

PUBLIC CONSULTATION / REGULATORY ENGAGEMENT

Comments on Draft Guidance Note on Artificial Intelligence

Submitted by the African Centre for AI Governance and Oversight (ACAIGO)

SUBMITTED TO	Office of the Data Protection Commissioner
CONSULTATION	ACAIGO/SE/1/2026
DATE	17 th August 2026
CONTACT	Email: info@acaigo.africa , website: https://www.acaigo.africa/

Executive summary

ACAIGO welcomes the opportunity to provide comments on the Draft Guidance Note on Artificial Intelligence. Our comments focus on the clarity, legal and technical coherence, and practical application of the proposed guidance, with a view to supporting a framework that provides meaningful guidance on the application of Kenya's data protection law to artificial intelligence.

KEY RECOMMENDATIONS

- Align the substantive content of the Guidance Note with its stated scope
- Strengthen the guidance on cross-border transfers of personal data
- Adopt a clear and technically sound definition of Artificial Intelligence
- Develop detailed guidance on Legitimate Interest Assessments (LIAs)
- Distinguish mandatory obligations from recommended AI governance measures

Detailed comments

Page No.	Statement as currently captured	Proposed Statement	Comments/Justification
4	“Artificial Intelligence (AI)” means a machine-based system that can, for a given set of objectives, make predictions, recommendations, or decisions influencing real or virtual environments, operating with varying levels of autonomy.	ACAIGO refrains from proposing an alternative definition of AI, recognising that multiple definitions exist across legal, scientific and technical contexts. We nevertheless recommend that the Office adopt a more conceptually and technically sound definition that clearly delineates the systems intended to fall within the scope of the Guidance.	The definition of AI in the Draft Guidance Note closely resembles the OECD’s 2019 definition of an AI system. However, its broad, capability-focused formulation does not establish a sufficiently clear distinction between AI systems and conventional automated software. Conventional software may also operate according to predefined objectives, produce predictions, recommendations or decisions, and influence physical or virtual environments. Consequently, the definition may encompass systems that would not ordinarily be regarded as AI, creating uncertainty regarding the intended scope of the Guidance
6	Training Data” means the dataset used to train an artificial intelligence or machine	“Training data” means a dataset used to train an artificial intelligence model, including a	The use of “or” is conceptually misleading because it presents artificial

	learning model, which may contain personal data.	machine-learning model, which may contain personal data.	intelligence and machine learning as separate, mutually exclusive categories, despite the Guidance Note defining machine learning as a subset of AI. It should instead refer to “an artificial intelligence model, including a machine-learning model” to maintain consistency with the established relationship between the two concepts.
8	Discriminatory profiling and bias: AI systems trained on historical data can perpetuate and amplify existing patterns of discrimination, resulting in profiling that disproportionately disadvantages individuals based on race, gender, ethnicity, disability, or economic status, characteristics that intersect with sensitive personal data categories under the Act.	Discriminatory profiling and bias: AI systems may perpetuate or amplify patterns of discrimination reflected in historical data. Bias may also arise from the way training data is collected, selected, represented or labelled, as well as from model design, testing, deployment and human interpretation of outputs. This may result in profiling that disproportionately disadvantages individuals based on race, gender, ethnicity, disability, economic status or other characteristics that intersect with sensitive personal data categories under the Act.	The provision appears to attribute discriminatory profiling and bias primarily to historical training data. However, bias may arise at multiple stages of the AI lifecycle, including data collection and labelling, dataset selection, model design, choice of proxy variables, system deployment and human interpretation of outputs. The wording should therefore be broadened to avoid presenting historical data as the only source of AI-related bias.

8	Cross-border data flows: AI services are frequently provided by offshore processors, and training data may be transferred to and processed in foreign jurisdictions without adequate data protection frameworks, contrary to the Act's transfer restrictions.	Cross-border data flows: AI services are frequently provided by offshore service providers, and training data may be transferred to and processed in foreign jurisdictions without compliance with the applicable transfer requirements and safeguards under the Act.	The provision assumes that offshore AI providers will act as processors, although their legal status will depend on their actual role in determining the purposes and means of processing. An offshore AI provider may act as a controller, joint controller or processor. Furthermore, the wording may suggest that personal data can only be transferred to a jurisdiction with an adequate data protection framework. However, Kenyan data protection law recognizes other mechanisms for transferring personal data outside Kenya, subject to the applicable legal requirements. The provision should therefore use role-neutral language and refer more broadly to compliance with the cross-border transfer requirements under the Act and Regulations. It should also recognize that cross-border processing may extend beyond training data to include user inputs, system
---	--	--	---


