

Algorithmic Sovereignty in Practice: A West African Model for Regulating High-Impact AI Systems

Emmanuel Nwonye

ABSTRACT

Artificial intelligence is increasingly shaping public and private decision-making across West Africa, particularly in sectors such as Education, healthcare, financial services, recruitment, agriculture, elections, border management, and public administration. However, regulatory responses in the region remain fragmented and largely centred on data protection, innovation policy, and broad ethical commitments. This paper argues that the most urgent AI regulatory gap in West Africa is not simply the absence of AI-specific legislation, but the absence of a regional mechanism to identify, document, and respond to AI harms after deployment. Many AI systems used in the region are imported, privately controlled, technically opaque, and trained outside the linguistic, social, economic, and institutional realities of West Africa. As a result, citizens may be subjected to automated decisions without clear evidence of local suitability testing, human accountability, appeal pathways, or cross-border regulatory oversight. To address this gap, the paper proposes the West African AI Harm Passport, a Pan-African policy model for governing high-impact AI systems deployed in Nigeria and the wider ECOWAS region. The proposed framework would require developers, vendors, and deploying institutions to provide verifiable evidence of system identity, data provenance, sectoral risk classification, contextual bias testing, local language and accessibility assessment, human oversight, complaint mechanisms, and post-deployment incident reporting. Unlike conventional AI ethics frameworks, the Harm Passport shifts regulation from abstract principles to traceable accountability at the point of social impact. The paper concludes that such a model could position West Africa as a global contributor to context-sensitive AI governance by transforming the region from a passive recipient of imported AI systems into an active regulator of algorithmic Power.

Keywords: Algorithmic sovereignty; AI regulation; high-impact AI systems; AI accountability; West African governance;

Emmanuel Nwonye, emmanuel@breatha.org, Cofounder/Chief Technology Officer, Breatha LTD/GTE, Port Harcourt, Nigeria

The Nigerian Journal of Business and Social Sciences, Volume 13, 2026 (Special Issue)
A Journal of the Faculty of Social Sciences, University of Lagos, Akoka, Lagos, Nigeria
© 2026

1. Introduction

Artificial intelligence is no longer a distant policy concern for West Africa. It is already entering the ordinary systems through which people learn, borrow, work, trade, receive health information, access public services, and participate in civic life (African Union [AU], 2024; Bakare, 2025; Ministry of Communications and Digitalisation, 2022; World Bank, 2025). For many citizens, the first encounter with AI will not look like a robot or a futuristic machine. It may appear as a credit score that denies a loan, a recruitment filter that removes a job application, a school platform that ranks performance, a health tool that interprets symptoms, or a public communication system that shapes what people believe during elections (World Bank, 2025; International Development Research Centre [IDRC], 2025; Bakare, 2025). These decisions may be technical in design, but their consequences are deeply human.

This is why AI governance in West Africa must be framed not only as a question of innovation but also as One of Power, accountability, and sovereignty. Nigeria has taken important steps through the Nigeria Data Protection Act 2023, which established the Nigeria Data Protection Commission (NDPC) and strengthened the legal basis for protecting personal data (Nigeria Data Protection Commission [NDPC], 2023). Nigeria's National Artificial Intelligence Strategy also identifies AI as a tool for economic growth, social inclusion, technological leadership, responsible development, and robust governance (Federal Ministry of Communications, Innovation and Digital Economy [FMCIDE], 2025). At the continental level, the African Union's (AU) Continental Artificial Intelligence Strategy calls for an Africa-centred approach to AI that promotes ethical, responsible, and equitable development while recognising risks such as bias, discrimination, privacy abuse, surveillance, and liability gaps (AU, 2024). More recently, calls within the Economic Community of West African States (ECOWAS) Parliament for a regional legal framework on AI show that West Africa is beginning to recognise the need for coordinated action beyond national borders (News Agency of Nigeria, 2025).

However, these developments also expose a deeper gap. Data protection laws can regulate how personal information is collected, processed, and shared. However, they do not fully answer what should happen when an AI system produces harmful outcomes even after data rules have been observed. An automated decision may be privacy-compliant and still be unfair (Kasirzadeh & Clifford, 2021). A system may carry an ethics statement and still fail to understand local languages, informal livelihoods, rural infrastructure, gendered realities, disability access, or the lived conditions of West African communities. The central regulatory problem, therefore, is not merely that West Africa lacks AI laws. It is that the region lacks a practical mechanism for testing, tracing, challenging, and correcting high-impact AI systems before and after they affect people's lives (Bakare, 2025; National Institute of Standards and Technology [NIST], 2023; Organization for Economic Co-operation and Development [OECD], 2024; United Nations Educational, Scientific and Cultural Organization [UNESCO], 2021).

This paper argues that algorithmic sovereignty must move from aspiration to practice. It does not require West Africa to reject foreign-built AI systems or build every technology locally. Rather, it requires that any high-impact AI system deployed in the region should be answerable to West African evidence, West African institutions, and West African citizens. To address this need, the paper proposes a West African AI Harm Passport: a regional accountability model requiring high-impact AI systems to disclose their identity, purpose, data provenance, local suitability testing, human oversight arrangements, complaint pathways, and post-deployment harm reports. By shifting the focus from data protection alone to harm protection, the paper offers a policy response that can help Nigeria and the wider ECOWAS region regulate algorithmic Power without shutting down innovation.

