



OFFICE OF THE DATA PROTECTION COMMISSIONER

Guidance Note on Artificial Intelligence

July 2026

Draft Guidance Notes on AI

CONTENTS

DEFINITION OF TERMS	4
FOREWORD	Error! Bookmark not defined.
1. OFFICE	7
2. INTRODUCTION	7
2.1. Background	7
2.2. Privacy and Data Protection Concerns in AI	8
2.3. Scope and Purpose of Guidance Note.....	9
3. LEGISLATIVE FRAMEWORK.....	9
a) The Constitution of Kenya, 2010:.....	9
b) The Data Protection Act, Cap 411C:	10
c) The Data Protection (General) Regulations, 2021	10
d) The Data Protection (Registration of Data Controllers and Data Processors) Regulations, 2021	11
e) The Data Protection (Complaints Handling Procedure and Enforcement) Regulations, 2021	11
f) Other Applicable Laws	11
4. APPLICATION OF DATA PROTECTION PRINCIPLES TO AI SYSTEMS.....	11
4.1. Lawfulness, Fairness, and Transparency.....	11
4.2. Purpose Limitation.....	12
4.3. Data Minimisation.....	12
4.4. Accuracy.....	13
4.5. Storage Limitation	13
4.6. Integrity and Confidentiality	13
4.7. Accountability	14
5. LAWFUL BASIS FOR PROCESSING PERSONAL DATA IN AI SYSTEMS	14
5.1. Consent.....	14
5.2. Contract	15
5.3. Legal Obligation	15
5.4. Vital Interests	15
5.5. Public Interest or Exercise of Official Authority	15
5.6. Legitimate Interests.....	15
6. RIGHTS OF DATA SUBJECTS IN AI CONTEXTS	16
6.1. Right to Be Informed	16
6.2. Right of Access	16
6.3. Right to Rectification	16
6.4. Right to Erasure	16
6.5. Right to Object.....	17
6.6. Right Not to Be Subject to Automated Decision-Making	17
6.7. Right to Data Portability	18

7.	OBLIGATIONS OF DATA CONTROLLERS AND DATA PROCESSORS IN AI SYSTEMS	19
7.1.	Registration	19
7.2.	Privacy by Design and Default	19
7.3.	Data Protection Impact Assessment (DPIA)	19
7.4.	Data Protection Officer (DPO).....	20
7.5.	Sensitive Personal Data and AI	20
7.6.	Children and AI Systems	20
7.7.	Notification and Communication of a Breach	20
7.8.	Data Sharing and Third-Party AI Processors.....	21
7.9.	Cross-Border Transfers of AI-Processed Data.....	22
8.	DATA PROTECTION AND THE AI LIFECYCLE	22
8.1.	Stage 1: Problem Definition and AI System Design.....	23
8.2.	Stage 2: Data Collection and Preparation.....	24
8.3.	Stage 3: Model Training and Development	25
8.4.	Stage 4: Testing, Validation, and Pre-Deployment Review.....	25
8.5.	Stage 5: Deployment and Integration	26
8.6.	Stage 6: Operation, Monitoring, and Maintenance	26
8.7.	Stage 7: Decommissioning and Data Disposal	27
9.	HIGH-RISK AI APPLICATIONS: SPECIFIC OBLIGATIONS.....	28
10.	PRE-PROCESSING CHECKLIST FOR AI SYSTEMS	29
11.	AI GOVERNANCE FRAMEWORK.....	30
12.	GENERATIVE AI.....	30
13.	ENFORCEMENT AND COMPLAINTS.....	31
	ANNEX 1: COMPLIANCE CHECKLIST FOR AI SYSTEMS.....	33
	ANNEX 2: DPIA TRIGGER TABLE FOR AI USE CASES	35

DEFINITION OF TERMS

"Act" means the Data Protection Act, Cap. 411C.

"AI Deployer" means an entity that uses an artificial intelligence system during its operations to process personal data or to produce outputs, recommendations, or decisions affecting data subjects.

"AI Developer" means an entity that designs, builds, trains, or substantially modifies an artificial intelligence system or artificial intelligence model, whether for its own use or for provision to others.

"AI Provider" means an entity that makes an artificial intelligence system or artificial intelligence model available to others, directly or through an intermediary, including through artificial intelligence-as-a-service platforms.

"AI User" means a natural person who interacts with or operates an artificial intelligence system under the authority of an AI deployer.

"Algorithmic Transparency" means the disclosure of meaningful information about the logic, significance, and envisaged consequences of artificial intelligence-driven decision-making to data subjects and the public.

"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.

"Automated Decision-Making (ADM)" means the process of deciding solely by automated means, without any meaningful human involvement.

"Biometric Data" means personal data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of a natural person which allows or confirms the unique identification of that person.

"Data Commissioner" means the person appointed pursuant to Section 6 of the Act.

"Data Controller" means a natural or legal person, public authority, agency, or other body which, alone or jointly with others, determines the purpose and means of processing of personal data.

"Data Processor" means a natural or legal person, public authority, agency, or other body which processes personal data on behalf of the data controller.

"Data Protection Impact Assessment (DPIA)" means a process designed to identify and mitigate data protection risks arising from a processing activity, required under Section 31 of the Act.

"Data Protection Officer (DPO)" means the officer designated pursuant to Section 24 of the Act to advise on and monitor compliance with data protection obligations.

"Data Subject" means an identified or identifiable natural person who is the subject of personal data.

"Explainability" means the degree to which the internal mechanics of an artificial intelligence system can be understood and communicated in meaningful human terms to those affected by its outputs.

"Foundation Model Provider" means an entity that develops or provides a large artificial intelligence model trained on broad datasets at scale, including a large language model, that is capable of being adapted to a wide range of downstream tasks.

"Generative AI" means artificial intelligence systems capable of generating new content, including text, images, audio, video, or code, based on training data.

"High-Risk AI Application" means an artificial intelligence system whose use is likely to result in significant risks to the fundamental rights, freedoms, or vital interests of data subjects, including in the areas of credit scoring, employment, healthcare, criminal justice, and children's services.

"Large Language Model (LLM)" means a type of artificial intelligence system trained on large volumes of text data, capable of generating, summarising, classifying, or translating natural language.

"Machine Learning (ML)" means a subset of artificial intelligence in which systems learn from data to improve their performance on a task without being explicitly programmed.

"Office" means the Office of the Data Protection Commissioner established under Section 5 of the Act.

"Personal Data" means any information relating to an identified or identifiable natural person.

"Privacy by Design and by Default" means the integration of data protection principles and technical measures into the design of artificial intelligence systems and processing activities from the outset.

"Processing" means any operation or set of operations performed on personal data, whether by automated means, including collection, recording, organisation, storage, adaptation, retrieval, use, disclosure, combination, restriction, erasure, or destruction.

"Profiling" means any form of automated processing of personal data to evaluate certain personal aspects relating to a natural person, to analyse or predict aspects concerning performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location, or movements.

"Regulations" means all regulations made in accordance with Section 71 of the Act.

"Sensitive Personal Data" means data revealing a natural person's race, health status, ethnic or social origin, conscience, belief, genetic data, biometric data, property details, marital status, family details, sex, or sexual orientation.

"Training Data" means the dataset used to train an artificial intelligence or machine learning model, which may contain personal data.

Draft Guidance Notes on AI

1. OFFICE

The Office of the Data Protection Commissioner (Office) is a government agency established to protect the privacy and security of personal data in our increasing digital world. It is responsible for enforcing data protection laws and policies to safeguard the privacy, dignity, and fundamental rights of individuals. The Office is mandated to oversee the implementation and enforcement of the Data Protection Act, 2019, which regulates the processing of personal data of persons located in Kenya by both private and public sector organisations.

