AISCA
AI SKILLS AND COMPUTE AFRICA FOUNDATION

POLICY BRIEF

Practical AI Governance for African Regulators

Africa AI Stakeholder Meeting on Innovative Governance and Capacity Building

Introduction

The peer exchange convened representatives from data protection authorities across eight African countries, domain regulators, software engineers building AI tools, policy specialists, and academics. Bringing these groups together ensured legal discussions engaged directly with technical realities from the outset. Held under the Chatham House Rule, this brief captures the room's collective analysis with- out attributing specific statements or positions to any individual or institution.

Discussions centered on a core premise: existing legal frameworks, particularly data protection laws, al- ready function as de facto AI governance. While foundational legal principles apply effectively to AI sys- tems, the operational mechanisms supporting them face practical implementation challenges. This document summarizes the key insights, debate points, and consensus outcomes from the meeting.

The findings below reflect the principal thematic areas explored during the exchange.

1 A shared technical baseline

Establishing precise terminology was critical because regulatory oversight depends on the specific component under review. Traditional AI models analyze existing data to score, rank, or flag information, characterizing most current deployments in the region. In contrast, generative AI creates new synthetic output. Distinguishing a standalone model from a complete AI system is equally vital. A model is the trained core algorithm, while a system includes data pipelines, user interfaces, decision rules, and oper- ational workflows. Standard definitions, including the Organisation for Economic Co-operation and Development's November 2023 revision, focus on full systems rather than isolated models. Consequently, supervisory authorities typically oversee local system deployments even when the un- derlying core model was developed elsewhere.

The discussion also differentiated standard conversational chatbots from autonomous AI agents. While chatbots respond directly to user queries, AI agents execute multi-step workflows across external data- bases, payment channels, and file systems without continuous human authorization. This creates a dis- tinct audit challenge: the execution trail needed to reconstruct automated decisions may be frag- mented across multiple system environments.

2 Where data protection law reaches, and where it strains

Regulators across several African jurisdictions have already taken formal action on AI implementations using existing statutory mandates, covering areas such as healthcare data, facial recognition, workplace biometrics, iris scanning, and synthetic media. Rather than asking whether data protection law applies broadly to AI, regulators evaluated the specific harms those statutory provisions address. For instance, cross-border transfer restrictions exist to preserve regulatory oversight over personal data. Identifying specific risks enables authorities to select and apply the appropriate mechanism effectively.

Participants analyzed whether personal data incorporated into model weights remains personal data in the legal sense. One perspective suggested that models merely store mathematical weights and refer- ence external databases, meaning personal data exists only in input and output streams. Another view noted that fine-tuned models can retain detailed feature representations, leaving model parameters vulnerable to extraction attacks that recover identifiable information. The consensus distinguished gen- eral foundation models from models fine-tuned on personal datasets, concluding that fine-tuned vari- ants are substantially more likely to retain personal data.

Regarding the right to erasure, deleting data from a training set does not automatically eliminate its influ- ence on an already trained model. Participants emphasized that technical complexity does not exempt organizations from statutory compliance. Where complete retraining is impractical, alternative mea- sures, such as output filtering or targeted machine unlearning, should be evaluated based on the specific outcome requested by the individual. Similar complexities apply to data access requests, as inspecting raw model weights does not provide intelligible information to data subjects.

Existing frameworks for international data transfers assume static transfers to identifiable recipients. While cross-border model training fits this framework, real-time inferencing across distributed cloud in- frastructure often routes queries dynamically without predefined server destinations. Current adminis- trative practices vary: some jurisdictions place compliance accountability primarily on data controllers, while others require prior risk assessments and formal authorization for transfers. Participants noted that rigid pre-approval requirements for every cross-border query could create administrative bottlenecks that hinder timely processing without improving protection.

Legal frameworks across the region treat automated decision-making differently, particularly regarding opt-in rights, notification duties, explanation rights, and options for human review. Common statutory terms such as “solely automated,” “significant effect,” and “meaningful human intervention” require clear operational definitions. Superficially reviewing and confirming an automated output does not constitute meaningful human oversight. Furthermore, individuals must first know that an automated decision oc- curred before they can exercise their rights to challenge it or request an explanation.

Navigating these operational challenges requires a clear understanding of how responsibility and ac- countability are distributed across the entire AI ecosystem.

3 Liability across the lifecycle

To examine liability allocation, participants reviewed a scenario where a diagnostic AI tool for tuberculo- sis screening missed a positive diagnosis. The deployment chain involved multiple entities: healthcare facilities providing training images, data labeling contractors, the base model developer, cloud infra- structure providers, downstream software vendors adapting the tool, public procurement authorities, purchasing hospitals, and attending clinicians evaluating final outputs.

The analysis confirmed that legal liability is phased across stages and depends on establishing fault, harm, and direct causation. Responsibility may be shared between system developers and deploying organizations. Determining liability requires comprehensive technical audits to reconstruct system op- erations, highlighting technical audit capabilities as a fundamental component of regulatory oversight.

4 One technology, many mandates

A single AI deployment can intersect with over twenty regulatory areas, including consumer rights, mar- ket competition, financial services, employment, intellectual property, sector licensing, and public pro- curement. Professional oversight bodies governing medical, legal, and engineering practices also inter- act directly with these deployments.

Participants clarified structural distinctions between different regulatory intersections. Overlapping mandates occur when multiple authorities have jurisdiction over the same activity, requiring clear referral protocols. Regulatory conflict arises when compliance with one authority's requirement violates another's, necessitating joint guidance. Regulatory gaps occur when emerging system harms fall out- side existing statutory definitions, requiring proactive inter-agency coordination.

Agencies reported handling AI-related incident notifications, complaints regarding automated recruit- ment filtering, and requests for compliance guidance. Effective oversight requires mapping personal data flows across every processing stage to clearly identify responsible entities. Establishing formal in- stitutional channels for inter-agency coordination enables regulatory bodies to align supervisory ap- proaches and publish joint administrative guidance efficiently.

5 Governance models beyond writing a new law

Designing AI governance strategies involves navigating rapid technological change, information asym- metries between developers and oversight bodies, and cross-sectoral operational impacts. Effective oversight depends on five key operational questions: what specific component is being regulated, where technical information resides, which regulatory instrument is suitable, how compliance will be ver- ified, and how rules adapt as technology evolves.

The group evaluated four governance frameworks using a digital credit scoring case study. Regulatory sandboxes allow authorities to observe live system behavior within controlled parameters. To be effec- tive, sandbox testing should conclude with clear evaluation reports and formal administrative determi- nations. Sandboxes foster mutual technical learning between regulators and developers, though author- ities must maintain independent evaluation capabilities.

Co-regulatory models combine statutory boundaries set by public authorities with technical standards developed alongside industry bodies. Technical standards and conformity assessments provide struc- tured baselines for evaluating compliance. Adaptive regulation relies on flexible, principle-based rules designed to remain applicable as underlying technologies evolve.

The session emphasized the importance of aligning legislative proposals with existing national digital strategies and economic policies to ensure statutory frameworks remain practical and coherent.

Conclusion

While foundational legal principles apply to AI implementations, specific administrative and technical mechanisms require adaptation to function effectively at scale. Regulatory bodies are increasingly called upon to oversee AI systems within their existing mandates, making technical capacity building es- sential. Addressing information asymmetries and cross-sector challenges requires strong inter-agency cooperation, shared technical baselines, and sustained institutional collaboration across the region.