2. Conceptual Framework: From Data Protection to Harm Protection

AI regulation often begins with data protection because data is the raw material from which many AI systems learn, infer, rank, predict, and decide (NIST, 2023; OECD, 2024). This starting point is necessary, especially in societies where citizens already face risks of surveillance, weak consent practices, identity misuse, and commercial exploitation of personal information. However, data protection alone is too narrow for the kind of Power AI now exercises. It asks important questions about how data is collected, stored, processed, and shared. However, it does not always ask whether an AI system is socially fit, locally accurate, contestable, or fair in its real-world consequences. Scholars have shown that automated inferences and data-driven decisions can create accountability gaps even where personal data is lawfully processed, especially when individuals have limited control over how conclusions about them are produced or used (Barocas & Selbst, 2016; Wachter & Mittelstadt, 2019).

This distinction matters for West Africa. A bank may process personal data lawfully and still use an automated scoring system that wrongly excludes informal workers because their income patterns do not fit imported credit models. A health platform may protect patient information and still misread symptoms because its training data does not reflect local disease burdens, skin tones, languages, or care-seeking behaviours. A recruitment tool may comply with privacy rules and still disadvantage applicants from certain universities, regions, speech patterns, or employment histories. Comparable risks have been documented in employment screening, healthcare allocation, facial analysis, and speech recognition systems, where automated tools produced unequal outcomes because of biased historical data, weak proxy variables, or unrepresentative testing environments (Buolamwini & Gebru, 2018; Dastin, 2018; Koenecke et al., 2020; Obermeyer et al., 2019). In such cases, the harm is not only a privacy breach. It is exclusion, misclassification, denial of opportunity, weakened dignity, and the loss of a meaningful chance to challenge the decision.

Global AI governance already recognises the need to move beyond privacy. UNESCO's Recommendation on the Ethics of Artificial Intelligence places human rights, dignity, transparency, fairness, and human oversight at the centre of AI governance. At the same time, the Organisation for Economic Co-operation and Development (OECD) AI Principles stress trustworthy AI, human-centred values, transparency, robustness, safety, and accountability (OECD, 2024; UNESCO, 2021). The European Union's (EU) AI Act also reflects this shift by adopting a risk-based approach, in which systems that may pose adverse effects on safety or fundamental rights are treated as high-risk and subject to stricter obligations (European Commission, 2025, 2026). These frameworks are useful, but they cannot simply be copied into West Africa without adaptation. They were built within institutional environments with stronger enforcement capacity, larger technical bureaucracies, and longer traditions of administrative redress than many West African states currently possess.

For this paper, harm protection refers to a regulatory approach that focuses on what AI systems do to people after they leave the laboratory, policy document, or vendor presentation. It asks whether affected persons can understand that AI was involved, identify who is responsible, question the outcome, seek correction, and trigger regulatory attention when harm occurs. This does not replace data protection. Rather, it builds on it. Data protection safeguards the conditions under which information is handled, while harm protection safeguards against the consequences of automated Power.

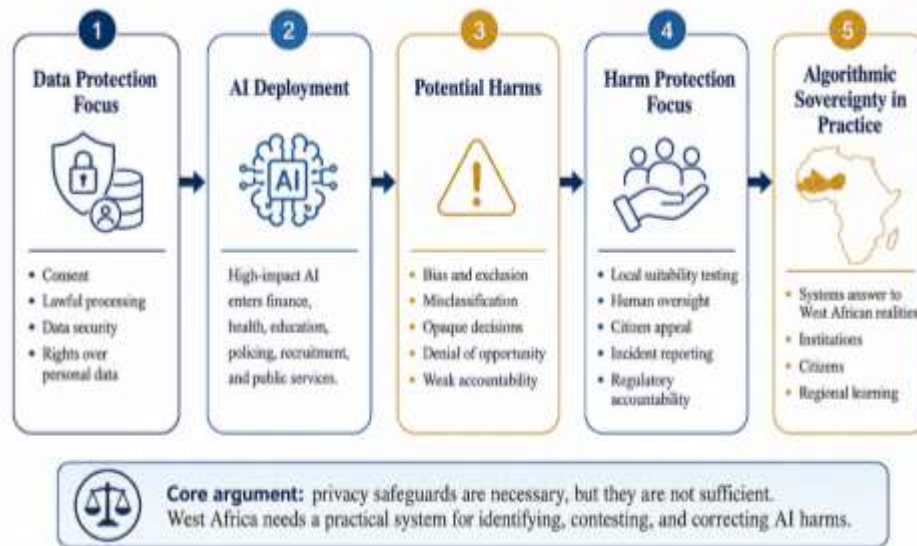


Figure 1: From Data Protection to Harm Protection

The concept of algorithmic sovereignty is therefore used here in a practical sense. It is not a call for technological isolation or a rejection of foreign innovation. It is the capacity of West African states and communities to set the terms under which high-impact AI systems operate within their societies. In practice, sovereignty means that a system built elsewhere must still answer locally. It must be tested against West African realities, supervised by accountable institutions, and open to challenge by the people whose lives it affects. This is the foundation for the West African AI Harm Passport proposed later in this paper.