			outputs, prompts, logs and other operational data.
13	Personal data used to train AI models shall not be retained longer than is necessary for the training purpose. Entities shall establish and enforce data retention policies that address AI training datasets specifically, including deletion or anonymization of training data upon completion of the relevant training cycle. Model weights and parameters that encode personal information should also be addressed in retention and deletion policies.	Personal data used to develop or deploy AI systems shall not be retained longer than is necessary for the relevant processing purpose. Entities shall establish and enforce data retention policies that address personal data processed throughout the AI lifecycle, including training datasets, user inputs, system outputs and interaction logs. Model weights and parameters that encode personal information should also be addressed in retention and deletion policies.	The storage-limitation provision should reflect the stated scope of the Guidance Note by addressing personal data processed during both the development and deployment of AI systems, rather than focusing exclusively on training data. This narrow focus on AI model training is also a recurring issue in the preceding sections (for instance the purpose limitation section) and is inconsistent with the stated purpose of the Guidance Note, which extends to the development, deployment and procurement of AI systems.
14	Entities shall identify the lawful basis before processing commences and shall document that identification.	Entities shall identify the applicable lawful basis before processing commences and document the selected basis and the reasons for relying on it.	The amendment clarifies that entities must document the lawful basis selected and the reasons for relying on it, rather than merely recording that an

			identification exercise was undertaken.
15	Where a statutory or regulatory obligation requires an entity to conduct a processing activity, and that activity is performed using AI, the legal obligation lawful basis may be relied upon. For example, a regulated financial institution using AI for anti-money laundering transaction monitoring may rely on the legal obligations imposed by the Proceeds of Crime and Anti-Money Laundering Act.	Where a statutory or regulatory obligation requires an entity to conduct a processing activity, and that activity is performed using AI, the legal obligation lawful basis may be relied upon, provided that the processing undertaken through the AI system is necessary and proportionate for compliance with that obligation. For example, a regulated financial institution using AI for anti-money-laundering transaction monitoring may rely on the legal obligations imposed by the Proceeds of Crime and Anti-Money Laundering Act, subject to satisfying these requirements.	A statutory obligation to undertake a processing activity does not necessarily require or automatically justify the use of AI as the means of performing that activity. The amendment clarifies that the entity must demonstrate that the particular processing undertaken through the AI system is necessary and proportionate for compliance with the relevant legal obligation.
15 and 16	Entities relying on legitimate interests for AI processing shall conduct and document a Legitimate Interests Assessment (LIA) that: • identifies the legitimate interest; • demonstrates necessity; and 15 • balances the interest against the data subject's rights, considering the intrusiveness of AI	This provision should provide more practical and AI-specific guidance on conducting a Legitimate Interests Assessment. At a minimum, it should explain the factors to be considered when assessing necessity and balancing the legitimate interest against the rights and interests of data subjects, including	The provision merely restates the conventional stages of a Legitimate Interests Assessment without explaining how they should be applied to AI processing. It should provide practical guidance on factors such as

	profiling and automated decision-making.	the nature and source of the personal data, the scale of processing, reasonable expectations of data subjects, the likelihood and severity of harm, the involvement of vulnerable persons, and available safeguards. However, an elaborate alternative framework is beyond the scope of these comments.	the nature and source of the personal data, the scale of processing, data subjects' reasonable expectations, potential harms, vulnerable persons and appropriate safeguards. Given the increasing reliance by AI developers globally on legitimate interests for activities such as AI model training, the Data Commissioner should also consider issuing a separate and comprehensive guidance note on conducting Legitimate Interests Assessments.
17	Data subjects have the right to object to the processing of their personal data for purposes of profiling or direct marketing, including profiling by AI systems.	Data subjects have the right to object to the processing of their personal data, including processing involving profiling by AI systems. The right to object applies as an absolute right where the processing is for direct marketing purposes, including profiling to the extent that it is related to such direct marketing.	The existing wording unduly confines the right to object to profiling and direct marketing, whereas section 36 of the Data Protection Act and Regulation 8(1) of the Data Protection General Regulations establishes a general right to object to the processing of personal data. The amendment also reflects regulation 8(4) of the Data Protection