The Office plays a vital role in ensuring that individuals have control over their personal data and that organisations respect their privacy rights. The Office's work involves monitoring and enforcing compliance with data protection regulations, investigating data breaches, and imposing sanctions on entities that violate data protection laws. In addition, the Office is responsible for raising public awareness about data protection issues and educating individuals and organisations on how to protect personal data. With the growing importance of data protection in our digital age, the Office is a critical institution in maintaining trust and confidence in our data-driven society.

The Office is uniquely positioned to facilitate both the government and private sector entities in achieving Government's strategic goals under the "Bottom-Up Economic Transformation Agenda" and, in particular, its digital superhighway initiative. As the digital landscape expands, the need for robust data protection mechanisms becomes paramount. The Office, with its mandate to oversee, regulate, and ensure lawful data processing, plays a pivotal role in this transformation.

Kenya remains at the cutting edge of digital transformation while maintaining stringent data protection standards. The Office serves as a key stakeholder and regulator in guiding the nation's digital superhighway journey by ensuring that as we advance technologically, the rights and privacy of individuals remain safeguarded.

2. INTRODUCTION

2.1. Background

AI is increasingly deployed across Kenya's public and private sectors. Financial institutions use AI-driven credit scoring and fraud detection. Hospitals and health insurers apply machine learning to diagnostic imaging and patient risk stratification. Law enforcement agencies explore predictive analytics. Employers use algorithmic tools to screen job candidates. Retailers and digital platforms deploy AI for personalised advertising and content recommendation. Government agencies process citizen data through AI-enabled service delivery and benefits administration systems.

This widespread adoption of AI is occurring against the backdrop of Kenya's ambitious digital transformation agenda and growing data infrastructure, including expanding mobile internet penetration, the proliferation of digital financial services, and increasing government digitisation. The volume of personal data being collected, aggregated, and fed into AI training pipelines and inference systems is growing exponentially.

AI systems by their nature introduce novel data protection challenges that existing sectoral guidance has not fully addressed: the opacity of algorithmic models, the aggregation of large training datasets, the generation of inferences and predictions about data subjects without their direct input, the risk of discriminatory outcomes from biased training data, and the reduced role of human decision-

makers in consequential choices. These characteristics require specific data protection guidance calibrated to the AI context.

AI systems process personal data in ways that extend beyond data directly provided by individuals. In addition to collected data, these systems generate inferred data such as behavioural profiles, risk scores, and predictions about individuals. Although not explicitly supplied by the data subject, such inferences relate to identifiable people and therefore constitute personal data under Section 2 of the Act. Their use raises important transparency and fairness obligations under Sections 25 and 29, particularly where individuals may be unaware that such conclusions are being drawn.

AI systems also produce outputs, including classifications, recommendations, and automated decisions that may significantly affect individuals. Where these outputs relate to identifiable persons, they fall within the scope of personal data and are subject to the protections of the Act. Organizations should therefore ensure compliance not only at the point of data collection, but across the full AI processing chain, including the use of outputs and the safeguards required for automated decision-making under Section 35.

2.2. Privacy and Data Protection Concerns in AI

The deployment of AI systems raises a distinct and significant set of privacy and data protection concerns:

- **Opacity and lack of explainability:** Many AI models, particularly deep learning systems, operate as "*black boxes*", making it difficult to explain how a particular output or decision was reached. This fundamentally challenges the transparency principle and the right of data subjects not to be subjected to solely automated decisions without explanation.
- **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.
- **Excessive data collection for training:** The training of large AI models often requires vast datasets of personal data, creating pressure to collect and retain more data than would be justified under the data minimisation and storage limitation principles.
- **Re-identification risk:** Ostensibly anonymised datasets used for AI training can be re-identified through linkage attacks or inference from model outputs, undermining the effectiveness of anonymisation as a de-risking measure.
- **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.
- **Children and vulnerable groups:** AI systems may process children's data or be deployed in contexts affecting vulnerable individuals, in education, child welfare, or social protection without appropriate safeguards calibrated to heightened risk.
- **Generative AI and synthetic data:** Large language models and generative AI systems may reproduce or recombine personal data from their training sets in their outputs and may generate realistic but false content about real individuals.⁹
- **Inferred data and AI-generated outputs:** AI systems may derive new information about individuals, including behavioural profiles, risk scores, and predictions, and generate outputs

such as classifications or decisions that significantly affect individuals. These inferences and outputs may remain opaque to data subjects despite influencing outcomes.

- **Emerging AI risks (bias, drift, and security):** AI systems may introduce risks such as proxy discrimination, feedback loops that reinforce disadvantage, degradation in accuracy over time (model drift), and exposure to attacks including data extraction, model inversion, or adversarial manipulation.

2.3. Scope and Purpose of Guidance Note

The purpose of this Guidance Note is to provide data controllers, data processors, AI developers, and deployers with an understanding of their obligations under the Data Protection Act and the attendant Regulations when developing, deploying, or procuring AI systems that process personal data of persons located in Kenya.

This Guidance Note applies to all entities, whether public or private, that:

- develop AI systems that are trained on personal data,
- deploy AI systems that process personal data in producing outputs, recommendations, or decisions,
- procure AI-as-a-service products or platforms that involve the processing of personal data, or
- use AI systems in automated decision-making that affects the rights or interests of data subjects.

This Guidance Note should be regarded as a minimum standard which can be supplemented by additional measures. Entities operating in regulated sectors should additionally comply with sector-specific data protection guidance issued by this Office and applicable sector regulators.

This Guidance Note considers:

- a) The Data Protection Act, Cap 411C,
- b) The Data Protection (General) Regulations, 2021,
- c) The Data Protection (Registration of Data Controllers and Data Processors) Regulations, 2021,
- d) The Data Protection (Complaints Handling Procedure and Enforcement) Regulations, 2021,
- e) The Constitution of Kenya, 2010,
- f) Other applicable laws and regulations, and
- g) International best practices and comparative AI governance frameworks.

3. LEGISLATIVE FRAMEWORK

The development, deployment, and use of AI systems in Kenya is governed by a layered legislative framework. The following instruments are of primary relevance:

a) The Constitution of Kenya, 2010:

Article 31 guarantees every individual the right to privacy, including in respect of information relating to their family or private affairs.

Article 27 guarantees equality and non-discrimination, providing a constitutional basis for challenging AI systems that produce discriminatory profiling or decisions.

Article 46 guarantees consumers' rights to safe and transparent goods and services, and provides for redress, including compensation, where loss or injury arises from defective or harmful outcomes. Article 35 guarantees access to one's personal data held by another person, requiring disclosure where rights are affected, and allowing individuals to correct or delete inaccurate data

b) The Data Protection Act, Cap 411C:

The Act is the primary statutory instrument governing AI processing of personal data. Key provisions include:

- Designation of a Data Protection Officer (DPO) for entities whose core activities involve large-scale systematic monitoring or the large-scale processing of sensitive personal data, including many high-risk AI deployments.
- Application of the core data protection principles to all AI processing activities.
- Protection and facilitation of data subject rights.
- Transparency obligations, including informing data subjects about automated processing and profiling.
- Lawful bases for AI processing, including consent, contract, legal obligation, vital interests, public task, and legitimate interests.
- Requirement to conduct a Data Protection Impact Assessment (DPIA) for high-risk AI processing, including automated decision-making and large-scale profiling.
- Enhanced protection of children's personal data, including parental or guardian consent and age-appropriate safeguards for AI systems accessible to children.
- Protection against solely automated decision-making that significantly affects a data subject, including the right to meaningful human review and an explanation of the logic underlying the decision.
- Right to object to the processing of personal data, including processing for AI profiling.
- Right to data portability, where AI systems rely on personal data provided by the data subject.
- Restrictions on cross-border transfers of personal data, including transfers involving AI cloud services and offshore AI providers.
- Enforcement powers of the Data Commissioner, including issuing enforcement notices, imposing administrative penalties, and ordering the cessation of unlawful AI-related processing.

c) The Data Protection (General) Regulations, 2021

These Regulations elaborate the principles and obligations under the Act. They set out the content requirements for privacy notices, the conditions for valid consent, the requirements for Data Protection Impact Assessments, and the obligations of data processors engaged through contractual arrangements.

