



CBA STANDARD

AI Governance Professional

Technology · 6 weighted domains · 31 learning objectives

Version 6.1

This edition published September 2026

About this standard

This standard defines what competent practice consists of in the discipline set out under Scope below. It divides the discipline into weighted domains and states, for each domain, the learning objectives a competent practitioner meets. It is one of the CBA Standards, which the Corporate Business Alliance maintains across finance, accounting, management and technology.


The standard is written for employers defining what a role requires, for training providers designing courses to a published specification, and for practitioners planning their own development.

Each domain carries a weight: its share of the standard, and of the examination for the CBA-AIG credential, which is assessed against it. The version number rises whenever a change would alter what competence in the discipline means, and every credential records the version its holder met.

Scope

A professional certification in organisational AI governance covering risk classification, the EU AI Act landscape, NIST AI RMF and ISO/IEC 42001 concepts, documentation and transparency, bias and robustness testing, and incident response structures.

Weighting by domain



Domains and learning objectives

Domain 1

15% of the standard

Foundations of AI Governance

- 1.1 Distinguish AI governance from adjacent disciplines such as data protection, information security and model development quality assurance
- 1.2 Identify the characteristics of AI systems that create distinct governance needs, including opacity, drift, scale of automated decisions and probabilistic outputs
- 1.3 Classify organisational AI activity into common categories such as procured AI features,

in-house models, general-purpose model use and shadow AI

- 1.4 Select appropriate first steps for establishing governance in an organisation with no existing AI oversight, such as inventory building and use-case intake
- 1.5 Interpret the roles of provider, deployer, and affected person as they recur across regulatory and framework language

Domain 2

25% of the standard

Risk Classification and the Regulatory Landscape

- 2.1 Classify described AI use cases into the EU AI Act's tiers of prohibited practice, high risk, transparency risk and minimal risk
- 2.2 Distinguish the obligations that fall on providers from those that fall on deployers of high-risk AI systems
- 2.3 Interpret the treatment of general-purpose AI models, including when additional obligations attach to models presenting systemic risk
- 2.4 Evaluate the extraterritorial reach of AI regulation and its practical effect on organisations in the Middle East, Asia and other regions outside the EU
- 2.5 Select the correct regulatory response to borderline cases, such as emotion recognition at work, biometric identification and credit scoring
- 2.6 Identify the broad direction of AI rules in other jurisdictions, including sectoral approaches, national strategies and voluntary codes

Domain 3

20% of the standard

Governance Frameworks and Management Systems

- 3.1 Interpret the purpose and structure of the NIST AI Risk Management Framework, including its four functions and its trustworthiness characteristics
- 3.2 Distinguish a certifiable management system standard such as ISO/IEC 42001 from a voluntary risk framework such as the NIST AI RMF
- 3.3 Select framework elements appropriate to a given organisational need, such as risk identification, policy setting or supplier assurance
- 3.4 Evaluate an organisation's AI policy against the elements a recognised framework would expect it to contain
- 3.5 Interpret how the plan-do-check-act cycle applies to an AI management system, including internal audit and management review

Domain 4

15% of the standard

Documentation, Transparency and Accountability

- 4.1 Select the appropriate documentation artefact for a given purpose, such as a model card, data sheet, AI inventory entry or impact assessment
- 4.2 Evaluate the completeness of a model card or system documentation against the questions a deployer or regulator would ask
- 4.3 Interpret transparency obligations owed to end users and affected persons, including disclosure of AI interaction and synthetic content labelling

- 4.4 Distinguish explainability of individual decisions from transparency about a system's existence, purpose and limits
- 4.5 Calculate simple documentation coverage measures from an AI inventory, such as the proportion of high-risk systems with a completed impact assessment

Domain 5

15% of the standard

Bias, Robustness and Testing Concepts

- 5.1 Distinguish sources of bias across the lifecycle, including historical data bias, sampling bias, label bias and deployment drift
- 5.2 Interpret basic fairness measures conceptually, such as comparing selection or error rates across groups, and recognise that fairness definitions can conflict
- 5.3 Calculate simple disparity measures from a small results table, such as a selection-rate ratio between two groups
- 5.4 Evaluate the design of a testing plan for an AI system, including test data independence, edge cases and human review checkpoints
- 5.5 Select appropriate ongoing monitoring practices for a deployed system, including drift indicators, performance thresholds and re-validation triggers

Domain 6

10% of the standard

Incident Handling and Organisational Structures

- 6.1 Classify AI-related events by severity and type, distinguishing incidents, near misses and complaints
- 6.2 Select the correct sequence of actions when a serious AI incident occurs, including containment, escalation, communication and regulator notification concepts
- 6.3 Evaluate the design of an AI governance operating model, including board oversight, an AI governance committee and operational ownership
- 6.4 Distinguish the responsibilities of first-line operators, second-line risk and compliance functions and third-line internal audit for AI systems
- 6.5 Interpret post-incident review findings and select corrective actions that address root causes rather than symptoms

Competence at award

What a holder of the CBA-AIG credential has demonstrated.