			(General) Regulations, which provides that the right is absolute where processing is for direct marketing purposes, including profiling to the extent that it relates to such direct marketing.
21	A private school deploys an AI system to monitor student behavior and performance using biometric data (facial recognition) and health-related information to identify at-risk students. The system processes data of minors without obtaining verifiable parental consent. Subsequently, a breach occurs at a third-party service provider, exposing student biometric and health data, but the school delays notifying affected parents and the regulator. This raises multiple violations. The school fails to: • establish a valid legal basis and implement heightened safeguards for processing sensitive personal data; • obtain parental consent and apply appropriate protections for children's data;	A private school deploys an AI system to monitor student behavior and performance using facial-recognition and health-related data to identify students considered at risk. The school does not establish that this intrusive processing is necessary and proportionate or that the same objective could not reasonably be achieved through less intrusive means. It also processes children's sensitive personal data without obtaining verifiable parental consent or demonstrating that the processing protects and advances the children's rights and best interests. Subsequently, a breach occurs at a third-party service provider, exposing the students' biometric and health data, but the school delays notifying affected parents and the Data Commissioner.	The existing example focuses on the absence of parental consent and safeguards, which may imply that these measures alone would legitimize AI-enabled biometric monitoring of students. Given the highly intrusive nature of continuous monitoring involving children's biometric and health data, the example should also address necessity, proportionality, data minimization, the child's best interests and the availability of less intrusive alternatives.

	and • notify the breach promptly and inform affected data subjects as required. As a result, the school fails to meet its obligations in handling sensitive personal data, protecting children, and responding to data breaches.		
22	Entities shall assess and document the adequacy of data protection in any jurisdiction to which personal data is transferred in connection with AI processing and shall implement contractual or other safeguards including binding corporate rules or contractual clauses where adequacy cannot be established. Entities may not transfer personal data to offshore AI processors without a lawful transfer basis.	Entities shall identify and document the applicable lawful basis for any transfer of personal data in connection with AI processing. A transfer may be based on an adequacy decision by the Data Commissioner, appropriate safeguards, including binding corporate rules or contractual clauses or any other transfer basis permitted under the Act and Regulations. Personal data shall not be transferred outside Kenya unless an applicable transfer basis has been established.	The existing provision conflates adequacy decisions, appropriate safeguards and other lawful bases for cross-border transfers. Under regulation 44 of the Data Protection (General) Regulations, determining whether a country, territory, specified sector or international organisation provides an adequate level of protection is reserved for the Data Commissioner; it is not a determination to be made independently by transferring entities. This is also consistent with international best practice, under which adequacy decisions are ordinarily made by the competent regulatory or public

		authority rather than by individual transferring entities. As at the date of these comments, the Data Commissioner has not published a list of jurisdictions or organisations deemed adequate. Consequently, in the absence of a relevant adequacy decision, entities cannot presently rely on adequacy as a transfer basis and must identify another applicable basis under sections 48 and 49 of the Act and Part VII of the Regulations. The amendment also avoids limiting the provision to transfers made to offshore AI processors, as personal data may be transferred to various categories of recipients in connection with AI processing. This comment also applies to the cross-border data transfer checklist on page 30 of the Guidance Note.
--	---	--

22	A Kenyan bank deploys a cloud-based AI fraud detection system provided by an overseas vendor. Customer transaction data is shared with the vendor for processing, but no formal data processing agreement is in place. The data is transferred and processed on servers located outside Kenya without assessing the adequacy of data protection in the receiving jurisdiction or implementing appropriate safeguards. This raises compliance failures. The bank fails to: • enter into a data processing agreement governing the vendor’s handling of personal data; and • assess and establish a lawful basis for cross-border data transfer and implement appropriate safeguards. As a result, the bank does not meet its obligations when engaging third-party AI processors and transferring personal data outside Kenya.	A Kenyan bank deploys a cloud-based AI fraud-detection system provided by an overseas vendor. Customer transaction data is shared with the vendor for processing, but no formal data-processing agreement is in place. The data is transferred to and processed on servers located outside Kenya without establishing an applicable lawful basis for the transfer and implementing appropriate safeguards. This raises compliance failures. The bank fails to: • enter into a data-processing agreement governing the vendor’s handling of personal data; and • establish and document a lawful basis for the cross-border transfer and implement appropriate safeguards where required. As a result, the bank does not meet its obligations when engaging a third-party AI processor and transferring personal data outside Kenya.	The example incorrectly suggests that the bank is responsible for assessing the adequacy of data protection in the receiving jurisdiction. Under regulation 44 of the Data Protection (General) Regulations, adequacy determinations are reserved for the Data Commissioner. The bank should instead establish whether a relevant adequacy decision exists and, in its absence, identify and document another applicable transfer basis and implement the safeguards required under the Act and Regulations.
----	---	--	---