All these obligations apply directly to AI processing activities.

d) The Data Protection (Registration of Data Controllers and Data Processors) Regulations, 2021

All entities deploying or operating AI systems that process personal data of data subjects in Kenya shall register with the Office as data controllers or data processors, subject to applicable thresholds and exemptions.

e) The Data Protection (Complaints Handling Procedure and Enforcement) Regulations, 2021

These Regulations set out the procedure for lodging, admitting, and determining complaints, the conduct of investigations, and the enforcement mechanisms available to the Office, including enforcement and penalty notices. They apply fully to AI systems that process personal data, including complaint handling, investigations, and enforcement of data protection obligations arising from AI-driven processing and decisions.

f) Other Applicable Laws

Entities deploying AI in specific sectors shall additionally comply with sector-specific legislation, in sectors such as health, payment systems, financial sector, communication, education, among others.

4. APPLICATION OF DATA PROTECTION PRINCIPLES TO AI SYSTEMS

All processing of personal data by or in connection with AI systems shall comply with the data protection principles set out in Section 25 of the Act. The following sub-sections address the specific application of each principle to AI contexts.

4.1. Lawfulness, Fairness, and Transparency

AI systems shall process personal data lawfully, which requires identifying and relying on a valid lawful basis under Section 30 of the Act.

Fairness requires that AI systems do not produce outcomes that arbitrarily disadvantage data subjects or use their data in ways that would be unexpected or harmful.

Transparency requires that the use of AI, the categories of personal data involved, and the logic of automated processing be disclosed to data subjects in accessible, plain-language terms.

Entities deploying AI systems shall ensure that their privacy notices specifically disclose:

- that personal data is processed by automated or AI-assisted means;
- the categories of data used as inputs;
- the nature of the outputs, inferences, or decisions produced; and
- the significance and envisaged consequences of such processing for the data subject.

Example

Jua Kali Credit, a digital lender, uses an AI-based credit scoring model that analyses mobile money transaction history and device usage patterns to determine loan eligibility. To comply with the lawfulness, fairness, and transparency principle, the company:

- identifies and documents a valid lawful basis before deployment (performance of the loan contract for application data, and a documented Legitimate Interests Assessment for fraud checks);

- ensures fairness by testing the model to confirm that it does not produce unjustified adverse outcomes for groups of applicants and that the data used is within applicants' reasonable expectations; and
- ensures transparency through a privacy notice that discloses the use of AI scoring, the categories of data used as inputs, the significance of the score for lending decisions, and the applicant's right to seek human review of an adverse automated decision under Section 35 of the Act.

4.2. Purpose Limitation

Personal data collected for a specific purpose shall not subsequently be used to train AI models for a different, incompatible purpose. Entities shall assess whether the use of existing datasets for AI training is compatible with the original collection purpose, considering the nature of the data, the reasonable expectations of data subjects, and the potential consequences for data subjects of further processing.

Where personal data is to be used to train AI models, entities should identify and document a specific, lawful purpose for that training activity. Repurposing production or operational data for AI training without fresh legal basis or compatibility assessment is not permitted.

Example

A public hospital collects patient records under a lawful basis of public interest for the delivery of healthcare services. The hospital cannot, without further assessment and legal authority, supply those records as training data to a commercial AI vendor developing a diagnostic tool for commercial licensing. The use of the data for commercial AI development is incompatible with the original public healthcare purpose.

4.3. Data Minimisation

AI systems shall be designed and operated to use only personal data that is adequate, relevant, and necessary to achieve the defined purpose. Entities shall resist the pressure to collect and process ever-larger datasets on the assumption that more data will produce better AI outcomes.

Privacy by design requires that data minimisation be addressed at the AI system design stage, not retrospectively.

Where aggregated, pseudonymised, or synthetic data can achieve the same training or operational purpose as full personal data, entities should prefer those less privacy-invasive approaches.

Example

A digital lender builds an AI model for credit scoring but collects additional data such as contact lists and social media activity that are not necessary for assessing creditworthiness. This violates the data minimization principle.

The lender should restrict data collection to relevant financial and transactional information.

4.4. Accuracy

AI systems can amplify inaccuracies in training data and generate inaccurate inferences at scale. Entities shall take reasonable steps to ensure that personal data used to train AI models is accurate and up-to-date and shall implement mechanisms to detect and correct inaccurate outputs that affect data subjects.

Erroneous outputs, particularly in credit scoring, health diagnosis, or fraud detection can cause serious harm. Entities shall have processes to rectify inaccurate AI-generated information about data subjects upon request under Section 40 of the Act.

Example

An insurance company uses an AI system trained on outdated data, resulting in incorrect risk assessments and higher premiums for some customers.

This undermines the accuracy principle. The company should regularly update data and allow correction of inaccurate information affecting decisions.

4.5. Storage Limitation

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 anonymisation 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.

Example

A telecommunications company retains personal data used to train an AI model even after the model has been deployed and the data is no longer needed.

This contravenes the storage limitation principle. The company should delete or anonymise the data once the training purpose is fulfilled.

4.6. Integrity and Confidentiality

AI systems shall incorporate appropriate technical and organisational security measures to protect personal data against unauthorised access, loss, or destruction at all stages of the AI lifecycle, including data collection, preprocessing, training, inference, and storage of outputs. Entities shall specifically address the security of:

- training datasets and the systems used to store and process them;
- model parameters and weights, which may be subject to model inversion or membership inference attacks capable of reconstructing training data;
- inference APIs and the personal data transmitted through them; and

- logs of AI-generated outputs that may contain or relate to personal data.

Example A healthcare provider uses an AI system to analyze patient data, but weak access controls allow unauthorized access to sensitive records.

This breaches the integrity and confidentiality principle. The provider should implement strong access controls and security measures to protect the data.

4.7. Accountability

Entities deploying AI systems bear accountability for ensuring that the system complies with data protection law throughout its lifecycle. Accountability obligations in the AI context include:

- maintaining records of processing activities that document AI systems and their data inputs;
- conducting DPIAs before deploying high-risk AI applications;
- implementing governance structures that assign clear responsibility for AI data protection compliance; and
- monitoring deployed AI systems on an ongoing basis for drift, bias, or emerging privacy risks.

Example

A financial institution deploys an AI fraud detection system without documenting the data used, system design, or measures to ensure fairness.

This undermines accountability. The institution should maintain records, conduct assessments, and implement governance structures to demonstrate compliance

5. LAWFUL BASIS FOR PROCESSING PERSONAL DATA IN AI SYSTEMS

Every processing operation conducted by or in connection with an AI system shall have a valid lawful basis under Section 30 of the Act. Entities shall identify the lawful basis before processing commences and shall document that identification. The following sub-sections address the application of each lawful basis to AI processing.

5.1. Consent

Consent, under the Act, shall be freely given, specific, informed, and unambiguous.

In the AI context, consent provides a valid lawful basis only where it is obtained in relation to the specific AI processing and decision-making concerned. Before giving consent, the subject of data shall be provided with clear information about:

- the use of AI;
- the categories of data processed by the AI system; and
- the nature of any automated decision-making or profiling.