3. Methodology

This paper adopts a qualitative policy analysis approach, supported by illustrative case analyses. Its purpose is not to measure public attitudes toward AI or test the performance of a specific AI system. Rather, it examines the regulatory environment in which high-impact AI systems are likely to operate in Nigeria and the wider West African region, and asks what kind of accountability mechanism would be needed to protect citizens from algorithmic harm. This methodology is appropriate because the study concerns a regulatory and institutional problem, not a statistical or experimental one. The central question is how existing legal and policy frameworks can address the risks posed by high-impact AI systems, especially when such systems are imported, opaque, and deployed across diverse social contexts. A qualitative policy analysis approach enables the paper to examine laws, strategies, governance principles, and institutional arrangements to identify gaps between policy ambition and practical accountability. It is also suitable for developing a normative policy proposal, since the paper seeks not only to describe existing frameworks but also to design a context-sensitive regulatory response for West Africa.

The analysis proceeds in four stages. First, the paper reviews legal and policy instruments relevant to AI governance in Nigeria, West Africa, and Africa more broadly, including the Nigeria Data Protection Act, Nigeria's National Artificial Intelligence Strategy, the AU's Continental Artificial Intelligence Strategy, and emerging calls for ECOWAS-level AI regulation. These instruments are examined to identify what they already address and where gaps remain. Second, the paper draws on international AI governance frameworks, including UNESCO's Recommendation on the Ethics of Artificial Intelligence, the OECD AI Principles, and the EU's risk-based AI regulatory approach. These frameworks are not treated as models for direct transplantation. Instead, they are used as

reference points for understanding global approaches to risk, accountability, human oversight, and post-deployment responsibility.

Third, the paper uses selected real-world cases to show how AI systems can misjudge people when training data, proxy variables, or testing environments do not reflect the populations the system affects. These cases include Amazon's experimental recruitment tool, documented demographic disparities in facial recognition systems, the wrongful arrest of Robert Williams following a false facial recognition match, racial bias in a healthcare allocation algorithm, and disparities in automated speech recognition (Buolamwini & Gebru, 2018; Dastin, 2018; Grother et al., 2019; Koenecke et al., 2020; Obermeyer et al., 2019). These examples are not presented as direct cases from West Africa. Rather, they serve as diagnostic evidence of risk pathways that could become more damaging in contexts with weaker redress systems, limited audit capacity, imported technologies, multilingual populations, and uneven institutional oversight. Finally, the paper uses institutional design analysis to propose the West African AI Harm Passport as a practical regulatory tool for high-impact AI systems. The proposed framework is assessed according to whether it can identify responsible actors, require local suitability testing, create citizen redress pathways, support post-deployment incident reporting, and enable cross-border regulatory learning within ECOWAS.

4. Analysis and Findings

4.1 AI Governance in Nigeria and West Africa: The Emerging Landscape

AI governance in West Africa is no longer starting from zero. Across the region and the continent, there is growing recognition that law, public policy, ethical standards, and institutional oversight must guide the development and use of artificial intelligence. Nigeria is especially important in this conversation because of its population, digital economy, startup ecosystem, and influence within ECOWAS. If Nigeria develops a workable model for governing high-impact AI systems, it could shape broader regional practice.

Table 1: Emerging AI Governance Instruments and Gaps in Nigeria and West Africa

Governance instrument or development	Main contribution to AI governance	Key limitation for high-impact AI regulation
Nigeria Data Protection Act 2023	Establishes the NDPC and strengthens rules on lawful processing, consent, data subject rights, and duties of data controllers and processors (NDPC, 2023).	It is primarily concerned with data processing and does not fully address broader AI harms such as algorithmic exclusion, misclassification, opaque automated decisions, and weak redress.
Nigeria's National Artificial Intelligence Strategy	Links AI to economic competitiveness, social development, inclusion, technological leadership, and responsible governance. It identifies sectors such as health, agriculture, Education, climate action, government services, and productivity as priority areas (FMCIDE, 2025).	It provides strategic direction but does not, by itself, create enforceable duties for high-impact AI vendors, developers, or deploying institutions.
African Union Continental Artificial Intelligence Strategy	Provides a continental framework for Africa-centred AI development, ethical governance, socioeconomic transformation, and responsible innovation (AU, 2024).	It depends on national and regional institutions for implementation and lacks a direct enforcement mechanism for AI systems deployed in specific sectors in West Africa.
Ghana's National AI Strategy	Demonstrates growing regional momentum by promoting responsible, human-centred AI, local talent development, AI adoption in key sectors, and ethical and legal frameworks (Ministry of Communication, Digital Technology and Innovations, 2026).	It remains a national strategy, while many AI systems operate across borders and may carry the same risks into several West African markets.
ECOWAS-level calls for AI regulation	Recent calls within the ECOWAS Parliament indicate growing recognition that West Africa needs coordinated rules for AI and digital transformation (News Agency of Nigeria, 2025).	The region still lacks a practical, harmonised accountability instrument to test, monitor, and respond to high-impact AI harms across member states.

