



CBA STANDARD

Artificial Intelligence Professional

Technology · 6 weighted domains · 36 learning objectives

Version 6.2

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-AIP 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 broad professional certification in applied artificial intelligence covering machine learning and deep learning concepts, data foundations, the full model lifecycle, application patterns including generative AI, and responsible AI practice. Conceptual and applied; no coding required.

Weighting by domain



Domains and learning objectives

Domain 1

20% of the standard

AI and Machine Learning Foundations

- 1.1 Distinguish artificial intelligence, machine learning, deep learning and rule-based automation by how each produces its outputs.
- 1.2 Classify a described problem as supervised, unsupervised or reinforcement learning and identify the target and features involved.
- 1.3 Select an appropriate model family (linear models, tree-based models, clustering, neural

networks) for a given problem and justify the choice by data type, interpretability need and scale.

- 1.4 Interpret the roles of parameters, training, loss and optimisation in how a model learns, using simple numeric illustrations.
- 1.5 Explain how core deep learning architectures (feedforward networks, convolutional networks, sequence models and transformers) differ in the data they suit.
- 1.6 Identify the bias-variance trade-off, underfitting and overfitting from descriptions of model behaviour on training and unseen data.

Domain 2

15% of the standard

Data Foundations for AI

- 2.1 Assess the suitability of a dataset for a stated modelling task in terms of volume, representativeness, label quality and freshness.
- 2.2 Select appropriate treatments for missing values, outliers, duplicates and inconsistent categories in a described dataset.
- 2.3 Distinguish training, validation and test splits and identify splitting mistakes such as leakage and temporal contamination.
- 2.4 Apply basic feature concepts, including encoding categorical variables, scaling numeric variables and simple feature engineering, to spreadsheet-scale examples.
- 2.5 Evaluate the effect of class imbalance and sampling choices on what a model will learn.
- 2.6 Identify data governance obligations relevant to AI datasets, including consent, minimisation and retention at concept level.

Domain 3

20% of the standard

Model Development and Evaluation

- 3.1 Calculate accuracy, precision, recall, F1 and error rates from a small confusion matrix or result table.
- 3.2 Select the evaluation metric that matches a stated business cost of false positives versus false negatives.
- 3.3 Interpret regression evaluation measures (mean absolute error, root mean squared error, R-squared) and compare candidate models with them.
- 3.4 Distinguish the purposes of validation and test data and evaluate whether a described experiment supports its performance claim.
- 3.5 Evaluate baseline comparisons, hyperparameter tuning descriptions and simple experiment designs for soundness.
- 3.6 Interpret model explainability outputs, such as feature importance summaries, and identify their limits.

Domain 4

15% of the standard

Deployment, Operations and the AI Lifecycle

- 4.1 Evaluate whether a described business problem is a sound candidate for an AI solution and define measurable success criteria for it.

-
- 4.2 Select an appropriate deployment pattern (batch scoring, real-time API, embedded on-device, human-in-the-loop) for a stated use case.
 - 4.3 Distinguish build, buy and API-based sourcing options for AI capability and assess their cost, control and data implications.
 - 4.4 Interpret monitoring signals to identify data drift, concept drift and performance degradation in a deployed model.
 - 4.5 Select appropriate responses to model degradation, including retraining, rollback and threshold adjustment.
 - 4.6 Identify the documentation and versioning practices that make an AI system auditable and reproducible.
-

Domain 5

15% of the standard

AI Application Patterns and Generative AI

- 5.1 Distinguish predictive, generative and retrieval-based AI applications by their inputs, outputs and failure modes.
 - 5.2 Explain how large language models produce text, including tokens, context windows and probabilistic generation, at concept level.
 - 5.3 Select prompt design techniques (role setting, examples, structured output, decomposition) appropriate to a described task.
 - 5.4 Evaluate when retrieval-augmented generation, fine-tuning or an unmodified foundation model is the right approach for a knowledge task.
 - 5.5 Identify hallucination, prompt injection and over-reliance risks in described generative AI workflows and select mitigations.
 - 5.6 Select appropriate human-in-the-loop checkpoints for AI-assisted work in fields such as drafting, translation, coding support and customer service.
-

Domain 6

15% of the standard

Responsible AI, Governance and Risk

- 6.1 Identify sources of unfair bias across the AI lifecycle, from historical data and labels to deployment context.
 - 6.2 Interpret simple fairness evidence, such as approval or error rates broken down by group, and identify the concern it raises.
 - 6.3 Distinguish transparency, explainability and accountability obligations for AI systems and match each to appropriate practices.
 - 6.4 Evaluate privacy and security risks specific to AI, including training-data exposure, model inversion at concept level, and misuse of generated content.
 - 6.5 Classify described AI systems into risk tiers of the kind used by the EU AI Act and identify the obligations that follow at concept level.
 - 6.6 Select appropriate governance controls for an AI deployment, including impact assessment, human oversight, documentation and incident response.
-

Competence at award

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

- Explain how the main families of machine learning and deep learning models learn from data and select an appropriate approach for a given business problem.
- Specify, assess and prepare the data an AI project needs, including labelling, splitting, feature treatment and data quality controls.
- Interpret model evaluation outputs such as confusion matrices, precision, recall, error measures and validation results, and detect overfitting and leakage.
- Plan the lifecycle of an AI solution from problem framing and success criteria through deployment, monitoring, drift detection and retraining.
- Apply modern AI application patterns, including generative AI, prompt design, retrieval-augmented generation and human-in-the-loop workflows, using concepts that transfer across tools.
- Evaluate AI systems for fairness, transparency, privacy and security risks, and relate them to major governance frameworks such as the EU AI Act and relevant international standards.

Recognition at Practitioner level

The CBA-AIP credential is awarded at Practitioner level of the CBA Competency Framework, which evidences that the holder can carry out the core tasks of the discipline.

Awarded on the examination and 4 assessed workbooks of applied work.

Who the standard is written for

Business and data analysts moving into AI work

Analysts who already work with data in Excel, SQL or BI tools and are increasingly asked to scope, evaluate or interpret machine learning solutions, and who want a structured, credible grounding in how models actually work.

Project, product and delivery leads on AI initiatives

Managers and product owners responsible for AI projects who need to frame problems correctly, challenge vendor claims, read evaluation results, plan deployment and monitoring, and manage AI-specific risks without writing code themselves.

IT professionals and consultants broadening into AI

Software, infrastructure, security and advisory professionals who want a recognised certification demonstrating fluency across the AI landscape, from data foundations and model lifecycle to generative AI applications and responsible AI governance.

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). Artificial Intelligence Professional: CBA Standard, version 6.2. 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