Transparency obligations under Section 29 of the Act apply irrespective of the lawful basis relied upon; the provision of a privacy notice does not, of itself, constitute consent. Where consent is relied upon, entities shall be able to demonstrate that it was freely given, specific, informed, and unambiguous in respect of the AI processing, including any automated decision-making or profiling.

Data subjects shall be informed of their right to withdraw consent at any time, and withdrawal shall not result in detriment.

Consent is not an appropriate lawful basis where there is a significant power imbalance between the entity and the data subject (for example, employer-employee AI monitoring) or where consent would need to be given as a condition of access to a service.

5.2. Contract

AI processing may be based on contractual necessity where the processing is objectively necessary for the performance of a contract with the data subject or to take steps at the request of the data subject prior to entering into a contract.

A lender using AI to assess a loan application may rely on contract as a lawful basis for processing the applicant's financial data. The contract basis does not, however, extend to all ancillary AI processing that the entity finds convenient, only to processing that is genuinely necessary for the contract.

5.3. Legal Obligation

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.

5.4. Vital Interests

AI systems in emergency health contexts or disaster response may process personal data based on vital interests where the processing is necessary to protect the life of a data subject or another natural person. This basis should be relied upon only where other bases are not available, and processing is genuinely necessary to protect life.

5.5. Public Interest or Exercise of Official Authority

Public sector entities and entities performing public functions may rely on public interest or official authority as a lawful basis for AI processing that is necessary for the performance of those functions.

AI-enabled service delivery by government agencies, AI-driven benefits eligibility assessment, and AI-based land records management are examples. Public bodies shall nonetheless adhere to all other data protection principles and obligations, including transparency and DPIA requirements.

5.6. Legitimate Interests

Private sector entities may rely on legitimate interests where the processing is necessary for the purposes of the legitimate interests pursued by the data controller or a third party, provided those interests are not overridden by the interests or fundamental rights of the data subject.

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

- balances the interest against the data subject's rights, considering the intrusiveness of AI profiling and automated decision-making.

Legitimate interests require careful assessment and may not be appropriate for certain high-risk AI applications, such as those that profile data subjects at scale, involve sensitive personal data, or produce significant automated decisions affecting data subjects' legal rights or opportunities, depending on the nature of the processing and its impact on data subjects.

6. RIGHTS OF DATA SUBJECTS IN AI CONTEXTS

The Act provides data subjects with enforceable rights that apply with full force to AI systems processing their personal data. Entities deploying AI systems shall implement technical and procedural mechanisms to give effect to each of the following rights.

The rights of data subjects are not absolute and may be limited in accordance with the Act. Entities shall give effect to data subject requests within the timelines prescribed under the Act and the Data Protection (General) Regulations, 2021. Where a request is declined, the entity shall notify the data subject of the refusal in writing, with reasons, within the prescribed period.

6.1. Right to Be Informed

Under Section 29 of the Act, data subjects shall be informed at the point of data collection of the use of AI systems, the logic of automated processing, the significance of any automated decisions, and the envisaged consequences of the processing for the data subject.

Privacy notices shall address AI processing in terms that are intelligible to a lay person.

Entities shall not rely on generic or technical language that obscures the nature of AI decision-making.

6.2. Right of Access

Under Section 26 of the Act, data subjects are entitled to access the personal data held about them, including data used as inputs to AI systems and the outputs, inferences, or scores generated about them.

Entities shall be able to supply data subjects with information about what data is held and how it is used in AI processing, including, where processing is automated, the categories of data used and the logic applied in the decision or output.

6.3. Right to Rectification

Under Section 40 of the Act, data subjects have the right to have inaccurate personal data rectified. In the AI context, this includes the right to correct inaccurate input data and, where an AI output is based on inaccurate data, to have the output reviewed and corrected.

Entities shall implement processes to promptly action rectification requests and to propagate corrections through AI systems where the inaccurate data has been used as an input.

6.4. Right to Erasure

Where the conditions for erasure under the Act are met, including where consent is withdrawn, where the data is no longer necessary for the purpose for which it was collected, or where processing is unlawful, entities shall erase personal data from training datasets and implement appropriate measures to give effect to erasure requests.

Where erasure from trained model parameters is not technically feasible, entities shall document the reasons and the basis for the feasibility assessment, assess and implement alternative mitigation measures, and maintain documented procedures for managing such requests.

Entities shall address the technical challenges of erasure from AI models in their data protection governance frameworks.

6.5. Right to Object

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.

Entities shall give effect to objections promptly and shall not continue profiling a data subject who has objected, unless compelling legitimate grounds for the processing override the interests of the data subject.

6.6. Right Not to Be Subject to Automated Decision-Making

Section 35 of the Act provides that a data subject has the right not to be subject to a decision based solely on automated processing, including profiling, which significantly affects the data subject. This right is of central importance in the AI context.

Entities whose AI systems produce decisions that significantly affect data subjects, including decisions on credit, employment, insurance, housing, healthcare, education, or access to public services, shall:

- ensure that such decisions are not made solely by automated means without meaningful human involvement;
- inform data subjects that a decision has been made through automated processing;
- provide data subjects with the right to request human review of any automated decision that significantly affects them;
- provide data subjects, upon request, with a meaningful explanation of the logic of the automated decision and its consequences; and
- implement mechanisms for data subjects to contest automated decisions and to have the contest considered by a human with authority to reverse or modify the decision.

In line with Regulation 22 of the Data Protection (General) Regulations, 2021, entities whose AI systems produce decisions based solely on automated processing shall additionally:

- ensure the prevention of errors in the automated decision-making;
- use appropriate mathematical or statistical procedures for the profiling or automated processing;
- put in place appropriate technical and organisational measures to correct inaccuracies and minimise the risk of errors; and
- process personal data in a manner that eliminates discriminatory effects and bias.

Example

Sokoni Insurance uses an AI model to assess claims and automatically decline those it classifies as high-risk for fraud. A claimant whose claim is declined by the AI model has the right under Section 35 of the Act to request human review of the decision.

Sokoni shall:

- inform the claimant that the decision was made by an automated system;
- explain the principal factors that influenced the AI's determination; and
- ensure a qualified human reviewer examines the claim afresh before a final decision is communicated.

Example 2

A public hospital uses an AI system to prioritize patients for specialized treatment. Patients assessed as low priority are automatically deferred without further review.

A patient is not informed that the decision was made using AI, is not given access to the data or reasoning behind the outcome, and is not provided with a way to request human review.

This violates the rights of data subjects. The hospital fails to:

- inform the patient of the use of AI (Right to be Informed);
- provide access to data and inferences used (Right of Access); and
- allow the patient to request human review of a significant decision (Right not to be subject to automated decision-making).

6.7. Right to Data Portability

Under Section 38 of the Act, data subjects have the right to receive personal data they have provided to a controller in a structured, commonly used, and machine-readable format, and to transmit that data to another controller.

In the AI context, entities that generate personalised AI outputs based on data provided by the data subject should consider whether those outputs fall within the scope of the portability right and ensure they have the technical capability to produce portable data exports.

Example

A university uses an AI-powered learning platform to track student performance and generate personalized learning profiles based on coursework, assessments, and engagement. A student requests a copy of their data to transfer to another institution.

The university declines to provide the data, stating that the performance analytics and profiles are internal.

This violates the right to data portability. The university should provide the student with their personal data, including submitted coursework, assessment results, and learning profiles—in a structured, commonly used, and machine-readable format to enable transfer to another institution.