Table 1 above shows that Nigeria, Ghana, the AU, and ECOWAS are all moving toward AI governance, but their efforts remain uneven and incomplete. Nigeria's Data Protection Act provides an important legal foundation, as many AI systems rely on personal and behavioural data. It establishes the NDPC and gives stronger legal force to the protection of personal data, including rules on lawful processing, consent, data subjects' rights, and the duties of data controllers and processors (NDPC, 2023). However, data protection does not fully answer what happens when an AI system classifies, ranks, excludes, predicts, or recommends decisions in ways that harm people.

An AI system may comply with data protection rules and still produce discriminatory or inaccurate outcomes.

Nigeria's National Artificial Intelligence Strategy extends the conversation beyond privacy by linking AI to economic competitiveness, social development, inclusion, technological leadership, and responsible governance (FMCIDE, 2025). The strategy recognises AI as a national development tool. It identifies sectors such as health, agriculture, Education, climate action, government services, and productivity as areas where AI could have transformative value. It also reflects an important shift: Nigeria is not only trying to consume AI but to participate in its design, deployment, and governance. However ambitious, a strategy does not by itself create enforceable duties for high-impact AI vendors, developers, or institutions. It sets direction, but the practical machinery of accountability still has to be built. At the continental level, the AU's Continental Artificial Intelligence Strategy provides a broader political and normative framework. The AU calls for an Africa-centred approach to AI that supports socioeconomic transformation while promoting ethical, responsible, and equitable practices (AU, 2024). This is important because it frames AI not merely as a technical opportunity but as part of Africa's long-term development, cultural, and governance agenda. However, continental strategies depend on national and regional institutions for implementation. Without mechanisms that can translate principles into enforceable obligations, the gap between policy ambition and lived protection may remain wide.

Other West African countries are also moving into the AI policy space. Ghana's National AI Strategy, for example, seeks to develop a responsible, human-centred AI ecosystem, strengthen local talent and research capacity, promote AI adoption in key sectors, and establish ethical and legal frameworks (Ministry of Communication, Digital Technology and Innovations, 2026). This shows that AI governance is becoming a regional concern rather than a single-country issue. Still, these national efforts remain uneven, and AI systems do not respect national borders. A system used in finance, education, recruitment, or identity verification may operate across several West African markets with the same underlying model, hidden assumptions, and risks. This is where ECOWAS becomes central. Recent calls within the ECOWAS Parliament for a legal framework on AI indicate an emerging recognition that West Africa needs coordinated rules for digital transformation and AI deployment (News Agency of Nigeria, 2025). Such coordination is necessary because fragmented national responses can leave citizens exposed, especially when harmful systems outpace regulators. The emerging landscape therefore reveals both progress and weakness. Nigeria, Ghana, the AU, and ECOWAS are all advancing AI governance, but the region still lacks a practical accountability framework for high-impact AI systems. That missing instrument is the space this paper seeks to fill through the proposed West African AI Harm Passport.

4.2. The Critical Regulatory Gap: Imported AI and the Accountability Vacuum

The most serious AI regulatory gap in West Africa is not simply that the region has too few AI laws. The deeper problem is that high-impact AI systems can enter social life faster than public institutions can understand, test, contest, or correct them. This is especially important because many systems likely to affect West African citizens are not designed primarily with West African realities in mind. They are often built by foreign companies, trained on non-local data, hosted on external infrastructure, sold through private vendors, and deployed by banks, schools, hospitals, employers, security agencies, or public institutions that may not fully understand how the systems work.

This creates an accountability vacuum. When an AI system causes harm, responsibility can be scattered among too many actors. The developer may claim it only built the tool. The vendor may claim it only supplied the product. The deploying institution may claim it relied on the system's

technical authority. The regulator may lack access to the model, training data, audit records, or performance reports. The affected citizen may not even know that AI was involved. In this chain, everyone has some connection to the decision, but no One is clearly answerable for the harm.

The first layer of this gap is contextual. AI systems do not make decisions in a social vacuum. They carry assumptions about language, documentation, identity, income, risk, behaviour, and credibility. In West Africa, these assumptions can fail quickly. A credit-scoring system may misunderstand informal income. A recruitment tool may reward employment histories that reflect privilege rather than ability. A health chatbot may fail to recognise local disease patterns or everyday descriptions of symptoms. A facial recognition or identity-verification system may perform unevenly across skin tones, lighting conditions, image quality, and documentation practices. These are not merely technical errors. They can shape who gets access to money, work, care, mobility, Education, public trust, or who gets to live or die.

The second layer is evidential. Current policy discussions often encourage ethical, responsible, and inclusive AI, but affected communities need more than assurance. They need proof. UNESCO's AI ethics recommendation emphasises oversight, impact assessment, auditability, traceability, and accountability across the AI life cycle. At the same time, the AU's Continental AI Strategy calls for AI systems to be adapted to Africa's context and guided by ethics, inclusion, human rights, dignity, and people's well-being (AU, 2024; UNESCO, 2021). These principles are important, but in West Africa, they must be translated into practical requirements. Before a high-impact system is deployed, there should be evidence that it has been tested against local realities, not only against the supplier's market claims.

