



CBA STANDARD

Generative AI Practitioner

Technology · 6 weighted domains · 32 learning objectives

Version 6.0

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-GAI 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 practical generative AI: how large language and diffusion models behave, prompt design, retrieval-grounded and multimodal workflows, output evaluation, and safe, well-governed use of AI in everyday work.

Weighting by domain



Domains and learning objectives

Domain 1

18% of the standard

Generative Model Foundations

- 1.1 Distinguish the pretraining, instruction-tuning and inference stages of a large language model lifecycle and what each contributes to behaviour
- 1.2 Interpret the effect of sampling controls such as temperature, top-p and maximum output length on variability and reliability
- 1.3 Identify how tokenisation and context window limits affect cost, truncation and the handling

of long or non-English text

- 1.4 Distinguish how diffusion models turn noise into images from how autoregressive models generate text token by token
- 1.5 Select an appropriate model class or capability tier for a stated task, balancing quality, speed, cost and access constraints
- 1.6 Recognise hallucination, knowledge cut-off effects and arithmetic weakness as predictable consequences of how models are built

Domain 2

22% of the standard

Prompt Design and Iteration

- 2.1 Select prompt components (role, context, task, constraints, output format, examples) appropriate to a stated goal
- 2.2 Distinguish zero-shot, few-shot and reasoning-eliciting prompting and identify when each improves results
- 2.3 Evaluate a weak prompt against a target outcome and select the revision most likely to improve the output
- 2.4 Interpret an unsatisfactory response to diagnose whether the fault lies in the instructions, the supplied context or the task framing
- 2.5 Apply output-format specifications, including tables, structured lists and JSON-like templates, to obtain reusable results
- 2.6 Construct multi-step prompt chains that decompose a large task into ordered, checkable stages

Domain 3

15% of the standard

Retrieval and Grounded Workflows

- 3.1 Distinguish tasks that need retrieval or grounding from tasks the model can answer reliably from its training alone
- 3.2 Interpret the stages of a retrieval-augmented pipeline: splitting documents into chunks, embedding, similarity search and prompt assembly
- 3.3 Evaluate how chunk size, retrieval quality and stale sources affect the accuracy of grounded answers
- 3.4 Select an appropriate grounding method (file upload, long-context prompting, web-connected search, retrieval pipeline) for a stated information need
- 3.5 Identify citation and traceability practices that let a reader verify a grounded answer against its sources

Domain 4

10% of the standard

Multimodal Creation and Editing

- 4.1 Select image-generation prompt elements, including subject, style, composition, lighting and negative guidance, for a stated creative brief
- 4.2 Distinguish the practical capabilities of text-to-image, image-to-image, vision-input, speech-to-text and text-to-speech tools

- 4.3 Evaluate generated media for common artefacts and defects such as distorted text, anatomical errors and inconsistent lighting
- 4.4 Interpret disclosure, provenance and licensing considerations that apply before publishing AI-generated media

Domain 5

15% of the standard

Output Evaluation and Quality Control

- 5.1 Evaluate an AI output against explicit criteria such as factual accuracy, completeness, relevance, tone and formatting
- 5.2 Distinguish types of error, including fabricated facts, false citations, subtle omissions and confident wrongness, and select a suitable verification method for each
- 5.3 Construct a simple rubric and small test set for a generative AI task that a team performs repeatedly
- 5.4 Calculate basic quality measures, such as pass rates and error rates, from a structured spot-check recorded in a spreadsheet
- 5.5 Interpret the limits of using one AI model to review another model's output

Domain 6

20% of the standard

Safe Use, Data Care and Work Integration

- 6.1 Identify categories of information, including personal data, client confidences and commercial secrets, that must not be entered into tools without appropriate safeguards
- 6.2 Interpret core data protection principles, such as purpose limitation, minimisation and individual rights, as they apply to workplace generative AI use
- 6.3 Distinguish consumer and enterprise tool arrangements, including training-on-your-data settings, retention and administrative controls
- 6.4 Evaluate a proposed workplace use case for risk level and select proportionate human oversight and disclosure
- 6.5 Recognise prompt injection, jailbreaking and data leakage risks in connected or automated workflows
- 6.6 Select integration patterns that embed generative AI into everyday work, such as drafting, summarising, meeting notes and spreadsheet analysis

Competence at award

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

- Explain, in practical terms, how large language models and diffusion models generate output and why this produces characteristic strengths and failure modes
- Design, test and refine prompts using structure, examples, constraints and iteration to obtain reliable, well-formatted results
- Choose and apply grounding strategies, including document upload, web-connected search and retrieval-augmented pipelines, to reduce hallucination and improve traceability
- Work across text, image, audio and vision inputs, selecting appropriate multimodal tools and

recognising quality defects in generated media

- Evaluate AI outputs systematically using verification workflows, rubrics and simple metrics rather than trusting fluency
- Apply data hygiene, confidentiality and disclosure practices that keep generative AI use lawful and professionally defensible
- Integrate generative AI into everyday tasks such as drafting, research, meeting summaries and spreadsheet-based analysis using broadly accessible tools

Recognition at Practitioner level

The CBA-GAI 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

Knowledge workers and administrators

Professionals in operations, HR, marketing, finance support or administration who use AI assistants for drafting, summarising and analysis and want to do so faster, more accurately and within data protection rules.

Early-career and transitioning professionals

Graduates, junior analysts and career changers who need a recognised credential showing they can work productively with generative AI tools from their first day in a role.

Team leads and process owners

Supervisors and small-business owners who are introducing AI into team workflows and need enough technical grounding to choose tools, set usage rules and review AI-assisted work with confidence.

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). Generative AI Practitioner: CBA Standard, version 6.0. 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