7. OBLIGATIONS OF DATA CONTROLLERS AND DATA PROCESSORS IN AI SYSTEMS

7.1. Registration

All entities that develop, deploy, or operate AI systems that process personal data of persons located in Kenya shall register with the Office as data controllers or data processors, subject to the thresholds and procedures set out in the Data Protection (Registration of Data Controllers and Data Processors) Regulations, 2021.

The Office has published a Guidance Note on Registration of Data Controllers and Data Processors which is accessible at www.odpc.go.ke.

Entities procuring AI-as-a-service products from third-party vendors shall assess whether the vendor is acting as a data processor and, if so, shall ensure the vendor is registered as required by the Act and enters into a compliant data processing agreement.

7.2. Privacy by Design and Default

Privacy by design and default, recognised under Section 41 of the Act and the General Regulations requires that data protection considerations be embedded into AI systems from the earliest stages of design and development, not addressed as an afterthought or add-on.

In the AI context, privacy by design obligations include:

- data minimisation at the data collection and feature selection stage;
- designing AI models to operate pseudonymised or aggregated data wherever feasible;
- implementing differential privacy or other privacy-enhancing techniques in AI training pipelines where technically appropriate;
- building in explainability and human review mechanisms as core system features;
- designing AI systems to default to the most privacy-protective settings available; and
- conducting privacy impact assessments at the system design stage.

7.3. Data Protection Impact Assessment (DPIA)

Section 31 of the Act and the General Regulations require a DPIA before commencing any processing activity that is likely to result in high risk to the rights and freedoms of data subjects.

AI systems that involve large-scale profiling, systematic monitoring, automated decision-making with significant effects, use of sensitive personal data, or processing of children's data will typically trigger the DPIA requirement.

A DPIA for an AI system shall address:

- a description of the AI system, its purpose, the personal data processed, and the envisaged outputs and decisions;
- an assessment of the necessity and proportionality of the processing in relation to its purpose;

- an assessment of the risks to data subjects, including discriminatory outputs, re-identification, data security risks, and the risk of erroneous or unexplainable automated decisions; and
- the measures proposed to address those risks, including technical controls, governance mechanisms, and data subject rights mechanisms.

The Office has published a Guidance Note on Data Protection Impact Assessments accessible at www.odpc.go.ke. Entities shall consult that Guidance Note when conducting DPIAs for AI systems.

7.4. Data Protection Officer (DPO)

The Act, under Section 24, provides that a data controller or data processor may designate or appoint a Data Protection Officer. The designation or appointment of a DPO is strongly advisable for entities whose core activities involve large-scale systematic monitoring of data subjects or large-scale processing of sensitive personal data.

Many AI deployments, including large-scale AI-driven customer monitoring, health AI, and AI-based credit assessment platforms make the designation or appointment of a DPO particularly advisable. The DPO shall be involved in all matters relating to the protection of personal data in AI systems, including DPIA processes, AI system procurement decisions, and data subject rights responses.

7.5. Sensitive Personal Data and AI

The processing of sensitive personal data, including health data, biometric data, racial or ethnic origin, and data revealing sexual orientation in AI systems requires explicit consent or another specific legal basis permitting such processing.

Entities shall not use sensitive personal data as AI training data or as inference inputs without a clear legal basis and without implementing heightened technical and organisational safeguards.

AI systems that generate inferences about sensitive personal data categories, for example, systems that infer health status or ethnicity from non-sensitive inputs may themselves produce sensitive personal data.

Entities shall assess whether AI outputs constitute sensitive personal data and apply the corresponding legal standards.

7.6. Children and AI Systems

Section 33 of the Act provides enhanced protections for children's personal data.

Entities deploying AI systems that are likely to be accessed by or that process data about children shall implement age-appropriate design measures, including age verification mechanisms, and shall obtain verifiable parental or guardian consent for the processing of children's data.

AI systems shall not be used to profile children for commercial purposes or to make automated decisions that significantly affect children without appropriate human oversight and the involvement of the child's parent or guardian.

7.7. Notification and Communication of a Breach

Entities shall notify the Office of any personal data breach affecting AI systems, including unauthorised access to training datasets, model extraction attacks, or adversarial manipulation of AI outputs within 72 hours of becoming aware of the breach, where the breach is likely to result in risk to data subjects. Where the breach is likely to result in high risk to data subjects, affected individuals shall also be notified without undue delay. Where a personal data breach occurs within a data

processor's systems, including those of a third-party AI service provider, the data processor shall notify the data controller within forty-eight hours of becoming aware of the breach, and the data controller remains responsible for notification to the Office and, where required, communication to affected data subjects.

Example

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; 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.

7.8. Data Sharing and Third-Party AI Processors

Where an entity engages a third party to provide AI services involving the processing of personal data, including cloud-based AI platforms, AI-as-a-service providers, and AI model vendors, the third party may act as a data processor, a data controller, or a joint controller depending on the circumstances of the processing, including whether it determines the purposes and means of the processing.

Entities shall assess and document the role of each third party in the AI ecosystem. Where the third party acts as a data processor, the entity shall enter into a written data processing agreement with the AI processor that:

- restricts the processor to processing personal data only on the instructions of the controller;
- requires the processor to implement appropriate technical and organisational security measures;
- prohibits the processor from engaging sub-processors without prior authorisation;
- requires the processor to assist the controller in fulfilling data subject rights requests and DPIA obligations; and
- requires the processor to delete or return all personal data upon termination of the agreement.

7.9. Cross-Border Transfers of AI-Processed Data

Many AI systems involve cross-border data transfers, including the transfer of training data to offshore model developers, the use of cloud-based AI inference services, and the sharing of AI outputs with offshore recipients. Part VI of the Act restricts transfers of personal data outside Kenya to jurisdictions that provide adequate data protection, or where appropriate safeguards are in place.

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.

Example

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.

8. DATA PROTECTION AND THE AI LIFECYCLE

An AI system does not come into existence at a single moment: it passes through a series of distinct stages from initial conception to eventual decommissioning. Personal data may be involved at every one of these stages, and different data protection obligations attach to each stage. Entities shall therefore approach compliance not as a one-time exercise at the point of deployment, but as a continuous obligation that tracks the entire lifecycle of the AI system.

This section maps the principal data protection obligations of the Act and the General Regulations to each stage of the AI lifecycle. It applies to both entities that develop AI systems internally and those that procure AI systems from third-party vendors, who shall apply equivalent scrutiny to the procurement and integration stages.

AI Lifecycle Stage	Principal Data Protection Obligations
--------------------	---------------------------------------

Stage 1: Problem Definition and AI System Design	DPIA scoping; privacy by design; lawful basis identification; DPO engagement
Stage 2: Data Collection and Preparation	Lawfulness, fairness, transparency; data minimisation; purpose limitation; consent or other lawful basis for training data; data quality controls
Stage 3: Model Training and Development	Purpose limitation; storage limitation on training data; security of training pipelines; bias assessment; pseudonymisation or anonymisation where feasible
Stage 4: Testing, Validation, and Pre-Deployment Review	Full DPIA completion; accuracy and fairness testing; explainability review; human review mechanism testing; data subject rights mechanism testing
Stage 5: Deployment and Integration	Privacy notice update; registration with the Office; data processor agreements; cross-border transfer assessment; breach notification procedures activation
Stage 6: Operation, Monitoring, and Maintenance	Ongoing accuracy and fairness monitoring; data subject rights fulfilment; breach detection and notification; periodic DPIA review; audit trail maintenance
Stage 7: Decommissioning and Data Disposal	Secure deletion of training datasets and operational data; erasure from model parameters where feasible; retention schedule compliance

8.1. Stage 1: Problem Definition and AI System Design

Data protection obligations begin before any personal data is processed. At the problem definition and system design stage, entities shall:

- define the purpose of the AI system with sufficient specificity to enable lawful basis identification and DPIA scoping;
- engage the DPO at the earliest practicable stage, where one has been designated;
- identify whether the proposed AI system falls within a high-risk category requiring a mandatory DPIA;
- determine, at the architectural level, whether the AI system can be designed to operate without personal data, or with minimised, pseudonymised, or synthetic data; and
- embed privacy by design requirements in the AI system specification, including explainability, human review mechanisms, and data subject rights interfaces.