The third layer is redress. Data protection laws can grant citizens rights over their personal information, but harms from AI are not always reducible to privacy violations. A person may be denied a loan, wrongly flagged as suspicious, excluded from a job shortlist, misclassified by a school platform, or affected by an automated public service decision without any clear appeal route. The EU's AI Act addresses this concern through a risk-based model that imposes obligations on high-risk AI systems and their deployers, including requirements for human oversight, monitoring, and risk reporting (European Commission, 2025). However, West Africa cannot simply import the European model. The region needs a lighter, context-sensitive mechanism that fits its institutional capacity while still making harm visible and contestable.

The fourth layer is regional. AI systems do not stop at national borders. A product tested or deployed in Nigeria may later be used in Ghana, Senegal, Côte d'Ivoire, Sierra Leone, or other ECOWAS states. If harm occurs in One country, other countries should not have to discover the same failure from the beginning. A regional warning and learning system is therefore essential. Without it, West Africa risks repeating harms across borders while each country responds in isolation. The accountability vacuum, then, is not just a legal absence. It is a missing infrastructure of visibility, evidence, responsibility, and repair. For algorithmic sovereignty to become practical, West Africa must be able to ask basic but powerful questions of every high-impact AI system: Who built it? Who deployed it? What was it trained on? Has it been tested here? Who is responsible if it fails? How can citizens challenge it? How will regulators learn from its harms? The West African AI Harm Passport proposed in this paper is designed to answer these questions.

4.3. Evidence from Real Cases: When AI Misjudges People and Context

The need for local testing and harm-based regulation is not theoretical. Several well-known cases show that AI systems can misjudge people when their training data, proxy variables, or testing environments fail to reflect the populations they affect. These cases are not presented as direct examples from West Africa, but as diagnostic warnings about what can happen when automated systems are deployed with weak accountability.

Table 2: Selected Cases of AI Misjudgment and Regulatory Lessons for West Africa

Case	AI system or context	Nature of misjudgment	Regulatory lesson for West Africa
Amazon recruitment tool	Experimental AI tool for screening job applicants	The system reportedly learned from previous hiring patterns in a male-dominated technology workforce and downgraded some applications associated with women. Amazon eventually abandoned the tool (Dastin, 2018).	Recruitment AI may reproduce historical exclusion if trained on past employment records that already favour certain schools, cities, social classes, accents, or career paths.
Facial recognition disparities	Commercial facial analysis systems and facial recognition algorithms	Buolamwini and Gebru (2018) found higher error rates for darker-skinned women, while Grother et al. (2019) found demographic differentials in facial recognition performance, including higher false-positive rates for some racial groups.	Biometric systems used in identity management, banking, elections, migration, and security should not be deployed without local performance testing on West African populations and under West African conditions.
Robert Williams wrongful arrest	Police use of facial recognition	Robert Williams, a Black man in Detroit, was wrongfully arrested after a false facial recognition match (American Civil Liberties Union, 2024).	Algorithmic error can become legal harm when institutions treat AI outputs as authoritative evidence without sufficient human review and appeal mechanisms.
Healthcare allocation algorithm	An algorithm used to identify patients needing additional care	Obermeyer et al. (2019) found that a health algorithm underestimated the needs of Black patients because it used healthcare cost as a proxy for medical need.	AI systems that rely on cost, documentation, or formal service-use data may misread need and vulnerability in contexts where people depend on informal care, delayed treatment, or out-of-pocket payment.
Automated speech recognition	Speech recognition systems used by major technology providers	Koenecke et al. (2020) found higher Word error rates for Black speakers than white speakers, tracing the disparity to underlying acoustic models.	Imported voice-based AI systems may struggle with Nigerian English, Ghanaian English, Pidgin, Hausa-accented English, Yoruba-accented English, Igbo-accented English, Wolof-influenced speech, and other regional speech patterns.

The first example from Table 2 above is Amazon's experimental AI recruitment tool. The system was reportedly trained on about ten years of previous job applications and learned patterns from a male-dominated technology workforce. As a result, it downgraded some applications associated with women, including résumés that contained terms such as "women's." Amazon eventually abandoned the tool (Dastin, 2018). The lesson for West Africa is clear: if recruitment systems learn from past employment records that already favor certain schools, cities, social classes, accents, or career paths, they may reproduce exclusion while appearing neutral. Reuters reported that Amazon scrapped the tool after it was found to be biased against women.

Facial recognition provides a second warning. Buolamwini and Gebru's *Gender Shades* study found that commercial facial analysis systems had the highest error rates for darker-skinned women, reaching 34.7%, compared with 0.8% for lighter-skinned men (Buolamwini & Gebru, 2018). NIST (2023) later found that many facial recognition algorithms showed demographic differentials, including higher false-positive rates for Asian and African American faces relative to Caucasian faces, with some differences ranging from 10 to 100 times depending on the algorithm (Grother et al., 2019). For West Africa, where biometric systems are used or proposed in identity management, banking, elections, migration, and security, this raises a direct policy concern: imported biometric tools should not be trusted without local performance testing.

The third example shows how a technical error can lead to legal harm. Robert Williams, a Black man in Detroit, was wrongfully arrested in 2020 after a false facial recognition match. His case is widely cited as the first publicly reported instance of a wrongful arrest caused by a false face-recognition match (American Civil Liberties Union, 2024). This case matters because it shows that AI harm often occurs when institutions treat an algorithmic output as authoritative evidence. In West African policing, migration control, fraud detection, or public security, a false match could result in arrest, denial of service, reputational damage, or movement restriction before the affected person has any meaningful opportunity to challenge the system.