24 and 25	Example Savannah Analytics, a data science firm, seeks to train a natural language processing model using social media posts. Before proceeding, the firm shall assess: • whether the personal data in those posts was made publicly available in a manner that permits use for AI training; • whether a lawful basis exists under Section 30 of the Act; and whether data minimization requires use of anonymized or synthetic text data instead; and whether a DPIA is required given the scale of the processing.	The Guidance Note should elaborate on how a controller or AI developer should determine whether personal data contained in online posts was made publicly available in a manner that permits its use for AI training. Although developing comprehensive criteria is beyond the scope of these comments, the Data Commissioner should provide sufficiently detailed and practical guidance to enable entities to make this assessment consistently and lawfully.	Although the Guidance Note correctly recognises that public availability does not, by itself, authorise the use of personal data for AI training, it does not provide sufficiently clear criteria for determining when publicly available personal data may lawfully be used for that purpose. The Data Commissioner should establish practical criteria to guide controllers and AI developers in conducting this assessment.
26	confirming registration with the Office, or updating an existing registration to reflect the new AI processing activity;	confirming registration with the Office, or updating an existing registration to reflect the new AI processing activity, where registration is required under the Act and applicable Regulation;	Not every entity is eligible to register with the Office. The amendment clarifies that the obligation applies only to entities that meet the prescribed registration thresholds and are not otherwise exempt under the applicable Regulations. This applies to page 33 of the Draft Guidance Note as well

27	Entities shall establish an AI governance committee or equivalent body with clear responsibility for monitoring AI system compliance on an ongoing basis.	Entities are advised to establish an AI governance committee or equivalent body with clear responsibility for monitoring AI system compliance on an ongoing basis	The use of “shall” appears to impose a mandatory requirement on every entity to establish an AI governance committee or equivalent body. Unless this requirement is grounded in a specific provision of the Act or Regulations, the Guidance Note should not prescribe a particular internal governance structure. Entities may assign responsibility for AI compliance through different arrangements depending on their size, resources, processing activities and the risks presented by the AI systems they use. The provision should therefore be framed as a recommended governance measure rather than a mandatory obligation.
29	Have you clearly defined the purpose of the AI system? Have you identified and documented a valid lawful basis? Is the use of data for AI training compatible with the	Have you clearly defined the purpose of the AI system? Have you identified and documented a valid lawful basis for each processing activity undertaken during the development,	The compliance checklist focuses on the use of personal data for AI training, despite the stated scope of the Guidance Note extending to the

	original purpose or supported by a new lawful basis?	deployment and procurement of the AI system? At each stage, is the processing compatible with the original purpose for which the personal data was collected or supported by a new lawful basis?	development, deployment and procurement of AI systems. The checklist should address each processing activity involving personal data throughout the AI lifecycle, rather than treating training as the principal or only relevant activity.
30	Entities that develop or deploy AI systems processing personal data shall establish and document an AI governance framework, proportionate to the nature, scale, and risk of their AI processing, that specifies:	Entities that develop or deploy AI systems processing personal data are advised to establish and document an AI governance framework, proportionate to the nature, scale and risk of their AI processing, that specifies:	The use of “shall” presents the establishment of a distinct AI governance framework as a mandatory requirement. Although such a framework may support accountability and ongoing compliance, the Guidance Note should not impose a specific governance structure unless that requirement is expressly grounded in the Act or Regulations. Entities should retain flexibility to incorporate AI-related risks and responsibilities into their existing data protection, risk management and compliance frameworks. The provision should

			therefore be framed as a recommended governance measure rather than a mandatory obligation.
33	Data minimization- Review AI training datasets and inference inputs to ensure only necessary data is collected and processed. Remove or anonymize data that is not required.	Review the personal data processed in connection with the development, deployment and procurement of AI systems to ensure that only data necessary for the relevant purpose is collected and processed. This should include training, validation and testing datasets, inference inputs, system outputs and operational logs. Remove or anonymize personal data that is not required.	The existing checklist limits its explanation of data minimization primarily to training datasets and inference inputs. This does not fully reflect the stated scope of the Guidance Note, which extends to the development, deployment and procurement of AI systems. Data minimization should be assessed in relation to each processing activity within that scope.

For ACAIGO

Jane Awuor, FIP, AIGP
LLM AI and the Law
Queen Mary University of London

AFRICAN CENTRE FOR AI GOVERNANCE AND OVERSIGHT



<https://www.acaigo.africa/> | info@acaigo.africa | Nairobi, Kenya