The privacy-by-design principle requires that data protection is not retrofitted after system design is complete but is a foundational design constraint from the outset.

Entities that procure AI systems from vendors shall include privacy-by-design specifications in procurement requirements and shall not accept AI systems that cannot demonstrate compliance with the Act.

Example

Fahari Health, a private hospital network, plans to deploy an AI clinical decision-support system to assist physicians with patient triage. At the design stage, the hospital's DPO is engaged to assess whether the system can operate on pseudonymised patient records rather than fully identified records, and to scope a DPIA.

The AI vendor is required to provide documentation of the model architecture, the categories of data required for inference, and the explainability mechanisms built into the system before the procurement contract is signed.

8.2. Stage 2: Data Collection and Preparation

The collection and preparation of data for AI training is a processing activity governed in full by the Act. Entities shall:

- identify the lawful basis for collecting and processing personal data as training data, applying the same rigour required for any other processing activity;
- ensure that personal data used as training data was collected with a purpose that is compatible with the AI training purpose, or obtain fresh legal authority;
- apply data minimisation: collect the minimum categories and volume of personal data sufficient for the training purpose, and prefer anonymised, pseudonymised, or synthetic data where the training objective can be met without full personal data;
- implement data quality controls to ensure training data is accurate, representative, and free from systematic biases that would produce discriminatory model outputs;
- maintain records identifying the source of training datasets, the categories of personal data they contain, the lawful basis relied upon, and the data subjects' reasonable expectations at the time of original collection; and
- where training data is sourced from third parties, conduct due diligence to ensure that the third party collected the data lawfully and that its transfer to the entity for AI training is permissible.

Entities shall not use personal data scraped from the internet or aggregated from public sources as AI training data without assessing whether the collection and use of that data for training purposes is lawful under the Act. The fact that personal data is publicly accessible does not, of itself, provide a lawful basis for processing it as AI training data.

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 minimisation requires use of anonymised or synthetic text data instead; and whether a DPIA is required given the scale of the processing.

8.3. Stage 3: Model Training and Development

During model training, personal data is processed by computational systems that learn statistical patterns from it. Entities shall:

- implement appropriate technical and organisational security measures for training infrastructure, including access controls, encryption of training datasets at rest and in transit, and audit logging of access to training data;
- apply storage limitation: training datasets shall not be retained beyond the period necessary for the training activity; upon completion of training, training data shall be deleted or anonymised unless a separate lawful basis for continued retention exists;
- conduct bias testing during model development to detect and mitigate the risk of discriminatory outputs, particularly where training data reflects historical patterns of discrimination based on sensitive personal data categories;
- apply privacy-enhancing techniques, including differential privacy, federated learning, or secure multi-party computation where technically feasible and appropriate to the scale and sensitivity of the training data; and
- document the model development process, including the training data used, the architecture of the model, the training parameters, and the results of bias and accuracy testing.

Where AI model development is performed by a third-party developer on behalf of the data controller, the developer is acting as a data processor. A compliant data processing agreement shall be in place before any personal data is shared with the developer for training purposes.

8.4. Stage 4: Testing, Validation, and Pre-Deployment Review

Before an AI system is deployed in a live environment where it will process personal data of actual data subjects or produce decisions affecting them, entities shall complete a pre-deployment review that includes:

- finalisation and sign-off of the DPIA, including documentation of the residual risks assessed as acceptable and the measures adopted to mitigate risks assessed as unacceptable;
- technical testing of the human review mechanism required under Section 35 of the Act to ensure data subjects can effectively request and receive human review of automated decisions;
- testing of data subject rights interfaces, including access request fulfilment, rectification workflows, and erasure mechanisms;
- accuracy and fairness validation using test datasets that are representative of the real-world population that will be subject to the AI system's outputs, with specific attention to differential performance across demographic groups;
- explainability review: verification that the AI system can generate intelligible explanations of its principal outputs and decisions in terms accessible to affected data subjects; and
- DPO sign-off, where a DPO has been designated, confirming that data protection obligations have been assessed and addressed.

Where testing requires the use of personal data, entities shall ensure that a lawful basis exists for processing that data for testing purposes and shall apply the same data minimisation and security obligations that apply to training data. The use of synthetic or anonymised test data should be preferred where it can provide meaningful test results.

8.5. Stage 5: Deployment and Integration

At the point of deployment, entities shall ensure that all compliance infrastructure is operational and that the AI system is integrated into the entity's broader data protection governance framework. Deployment obligations include:

- updating privacy notices to disclose the deployment of the AI system, the categories of personal data processed, the nature of any automated decision-making or profiling, and data subject rights;
- confirming registration with the Office, or updating an existing registration to reflect the new AI processing activity;
- executing data processing agreements with all third-party AI service providers, cloud infrastructure providers, and any other processors engaged in connection with the AI system;
- assessing and documenting the lawful basis for any cross-border transfers of personal data made in connection with the AI system, and implementing contractual safeguards where required;
- activating breach detection, logging, and notification procedures specific to the AI system; and
- conducting staff training on the AI system's data protection obligations, including how to handle data subject rights requests related to the AI system.

Example

Harambee Bank deploys an AI-driven fraud detection system that monitors customer transactions in real time and automatically flags and suspends accounts suspected of fraudulent activity. Prior to deployment:

- the bank updates its privacy notice to disclose the AI monitoring system;
- signs a data processing agreement with the AI vendor;
- conducts a cross-border transfer assessment (the AI vendor processes data on servers in the European Union); and
- trains customer-facing staff on how to handle customer enquiries about automated account suspensions, including how to facilitate human review requests under Section 35 of the Act.

8.6. Stage 6: Operation, Monitoring, and Maintenance

Once deployed, AI systems shall be subject to ongoing monitoring and governance. Data protection compliance is not a point-in-time obligation: AI systems change over time through retraining, fine-tuning, and exposure to new data distributions, and the risk profile of an AI system may change materially without a new deployment event. Operational obligations include:

- ongoing monitoring of the AI system's outputs for accuracy, fairness, and consistency with the stated purpose, including detection of model drift and discriminatory output patterns;
- periodic review and, where necessary, revision of the DPIA, at intervals proportionate to the risk profile of the AI system and triggered by any material change in the system, the data processed, or the deployment context;
- maintenance of records of processing activities, including logs of AI system decisions and the data inputs used, to the extent necessary to demonstrate accountability and to respond to data subject rights requests and regulatory enquiries;
- fulfilment of data subject rights requests relating to the AI system within the statutory timeframes, including access requests, rectification requests, erasure requests, and requests for human review of automated decisions;
- prompt detection, containment, and notification of personal data breaches involving the AI system;
- periodic algorithmic auditing by a suitably qualified internal or independent auditor to assess model performance, bias, and compliance with data protection obligations;
- maintenance of a documented AI change management procedure covering model retraining, material changes to training datasets, deployment of the AI system in new contexts, and changes of third-party AI providers, with reassessment of privacy risks and review of the DPIA before material changes take effect; and
- maintenance of a documented AI incident management and remediation procedure, including a register of AI-related incidents such as discriminatory outcomes, systemic inaccuracies, and privacy harms, together with the remedial action taken and the notifications made.