The fourth example comes from healthcare. Obermeyer et al. (2019) found that a widely used health algorithm underestimated the needs of Black patients because it used healthcare cost as a proxy for medical need. Since Black patients had historically incurred lower health costs despite serious health conditions, the algorithm wrongly classified them as less in need of additional care. Correcting the bias would have increased the proportion of Black patients receiving extra support from 17.7% to 46.5%. This is especially relevant to West Africa because many people rely on informal care, delayed treatment, underfunded clinics, or out-of-pocket payments. A cost-based or documentation-based AI system could easily misread need, risk, and vulnerability.

The fifth example is speech recognition. Koenecke et al. (2020) found that five major automatic speech recognition systems produced substantially higher Word error rates for Black speakers than for white speakers, with an average Word error rate of 0.35 for Black speakers compared with 0.19 for white speakers. The researchers traced the disparity to the underlying acoustic models. For West Africa, this is crucial. Imported voice-based AI systems may struggle with Nigerian English, Ghanaian English, Pidgin, Hausa-accented English, Yoruba-accented English, Igbo-accented English, Wolof-influenced speech, and other regional speech patterns. Where such systems are used in banking, Education, healthcare, customer service, emergency lines, or public complaint channels, poor recognition can quietly become exclusion from services.

Together, these cases show that AI harms often begin before deployment, in the data, assumptions, proxy variables, and testing environments that shape a system's behavior. They also show why data protection alone is insufficient. A system may follow privacy rules and still deny opportunity, misread identity, misunderstand speech, or underestimate need. For West Africa, the central

regulatory question must therefore be: has this system been tested against the people, languages, institutions, and social realities it will affect?

This is where the West African AI Harm Passport becomes necessary. The Passport would require high-impact AI systems to carry evidence of local suitability, bias testing, human accountability, citizen appeal, and incident reporting. It would not be reasonable to assume that a system is safe just because it works elsewhere. It would require proof that the system can operate responsibly in West African conditions.

5. Recommendation

5.1 Proposed Framework: The West African AI Harm Passport

To close the accountability gap identified above, this paper proposes the West African AI Harm Passport as a practical regulatory instrument for high-impact AI systems. The Passport would serve as a mandatory accountability document for any AI system deployed in sectors where automated decisions can significantly affect people's rights, opportunities, safety, dignity, or access to essential services. Its purpose is not to ban AI or discourage innovation. Rather, it is to ensure that systems with serious social consequences can be identified, tested, monitored, challenged, and corrected within West African realities.

The central idea is simple: no high-impact AI system should operate in West Africa without a clear record of what it is, who controls it, what risks it carries, how it was tested, who is accountable for it, and how affected people can seek redress. This would move AI regulation away from abstract ethical promises toward documented responsibility. It would also help regulators distinguish between low-risk innovation and systems that require stronger scrutiny because they affect health, Education, finance, employment, elections, migration, policing, public administration, agriculture, or access to social protection.

The first component of the Passport should be system identity and ownership disclosure. Developers, vendors, and deploying institutions would be required to state the name of the AI system, its developer, country of origin, local distributor, if any, deploying institution, intended purpose, and sector of use. This matters because algorithmic harm is often difficult to address when layers of procurement, outsourcing, and technical complexity obscure responsibility. A citizen harmed by an automated decision should not have to guess whether responsibility lies with a foreign developer, a local vendor, a bank, a hospital, a school, or a government agency.

The second component should be risk classification. The Passport should classify systems as low, medium, high, or prohibited risk, depending on their likely effect on human life, rights, and access to essential opportunities. High-impact systems would include AI used in credit scoring, recruitment, health advice, biometric identification, educational placement, election communication, border control, predictive policing, public benefits, and other sensitive areas. Prohibited systems should include uses that are incompatible with human dignity, democratic participation, or fundamental rights, such as manipulative political targeting, unjustified mass surveillance, or systems that deny essential services without meaningful human review.

The third component should be data provenance and contextual suitability. A deploying institution should disclose, at least to the regulator, the broad categories of data used to train, test, or fine-tune the system. The goal is not to force companies to reveal every trade secret, but to establish whether the system has any credible relationship with the population it will affect. For West Africa, this is crucial. A model trained mainly on data from Europe, North America, or Asia may fail when applied to informal economies, multilingual communication, local names, rural documentation practices, low-bandwidth environments, and culturally specific expressions of need or risk. The Passport should therefore require evidence of local suitability testing before deployment.

The fourth component should be bias and accessibility testing. High-impact AI systems should be assessed for uneven performance across gender, age, disability, language, region, income level, skin tone, where relevant, and urban or rural location. This requirement should not be treated as a technical ritual. It should be connected to lived consequences. If a system performs poorly for speakers of Nigerian English, Pidgin, Hausa, Yoruba, Igbo, Wolof, Twi, Ewe, Fulfulde, or other West African languages, that weakness must be documented before the system is trusted in public-facing decisions. If a biometric or identity system performs poorly in low light, on older phone cameras, or due to inconsistent documentation, regulators should know before citizens are harmed. The fifth component should be human accountability and oversight. Every high-impact AI system should have a named institutional officer or accountable unit within the deploying organization. This person or unit would not serve as a symbolic point of contact. It would be responsible for ensuring that the system is used only for its approved purpose, that human review remains available, that complaints are investigated, and that serious incidents are reported. This is where algorithmic sovereignty becomes practical. A system may be foreign-built, but its use within West Africa must have a local chain of responsibility.

