiar, E., Bhanpuri, N., Miller, D., & Ghani, R. (2015). A machine learning framework to identify students at risk of adverse academic outcomes. *Proceedings of KDD '15* (pp. 1909-1918). https://doi.org/10.1145/2783258.2788620

Lakkaraju, H., Aguiar, E., Bhanpuri, N., Miller, D., & Ghani, R. (2015). A machine learning framework to identify students at risk of adverse academic outcomes. In *Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (pp. 1909-1918). https://doi.org/10.1145/2783258.2788620

Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., & Hardt, M. (2023). Reimagining the machine learning life cycle to improve educational outcomes. *Proceedings of the National Academy of Sciences, 120*(20), e2204781120. https://doi.org/10.1073/pnas.2204781120

Lynham, S. A. (2002). The general method of theory-building research in applied disciplines. *Advances in Developing Human Resources, 4*(3), 221-241. https://doi.org/10.1177/1523422302043002

Medina, J. (2013). *The epistemology of resistance: Gender and racial oppression, epistemic injustice, and resistant imaginations*. Oxford University Press.

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press. Retrieved from https://nyupress.org/9781479837243/algorithms-of-oppression/

Pane, J. F., Steiner, E. D., Baird, M. D., & Hamilton, L. S. (2015). *Continued progress: Promising evidence on* *personalized learning*. RAND Corporation. Retrieved from https://www.rand.org/pubs/research_reports/RR1365.html

Perrotta, C., & Williamson, B. (2018). The social life of learning analytics: Cluster analysis and the "performance" of algorithmic education. *Learning, Media and Technology, 43*(1), 3-16. https://doi.org/10.1080/17439884.2016.1182927

Reich, J. (2020). *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press.

Selwyn, N. (2016). *Is technology good for education?* Polity.

Sen, A. (1999). *Development as freedom*. Oxford University Press.

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist, 57*(10), 1510-1529. https://doi.org/10.1177/0002764213479366

Stone, C., O'Shea, S., May, J., Delahunty, J., & Partington, Z. (2016). Opportunity through online learning: Experiences of first-in-family students in online open-entry higher education. *Australian Journal of Adult Learning, 56*(2), 146-169. Retrieved from https://eric.ed.gov/?id=EJ1107574

Tait, A. (2013). Distance and e-learning, social justice, and development: The relevance of capability approaches to the mission of open universities. *The International Review of Research in Open and Distributed Learning, 14*(4), 1-18. https://doi.org/10.19173/irrodl.v14i4.1526

UNESCO. (2019). *Beijing consensus on artificial intelligence and education*. Retrieved from https://unesdoc.unesco.org/ark:/48223/pf0000368303

UNESCO. (2021). *AI and education: Guidance for policy-makers*. Retrieved from https://unesdoc.unesco.org/ark:/48223/pf0000376709

Veletsianos, G. (2020). *Learning online: The student experience*. Johns Hopkins University Press.

Walker, L. O., & Avant, K. C. (2019). *Strategies for theory construction in nursing* (6th ed.). Pearson.

Wang, L. (2023). Who gets the highest return to distance higher education? Evidence from China. *The International Review of Research in Open and Distributed Learning, 24*(4), 19-37. https://doi.org/10.19173/irrodl.v24i4.7143

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.

Williamson, B. (2017). *Big data in education: The digital future of learning, policy and practice*. SAGE.

Wladis, C., Hachey, A. C., & Conway, K. (2015). The online STEM classroom—Who succeeds? *Community College Review, 43*(2), 142-164. https://doi.org/10.1177/0091552115571729

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—Where are the educators? *International Journal of Educational Technology in Higher Education, 16*, 39. https://doi.org/10.1186/s41239-019-0171-0

Acknowledgments

We gratefully acknowledge the University of South Africa (UNISA) for providing financial support toward the Article Processing Charges (APCs) for this publication.

Authors contributions

IBO and SAN were responsible for study design and conceptualization. IBO and SAN drafted the manuscript and IBO revised it. All authors read and approved the final manuscript.

Funding

This study did not receive any external funding.

Competing interests

The authors declare no conflicts of interest.

Informed consent

Not applicable

Ethics approval

Ethical approval is not applicable as this is purely argumentative in nature and does not involve human participants, human data, or any form of human interaction.

Provenance and peer review

Not commissioned; externally double-blind peer reviewed.

Data availability statement

Data availability is not applicable to this study, as it is theoretical in nature and did not involve the generation or analysis of empirical data.

Data sharing statement

No additional data are available.