Entities shall establish an AI governance committee or equivalent body with clear responsibility for monitoring AI system compliance on an ongoing basis. The DPO, where designated, shall be included in that governance structure with appropriate access to AI system documentation and outputs.

8.7. Stage 7: Decommissioning and Data Disposal

When an AI system is retired or replaced, entities shall ensure that the decommissioning process addresses all personal data processed in connection with the system. Decommissioning obligations include:

- secure deletion or anonymisation of all personal data held in training datasets, test datasets, and operational data stores associated with the AI system, in accordance with the entity's documented retention schedule;
- assessment of whether trained model weights or parameters encode personal data in a form that is recoverable through model inversion or membership inference techniques, and implementation of appropriate erasure measures where this risk is present;
- deletion or secure archiving of logs of AI system outputs that contain or relate to personal data, subject to any statutory minimum retention obligations;
- termination or amendment of data processing agreements with third-party AI processors upon cessation of the processing relationship, and confirmation of data deletion by those processors;
- updating the records of processing activities to record the cessation of the AI processing activity; and

- issuing a decommissioning notice to the Data Protection Officer and, where the cessation of processing is material, informing the Office.

Where an AI system is replaced by a successor system, entities shall treat the successor system as a new AI deployment and apply the full lifecycle compliance process from Stage 1, including a fresh DPIA assessment that addresses any differences in architecture, data processing, or deployment context between the predecessor and successor systems.

9. HIGH-RISK AI APPLICATIONS: SPECIFIC OBLIGATIONS

Certain AI applications, by reason of the nature of the processing, the categories of data involved, or the significance of the decisions produced, pose heightened risks to data subjects and attract additional obligations.

The following categories are considered high-risk for the purposes of this Guidance Note and require, at a minimum:

- mandatory DPIA;
- DPO involvement, subject to Section 24 of the Act and applicable legal requirements;
- meaningful human review of automated decisions; and
- regular algorithmic audit.

High-Risk AI Application	Primary Risks	Specific Obligations
AI Credit Scoring and Financial Decisioning	Discriminatory profiling; exclusion of vulnerable groups; opaque adverse decisions on loan eligibility, limits, and terms	DPIA required; Section 35 human review mechanism mandatory; explanation of principal scoring factors to be provided upon request; regular bias audit of model outputs
AI in Health Diagnostics and Clinical Decision Support	Erroneous clinical recommendations; processing of health data at scale; re-identification from de-identified clinical datasets	DPIA required; Health data processed under explicit consent or specific legal basis; DPO mandatory; compliance with Digital Health Act.
AI-Enabled Biometric Recognition (Facial, Fingerprint, Voice)	Mass surveillance; re-identification; processing of sensitive biometric data at scale	DPIA required; Explicit consent or statutory basis; DPO mandatory; strict access controls; prohibition on real-time biometric surveillance absent specific legal authority
AI in Employment Screening and Monitoring	Discriminatory filtering of candidates; covert monitoring; automated disciplinary decisions	DPIA required; Disclosure to candidates and employees of AI screening tools; human review of AI-recommended hiring or termination decisions; right to contest
AI in Education (Adaptive Learning,	Profiling of children; automated grading or progression decisions; long-term data retention	Section 33 children's data protections apply; parental consent required; no commercial

Assessment, Student Monitoring)		profiling of children; DPIA required
Generative AI Platforms (LLMs, Image Generation)	Training data privacy; reproduction of personal data in outputs; generation of misleading content about real persons	Privacy notice to disclose generative AI processing; training data audit for personal data compliance; output monitoring for personal data reproduction
AI in Law Enforcement and Public Safety	Discriminatory targeting; automated decisions affecting liberty; surveillance at scale	Specific statutory authority required; DPIA mandatory; independent oversight; restrictions on automated profiling for law enforcement purposes

Case Study – AI Credit Scoring: Comparative Benchmark

A Kenyan digital lender deploys an AI credit scoring model using mobile transaction data and social metadata. Under the DPA, the lender shall:

- disclose the AI scoring system in its privacy notice;
- identify a lawful basis for the profiling;
- conduct a DPIA;
- provide data subjects with the right to human review under Section 35; and
- be able to explain the principal factors in an adverse score on request.

10. PRE-PROCESSING CHECKLIST FOR AI SYSTEMS

Before commencing any processing of personal data using AI systems, organizations shall complete the following checklist. Where any requirement is not met, the processing shall not proceed until the deficiency is addressed, or appropriate safeguards are implemented and documented. Where high risks remain, the organization shall conduct a Data Protection Impact Assessment and, where required, seek prior consultation with the Office of the Data Protection Commissioner before proceeding.

Compliance Area	Key Questions / Actions
Lawful Basis and Purpose	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 original purpose or supported by a new lawful basis?
Data Identification and Classification	What categories of personal data will be processed? Does the system involve sensitive personal data? Will it process data relating to children or vulnerable groups? Will inferred data or AI outputs constitute personal data?

Data Minimization and Design	Are you collecting only data that is necessary? Can anonymized, pseudonymized, or synthetic data be used? Has privacy by design been incorporated into the system?
Risk Assessment (DPIA)	Has a Data Protection Impact Assessment been conducted where required? Have risks such as bias, discrimination, re-identification, and security threats been identified? Are mitigation measures defined?
Transparency and Data Subject Rights	Have privacy notices been updated to disclose AI processing? Are mechanisms in place to support access, rectification, and human review? Can meaningful explanations of AI decisions be provided?
Governance and Accountability	Has responsibility for AI oversight been assigned? Is the Data Protection Officer involved where required? Are records of processing activities maintained?
Third-Party Processing	Will third-party AI providers process personal data? Are data processing agreements in place with appropriate safeguards?
Cross-Border Transfers	Will personal data be transferred outside Kenya? Has adequacy been assessed and appropriate safeguards implemented?
Security and System Readiness	Are appropriate security measures in place (access control, encryption, monitoring)? Has the system been tested for accuracy, fairness, and reliability? Are breach detection and notification procedures established?

11. AI GOVERNANCE FRAMEWORK

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:

- the governance structures for AI oversight and reporting lines to senior management;
- the roles and responsibilities of each actor across the AI lifecycle, including the AI developer, AI provider, foundation model provider, AI deployer, and AI user, and the allocation of data controller, joint controller, and data processor responsibilities among them;
- the role of the DPO, where designated or appointed, in AI oversight, with appropriate access to AI system documentation and outputs;
- the standards and frequency for algorithmic audits, including the qualification and independence of the auditor, the matters to be assessed (model performance, bias, drift, and compliance with data protection obligations), and the documentation of findings and remediation;
- the standards for meaningful human oversight of automated decisions, including the authority, competence, and information available to human reviewers and the escalation procedures for contested decisions; and
- the AI change management and incident management procedures described in this Guidance Note.

12. GENERATIVE AI

Entities that develop, provide, procure, deploy, or otherwise use generative artificial intelligence systems that process personal data, including large language models (LLMs) and systems capable of generating text, images, audio, video, code, or other synthetic content, shall ensure that such systems are designed, developed, deployed, and operated in a manner that complies with the Data

Protection Act, the Data Protection Regulations, and the principles of privacy by design and by default.