- Classify AI systems and use cases by risk level using criteria drawn from the EU AI Act and comparable regulatory regimes, and assign provider and deployer responsibilities correctly
- Apply the core structures of the NIST AI Risk Management Framework and ISO/IEC 42001 to design an organisational AI governance programme at a conceptual level
- Produce and evaluate AI governance documentation, including AI inventories, use-case risk registers, model cards, data sheets and transparency notices
- Interpret the purpose, methods and limits of bias, fairness and robustness testing, and select appropriate metrics and review checkpoints for a given deployment

- Design AI incident handling, escalation and post-incident review routines, and allocate governance roles across boards, committees and operational teams
- Evaluate third-party and procured AI systems using contractual, documentation and assurance evidence available to a deploying organisation

Recognition at Professional level

The CBA-AIG credential is awarded at Professional level of the CBA Competency Framework, which evidences that the holder is a practising professional in the discipline.

Awarded on the examination, 6 assessed workbooks of applied work and 3 years' relevant professional experience (2 with a relevant degree), confirmed by a named referee, with a signed undertaking to the Code of Professional Conduct.

Who the standard is written for

Compliance and risk professionals

Officers in compliance, data protection, internal audit or enterprise risk who are being handed responsibility for AI oversight and need a structured, defensible approach to classifying and controlling AI use across the business.

IT and data leaders

IT managers, data managers and solution architects who select, integrate or operate AI-enabled systems and must translate governance requirements into inventories, documentation, testing routines and monitoring practices.

Consultants and project professionals

Advisers, analysts and project or programme managers who guide clients or internal stakeholders through AI adoption and want recognised evidence that they can assess regulatory exposure and stand up governance structures.

Appendix: glossary of technology terms

Terms used across the technology standards, as CBA defines them.

A/B test

A controlled experiment comparing two versions of something by assigning users at random to each, so that a difference in outcome can be attributed to the change.

Acceptable use policy

An organisation's rules on which AI tools staff may use, for what purposes, and with what data — including what must never be entered into an external service.

AI impact assessment

A structured assessment, made before deployment, of how an AI system could affect individuals, groups and the organisation, and of the controls needed in response.

AI risk classification

Assigning each AI system a risk level according to the consequence of it being wrong — who could be harmed, how seriously and how easily an error would be noticed — so that controls are proportionate.

AI system inventory

A register of every AI system an organisation uses or builds, recording what each does, the decisions it informs, its owner, its data and how its outputs are checked.

Algorithm

A defined sequence of steps for solving a problem or performing a calculation.

Algorithmic bias

Systematic error in a model's outputs that disadvantages particular groups, often arising from unrepresentative training data or from features that act as proxies for protected characteristics.

Anonymisation

Processing personal data so that individuals can no longer be identified by any means reasonably likely to be used. Truly anonymised data falls outside data protection law; pseudonymised data does not.

Artificial intelligence

Computer systems that perform tasks associated with human intelligence, such as recognising patterns, making predictions, understanding language or generating content.

Classification

A machine learning task that assigns inputs to categories, such as flagging a transaction as fraudulent or not.

Confusion matrix

A table comparing a classification model's predictions with actual outcomes, showing true and false positives and negatives.

Context window

The maximum amount of text, measured in tokens, that a language model can take into account at once. Material beyond it is not considered, which can cause instructions or facts to be lost in long tasks.

Correlation and causation

Correlation is a statistical association between two variables; causation means one brings about the other. Correlation alone does not establish causation.

Dashboard

A visual display of the key measures needed to monitor a process or objective, arranged so the most important information is read first.

Data cleansing

Detecting and correcting errors, duplicates, inconsistent formats and missing values in a dataset before analysis.

Data governance

The roles, policies and standards that determine who owns, may use and is accountable for an organisation's data, and how its quality and security are maintained.

Data lineage

The documented path of data from its source through each transformation to where it is used, so that a figure in a report can be traced back.

Data minimisation

The data protection principle that personal data must be adequate, relevant and limited to what is necessary for the purpose.

Data quality

The degree to which data is accurate, complete, consistent, timely and fit for its intended use.

Data visualisation

Representing data graphically — charts, maps, tables — so that patterns and comparisons can be understood quickly and accurately.

Data warehouse

A central store of structured data integrated from several sources and organised for reporting and analysis.

Descriptive analytics

Analysis that summarises what has happened, through measures, aggregations and reports.

Diagnostic analytics

Analysis that investigates why something happened, by drilling into data, comparing segments and testing explanations.

ETL

Extract, transform, load: the process of taking data from source systems, converting it into a consistent structure and loading it into a target store such as a data warehouse.