The sixth component should be citizen notice, appeal, and redress. People affected by high-impact AI decisions should be informed when AI has played a significant role in an outcome. They should have access to a plain-language explanation of the decision, the right to request human review, and a complaint pathway that does not require technical expertise or expensive legal support. For example, a loan applicant, job seeker, patient, student, farmer, or welfare applicant should be able to ask, "Was AI used?" What information influenced the outcome? Who reviewed it? How can I challenge it? Without this, accountability remains theoretical.

The seventh component should be post-deployment incident reporting. AI systems can behave differently after deployment because real-world conditions are more complex than test environments. The Passport should therefore require institutions to report serious errors, discriminatory outcomes, security failures, wrongful denials, and repeated complaints to a national AI registry. This registry would allow regulators to identify patterns early rather than waiting for harm to become a public scandal.

The final component should be cross-border notification within ECOWAS. If a system causes serious harm in Nigeria, Ghana, Senegal, Côte d'Ivoire, Sierra Leone, or any other West African country, regulators across the region should be alerted. This would prevent the same defective system from moving quietly from one market to another. It would also allow ECOWAS to build a shared evidence base on AI risks, vendor behavior, and sector-specific vulnerabilities.



Figure 2: The West African AI Harm Passport

In this sense, the West African AI Harm Passport would serve three purposes at once. It would be a screening tool before deployment, an accountability record during use, and a learning mechanism after harm occurs. Its value lies in its practicality. It does not require West Africa to wait until every country has a fully developed AI regulator. It creates a common accountability template that national authorities can begin using immediately and that ECOWAS can later harmonize into a regional standard. By making AI systems answerable to local evidence and public responsibility, the Passport turns algorithmic sovereignty from a political slogan into a regulatory practice.

5.2 Implementation Pathway for Nigeria and ECOWAS

The West African AI Harm Passport should begin as a phased regulatory model rather than a heavy, one-step legal burden. Its first point of implementation should be Nigeria, not because Nigeria represents the whole region, but because its population size, digital economy, and emerging AI policy direction make it a practical testing ground. Nigeria already has institutions that can support the first version of the model, especially the NDPC and the national AI governance structures anticipated under the country's AI strategy (NDPC, 2023; FMCIDE, 2025). The immediate task would be to create a mandatory disclosure and risk-assessment template for high-impact AI systems used in sensitive sectors.

The first phase should focus on public-sector and essential-service AI systems. This would include AI used in healthcare, public Education, identity management, welfare delivery, taxation, elections, migration, agriculture support, and public security. Beginning with these sectors is important because citizens often have limited choice when public institutions deploy digital systems. Where the state or an essential service provider uses AI, the duty of accountability should be higher.

The second phase should extend the Passport to private high-impact sectors, especially banking, insurance, recruitment, digital lending, telecommunications, and large educational platforms. At this stage, compliance should be proportionate. Small local innovators should not face the same burden as multinational vendors deploying automated systems at scale. Low-risk AI tools could be registered through a simplified process, while high-impact systems would require stronger evidence of local testing, appeal mechanisms, and incident reporting.

The third phase should involve regional harmonisation through ECOWAS. A Nigerian pilot would produce lessons on cost, enforcement, institutional capacity, and sector-specific risks. ECOWAS could then adapt the Harm Passport into a regional protocol, model law, or regulatory standard for

high-impact AI systems. This would allow member states to share incident reports, recognise risk classifications, and warn One another when a system causes serious harm in One market. Such coordination would also prevent regulatory arbitrage, in which companies avoid stricter jurisdictions and move harmful technologies into less prepared jurisdictions.

Civil society, universities, consumer protection bodies, professional associations, and technical experts should be involved from the beginning. The Passport should not become a closed bureaucratic form that only regulators and companies understand. Its public-facing elements should be written in clear language so that citizens know when AI is being used, who is responsible, and how to complain. In this way, implementation would not only strengthen state regulation. It would also build public trust, democratic oversight, and regional learning around AI.




Figure 3: Implementation Pathway for Nigeria and ECOWAS

5.3 Challenges, Safeguards, and Policy Risks

The West African AI Harm Passport would face practical challenges. The first is institutional capacity. Many regulatory bodies in the region already struggle with limited funding, technical expertise, enforcement tools, and competing mandates. If the Passport is introduced without capacity-building, it may become a paper requirement rather than a living accountability mechanism. This risk can be reduced through a phased rollout, beginning with high-impact public-sector systems and a small number of sensitive private sectors before expanding to wider commercial use.

The second challenge is regulatory capture. Powerful technology vendors, financial institutions, and public contractors may seek to shape the process in their favour, especially when regulators rely on them for technical explanations. To prevent this, the Passport should include independent review by universities, civil society organisations, consumer protection bodies, and professional associations. Public reporting should also be required for non-sensitive information, so that accountability is not hidden inside closed procurement or licensing processes.