In addition to the obligations set out in this Guidance Note, such entities should, where applicable:

- (a) identify and document the lawful basis for all personal data used to train, fine-tune, evaluate, or ground the generative AI system, including personal data obtained from publicly available sources, third-party datasets, user interactions, or other external sources;
- (b) implement appropriate technical and organisational measures to minimise the risk of the AI system reproducing, disclosing, or otherwise revealing personal data contained in training data, prompts, retrieved data, or system memory, including through appropriate testing before deployment and periodic review thereafter;
- (c) ensure that appropriate safeguards are implemented to reduce the risk of hallucinations, inaccurate personal information, or misleading outputs concerning identifiable individuals, particularly where such outputs may be relied upon for decisions affecting data subjects;
- (d) provide clear and accessible information to users and, where appropriate, data subjects regarding the use of generative AI, including where AI-generated content may reasonably be mistaken for content produced solely by a human or may materially affect the rights or interests of a data subject;
- (e) establish effective mechanisms through which data subjects may request access, correction, deletion, restriction, or other appropriate remedies where their personal data appears in AI-generated outputs or has otherwise been processed in a manner inconsistent with the Act;
- (f) assess and document, as part of the Data Protection Impact Assessment the specific privacy risks associated with the use of generative AI, including risks relating to memorisation, data leakage, re-identification, inference, profiling, inaccurate outputs, and the generation or disclosure of personal data;
- (g) maintain appropriate governance, oversight, and human review mechanisms to monitor the operation of the generative AI system, investigate privacy-related incidents, and implement corrective measures where necessary; and
- (h) where generative AI systems or foundation models are procured from or operated by third parties, ensure that appropriate contractual, technical, and organisational measures are in place to allocate responsibilities for compliance with the Data Protection Act and to safeguard the rights and freedoms of data subjects.

13. ENFORCEMENT AND COMPLAINTS

The Office is mandated to oversee the implementation and enforcement of the Act. The Office is empowered to receive and investigate complaints, conduct inspections and audits of data controllers and data processors, issue enforcement and penalty notices, impose administrative fines, and make referrals for criminal prosecution where applicable. These enforcement powers apply in full to entities that process personal data through AI systems.

Data subjects who believe that an entity has violated their data protection rights in connection with an AI system may lodge a complaint with the Office. Complaints may relate to, among other matters:

- failure to disclose AI processing in a privacy notice;
- use of AI in automated decision-making without providing a human review mechanism;
- failure to provide an explanation of an automated decision upon request;
- discriminatory AI profiling producing adverse outcomes;

- processing of sensitive personal data by an AI system without lawful basis;
- failure to conduct a DPIA for a high-risk AI application; or
- unlawful cross-border transfer of personal data to an AI processor.

Entities found to have violated the Act in connection with AI processing are liable to enforcement action, including administrative fines and civil remedies, in accordance with the Data Protection (Complaints Handling Procedure and Enforcement) Regulations, 2021.

The Office encourages entities to engage proactively with the Office on AI data protection questions, including by seeking advisory opinions and participating in ODPC-facilitated stakeholder engagement processes on AI governance. The Office may consider cooperation and proactive engagement among relevant factors during enforcement processes.

To lodge a complaint or seek further guidance, contact the Office at:

Office of the Data Protection Commissioner

Britam Towers, 12th Floor, Hospital Road, Upperhill, Nairobi

P.O. Box 30920-00100, Nairobi, Kenya

Email: complaint@odpc.go.ke (for lodging complaints) | info@odpc.go.ke / enquiries@odpc.go.ke (for enquiries and further guidance)

Telephone: 0207801800 | 0796954269 / 0752896867

Website: www.odpc.go.ke

ANNEX 1: COMPLIANCE CHECKLIST FOR AI SYSTEMS

The following checklist is a practical tool to assist entities in self-assessing compliance with data protection obligations applicable to AI systems. It does not constitute legal advice and does not substitute for a full DPIA or legal review.

Compliance Requirement	Action Required
Registration	Register as a data controller or processor with the ODPC at www.odpc.go.ke before deploying any AI system that processes personal data.
Privacy Notice	Update privacy notices to disclose: use of AI; categories of data processed; nature of automated decisions or profiling; significance and consequences for data subjects; and data subject rights, including the right to human review under Section 35.
Lawful Basis	Identify and document a lawful basis under Section 30 of the Act for each AI processing activity, including training, inference, and output generation.
Sensitive Personal Data	Identify whether the AI system processes sensitive personal data (including biometric, health, or inferred sensitive data). If so, ensure explicit consent or specific legal basis and implement heightened safeguards.
Children's Data	Assess whether the AI system may process children's data. If so, implement age verification, parental consent mechanisms, and age-appropriate design measures per Section 33 of the Act.
DPIA	Conduct a DPIA before deploying any high-risk AI application. Document the assessment and retain it for inspection by the Office. Refer to the ODPC DPIA Guidance Note.
DPO	Designate or appoint a DPO where the AI deployment involves large-scale systematic monitoring or large-scale processing of sensitive personal data.
Automated Decision-Making	Implement a human review mechanism for all AI systems that make or significantly contribute to decisions affecting data subjects. Ensure data subjects can request human review and receive an explanation of the decision logic.
Data Minimisation	Review AI training datasets and inference inputs to ensure only necessary data is collected and processed. Remove or anonymise data that is not required.
Data Processor Agreements	Ensure written data processing agreements are in place with all third-party AI vendors and cloud AI service providers, containing the mandatory provisions under the General Regulations.
Cross-Border Transfers	Assess whether AI processing involves cross-border data transfers. Ensure adequate legal basis for any transfer under Section 48 of the Act.

Breach Notification	Implement incident response procedures for AI-related data breaches. Ensure the Office is notified within 72 hours and affected data subjects are notified where required.
Ongoing Monitoring	Establish governance procedures for ongoing monitoring of AI systems for bias, drift, and emerging privacy risks. Conduct periodic algorithmic audits.
Staff Training	Train staff responsible for AI governance, including DPOs, data protection leads, and AI system administrators, on their obligations under this Guidance Note and the Act.

Draft Guidance Notes on AI

ANNEX 2: DPIA TRIGGER TABLE FOR AI USE CASES

The following table provides guidance on whether a DPIA is required for common AI use cases. A "Yes" indicates that a DPIA is mandatory before deployment. A "Review" indicates that a DPIA should be conducted unless a documented assessment determines the risk to be low. The use cases and obligations listed are illustrative and not exhaustive.

AI Use Case	DPIA Required?	Key Risk Factors
AI credit scoring for loan decisions	Yes	Large-scale profiling; automated decisions affecting financial rights; potential discriminatory outcomes
AI-driven health diagnostics using patient records	Yes	Sensitive health data at scale; clinical risk if AI errs; re-identification risk
Facial recognition for customer authentication	Yes	Biometric sensitive data; risk of mass surveillance; high re-identification risk
Generative AI chatbot with no personal data input	Review	Risk depends on whether personal data is incidentally processed in user inputs
AI email spam filter (enterprise internal)	Review	Lower risk if processing limited to internal communications without sensitive data
AI content recommendation engine (consumer platform)	Yes	Large-scale profiling of consumer preferences; potential sensitive data inference
AI student assessment and adaptive learning	Yes	Children's data; automated decisions affecting educational outcomes; long retention
AI in recruitment screening and candidate ranking	Yes	Profiling; potential discrimination; automated decisions affecting employment rights
AI fraud detection in payment systems	Yes	Large-scale transaction monitoring; false positives affecting access to financial services
AI predictive maintenance (industrial, no personal data)	No	No personal data processed
AI translation tool with anonymised inputs	Review	Low risk if inputs are genuinely anonymised; risk of re-identification from context
AI-based employee productivity monitoring	Yes	Systematic employee monitoring; potential sensitive inference; power imbalance

Entities uncertain whether a DPIA is required for a particular AI application are encouraged to contact the Office for guidance. The Office may, in appropriate cases, issue advisory opinions on DPIA requirements.

Britam Towers, 12th Floor, Hospital Road, Upperhill, Nairobi, Kenya
P.O. Box 30920-00100, G.P.O. Nairobi

Draft Guidance Notes on AI