EU AI Act

Regulation (EU) 2024/1689, the European Union's law on artificial intelligence. It classifies AI systems by risk, prohibits certain practices and places obligations on providers and deployers of high-risk systems.

Explainability

The extent to which the reasons for an AI system's output can be understood by the people who rely on it or are affected by it.

Feature

An input variable used by a machine learning model, such as transaction amount or customer tenure.

Fine-tuning

Further training of a pre-trained model on a smaller, task-specific dataset to adapt its behaviour.

Generative AI

AI systems that produce new content — text, images, code, audio — in response to a prompt, based on patterns learned from their training data.

Hallucination

Output from a generative AI system that is fluent and plausible but factually wrong or unsupported, such as an invented figure or citation.

Human oversight

Arrangements that allow people to monitor, question and override an AI system. It is real only where the reviewer has the time, information and authority to intervene.

ISO/IEC 42001

The international standard, published in 2023, specifying requirements for an artificial intelligence management system within an organisation.

Large language model (LLM)

A generative AI model trained on large volumes of text to predict and produce language, used for drafting, summarising, answering questions and similar tasks.

Machine learning

A branch of AI in which systems learn patterns from data to make predictions or decisions, rather than following explicitly programmed rules.

Median

The middle value of a set of numbers arranged in order. Unlike the mean, it is not pulled by extreme values.

Model card

A short document describing an AI model's intended use, training data, performance, limitations and known risks.

Model drift

The decline in a model's accuracy over time as the data it meets in use departs from the data it was trained on.

NIST AI Risk Management Framework

A voluntary framework published by the US National Institute of Standards and Technology in 2023 to help organisations identify, assess and manage the risks of AI systems.

Outlier

A data point far from the others. It may be an error to correct or a genuine and important observation, and should be investigated rather than removed by default.

Overfitting

When a model learns the noise in its training data as well as the signal, performing well on that data but poorly on new data.

Personal data

Under UK and EU data protection law, any information relating to an identified or identifiable living individual.

Pivot table

A spreadsheet tool that summarises a table of records by grouping and aggregating them along chosen dimensions.

Precision and recall

Precision is the share of cases a model flags that are truly positive; recall is the share of truly positive cases the model flags. Improving one usually costs the other.

Predictive analytics

Analysis that estimates what is likely to happen, using statistical or machine learning models built on historical data.

Prescriptive analytics

Analysis that recommends a course of action, typically by combining predictions with optimisation or decision rules.

Prompt

The instruction and context given to a generative AI system. A well-specified prompt states the task, the relevant context, the constraints and the expected form of the output.

Pseudonymisation

Replacing identifying details with a code, so that data cannot be attributed to an individual without separately held information. Pseudonymised data remains personal data.

Red-teaming

Deliberately attempting to make an AI system fail — to produce harmful, wrong or prohibited output — in order to find weaknesses before users do.

Regression

A statistical or machine learning method for predicting a continuous value, such as sales, from one or more input variables.

Retrieval-augmented generation (RAG)

A technique in which relevant documents are retrieved and supplied to a language model with the prompt, so that its answer is grounded in that material rather than in its training alone.

Robustness

An AI system's ability to keep performing correctly when inputs are noisy, unusual or deliberately manipulated.

Shadow AI

The use of AI tools by staff without the organisation's knowledge or approval, often with data that should not leave it.

SQL

Structured Query Language, the standard language for querying and managing data in relational databases.

Standard deviation

A measure of how widely values are spread around their mean. Roughly two thirds of a normally distributed set lie within one standard deviation of the mean.

Statistical significance

An indication that an observed result would be unlikely if there were no real effect. It does not show that the effect is large or important.

Supervised learning

Machine learning from examples labelled with the correct answer, so that the model learns to predict the label for new cases.

Test set

Data held back from training and used only to measure how well a finished model performs on data it has not seen.

Token

The unit of text a language model processes — often a word or part of a word. Context limits and usage costs are measured in tokens.

Training data

The data from which a machine learning model learns. Its quality, coverage and representativeness largely determine the model's behaviour.

Unsupervised learning

Machine learning that finds structure in unlabelled data, such as grouping similar customers into segments.

Using this publication

Citation

Corporate Business Alliance (2026). AI Governance Professional: CBA Standard, version 6.1.
London: Corporate Business Alliance.

Terms of use

© Corporate Business Alliance 2026. You may download, print and use this publication, and share it unaltered and free of charge, for personal study and within your organisation. It may not be sold, or republished in whole or in part, without the written permission of the Corporate Business Alliance.

Status

The CBA Standards and CBA credentials are the professional standards and certifications of the Corporate Business Alliance. CBA credentials are not regulated qualifications, and CBA is not accredited to ISO/IEC 17024; CBA claims no equivalence to the qualifications of any other body.

Publisher

Corporate Business Alliance Ltd

Registered office: 167–169 Great Portland Street, 5th Floor, London, W1W 5PF, United Kingdom
support@cba-uk.com · cba-uk.com