The third challenge is cost. If compliance is too expensive, it may discourage local innovation and strengthen foreign firms that can afford legal and technical audits. This would defeat the purpose of algorithmic sovereignty. The solution is proportional regulation. Low-risk tools developed by small local innovators should face simplified registration. At the same time, high-impact systems that affect rights, safety, finance, health, Education, employment, elections, or public services should carry stronger obligations.

The fourth challenge is corruption and selective enforcement. A Harm Passport could become another approval document that institutions obtain without meaningful testing. To avoid this, the framework should require renewal, random audits, incident reporting, and penalties for false declarations. It should also allow affected citizens and civil society groups to trigger regulatory review when repeated complaints arise.

The final risk is overdependence on foreign models of AI governance. UNESCO, the OECD, the EU, and the AU provide important guidance on trustworthy, ethical, risk-based, and human-centred AI governance (AU, 2024; European Commission, 2025; OECD, 2024; UNESCO, 2021). However, West Africa's model must remain grounded in its own realities: imported systems, informal economies, multilingual societies, uneven infrastructure, weak redress pathways, and cross-border digital markets. The Harm Passport is valuable because it adapts global principles into a practical regional accountability tool.

6. Conclusion

AI governance in West Africa must move beyond broad ethical commitments and privacy-focused regulation toward practical systems of accountability. The region's most urgent challenge is not simply that AI laws are still emerging. It is that high-impact AI systems can already shape access to credit, work, Education, healthcare, security, public services, and civic participation without clear evidence of local testing, human responsibility, or citizen redress.

This paper has argued that algorithmic sovereignty should be understood as the Power to make AI systems answerable to West African realities. The proposed West African AI Harm Passport offers One way to make this possible. By requiring system disclosure, risk classification, local suitability testing, human oversight, complaint pathways, incident reporting, and regional notification, the Passport turns accountability into a practical regulatory process.

For Nigeria and ECOWAS, the goal should not be to reject AI or copy external models. The goal should be to build a governance framework that protects people while allowing innovation to serve the public good. In that balance lies the future of responsible AI in West Africa.

REFERENCES

- African Union. (2024). *Continental artificial intelligence strategy: Harnessing AI for Africa's development and prosperity*. African Union. https://au.int/sites/default/files/documents/44004-doc-EN-Continental_AI_Strategy_July_2024.pdf
- American Civil Liberties Union. (2024). *Williams v. City of Detroit: Face recognition false arrest*. <https://www.aclu.org/cases/williams-v-city-of-detroit-face-recognition-false-arrest>
- Bakare, M. A. (2025). *Artificial intelligence (AI) and democratic stability in West Africa: A game-changer or a passing disruption?* Centre for Democracy and Development West Africa. <https://www.cddwestafrica.org/reports/artificial-intelligence-ai-and-democratic-stability-in-west-africa-a-game-changer-or-a-passing-disruption/>
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 1–15. <https://proceedings.mlr.press/v81/buolamwini18a.html>
- Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. <https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/>
- European Commission. (2025). *AI Act*. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- Federal Ministry of Communications, Innovation and Digital Economy. (2025). *National artificial intelligence strategy*. National Centre for Artificial Intelligence and Robotics. <https://ncair.nitda.gov.ng/wp-content/uploads/2025/09/National-Artificial-Intelligence-Strategy-19092025.pdf>
- Grother, P., Ngan, M., & Hanaoka, K. (2019). *Face recognition vendor test part 3: Demographic effects* (NISTIR 8280)—National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.IR.8280>
- International Development Research Centre. (2025, May 16). *From commitment to action: Advancing the use of AI in Education in Africa*. <https://idrc-crdi.ca/en/research-in-action/commitment-action-advancing-use-ai-education-africa>
- Kasirzadeh, A., & Clifford, D. (2021). Fairness and data protection impact assessments. In *AIES '21: Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 146–153). Association for Computing Machinery. <https://doi.org/10.1145/3461702.3462528>
- 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 automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14), 7684–7689. <https://doi.org/10.1073/pnas.1915768117>
- Ministry of Communication, Digital Technology and Innovations. (2026, April 24). *President Mahama launches Ghana's national AI strategy, calls for AI that reflects Ghanaian values* <https://www.mint.gov.gh/president-mahama-launches-ghanas-national-ai-strategy-calls-for-ai-that-reflects-ghanaian-values/>
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- News Agency of Nigeria. (2025, July 1). *ECOWAS Parliament seeks an AI legal and Education framework*. <https://nannews.ng/ecowas-parliament-seeks-ai-legal-education-framework/>
- Nigeria Data Protection Commission. (2023). *Nigeria Data Protection Act, 2023*. <https://ndpc.gov.ng/download/nigeria-data-protection-act-2023>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>

- Organisation for Economic Co-operation and Development. (2024). *AI principles*.
<https://www.oecd.org/en/topics/sub-issues/ai-principles.html>
- United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the ethics of artificial intelligence*. UNESCO.
<https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>
- Wachter, S., & Mittelstadt, B. (2019). A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI. *Columbia Business Law Review*, 2019(2), 494–620.
<https://journals.library.columbia.edu/index.php/CBLR/article/view/3424>
- World Bank. (2025). *The use of alternative data in credit risk assessment: Opportunities, risks, and challenges*. <https://doi.org/10.1596/42970>