

BCS Foundation Certificate In Artificial Intelligence

BCS AIF

Version Demo

Total Demo Questions: 18

Total Premium Questions: 189

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to Artificial Intelligence	18
Topic 2, The History of Artificial Intelligence	4
Topic 3, Benefits of Artificial Intelligence	5
Topic 4, Challenges and Risks of Artificial Intelligence	37
Topic 5, Artificial Intelligence in the Organisation	14
Topic 6, Machine Learning	60
Topic 7, Robotics	5
Topic 8, Artificial Intelligence in Practice	46
Total	189

QUESTION NO: 1

A financial company is developing a generative AI application for loan approval decisions. The company needs the application output to be responsible and fair.

Which solution meets these requirements?

- A. Review the training data to check for biases. Include data from all demographics in the training data.
- B. Use a deep learning model with many hidden layers.
- C. Keep the model 's decision-making process a secret to protect proprietary algorithms.
- D. Continuously monitor the model's performance on a static test dataset.

ANSWER: A

Explanation:

Review the training data to check for biases. Include data from all demographics in the training data. This is the appropriate solution because the data used to develop an AI system strongly influences the fairness of its outputs. In a loan-approval context, historical data can contain patterns arising from past discriminatory decisions, unequal access to credit, or under-representation of particular demographic groups. Reviewing the data helps identify such issues before they are embedded in the application's decision logic. Ensuring that relevant groups are appropriately represented also supports evaluation of whether outcomes are equitable across those groups, rather than performing well only for the populations most prevalent in the original dataset.

Responsible AI requires active governance over data quality, representativeness, bias, and potential harmful impacts. These measures are particularly important for lending, where automated decisions can materially affect individuals' financial opportunities. Fairness work should continue throughout the system lifecycle, including documented dataset governance, suitable fairness measures, and human oversight for consequential decisions. The UK Information Commissioner's Office explains the importance of addressing bias and discrimination in AI systems in its [Guidance on AI and data protection](#). The [NIST AI Risk Management Framework](#) similarly identifies harmful bias as a risk that organisations should govern and manage.

QUESTION NO: 2

A financial company stores patterns of fraudulent behavior in a database. The company uses this data to conduct investigations.

The company wants to use a graph-based ML solution to develop an AI tool that helps with these investigations.

Which AWS service will meet these requirements?

- A. Amazon OpenSearch Service
- B. Amazon Aurora
- C. Amazon Neptune
- D. Amazon MemoryDB

ANSWER: C

Explanation:

Amazon Neptune is correct because it is AWS's purpose-built graph database service, designed to model and query highly connected data such as relationships among customers, accounts, transactions, devices, and known fraud patterns. Fraud investigation commonly depends on identifying indirect connections and suspicious relationship patterns that are difficult to represent or investigate efficiently with conventional relational tables. Neptune supports graph models and graph query languages, enabling investigators or applications to traverse these connections and discover patterns across many linked entities.

It also provides Neptune ML, which integrates graph data with machine learning workflows to support graph-based predictions. For example, a model can use the structure of relationships and entity attributes to help identify potentially fraudulent accounts, classify suspicious transactions, or detect anomalous links. This directly satisfies the requirement for a graph-based ML solution that assists investigations, while retaining the graph database as the system for storing and querying the relationship data. AWS describes Neptune and its graph use cases at [Amazon Neptune](#) and documents its machine-learning capability at [Neptune ML](#).

QUESTION NO: 3

Ensemble learning methods do what with the hypothesis space?

- A. Select a combination of hypothesis to combine their predictions
- B. Use stochastic gradient descent to optimise a network.
- C. Extract ergodic solutions.
- D. Test multiple hypotheses simultaneously.

ANSWER: A

Explanation:

“Select a combination of hypothesis to combine their predictions” is correct because ensemble learning builds a stronger predictive model by using multiple candidate hypotheses (base learners) and aggregating what they predict. Rather than relying on one learned model to represent the entire problem, an ensemble selects or constructs several models from the available hypothesis space. Their outputs are then combined, for example through majority voting in classification, averaging in regression, or a learned weighted combination. This approach can improve predictive accuracy and generalisation because the constituent hypotheses may make different errors or capture different patterns in the training data. Common ensemble techniques include bagging, random forests, boosting, and stacking; each uses multiple base models and a defined method to combine their predictions into one final result. The important relationship to hypothesis space is therefore the use of a collection or combination of hypotheses, not merely optimisation of one model. Scikit-learn’s ensemble-learning documentation describes ensemble methods as combining the predictions of several base estimators to improve robustness and generalisation: [scikit-learn Ensemble methods](#). For a broader definition of ensemble methods and their combination of classifiers, see the [Ensemble Methods reference entry](#).

QUESTION NO: 4

Sustainability focuses on which three core areas?

- A. Scientific, Environmental and Economic.
- B. Social, Economic and Environmental.
- C. Social, Economic and Entrepreneurial.
- D. Social, Entrepreneurial and Environmental.

ANSWER: B

Explanation:

Social, Economic and Environmental. is correct because these are the three widely recognised, interdependent dimensions of sustainable development. Environmental sustainability concerns the protection and responsible management of natural systems, resources, biodiversity, and the climate. Economic sustainability concerns maintaining long-term prosperity and using resources in a manner that remains viable over time, rather than creating short-term value at unacceptable future cost. Social sustainability concerns human wellbeing, fairness, inclusion, safety, communities, and the distribution of benefits and impacts among people.

Considering all three dimensions together is particularly important when assessing artificial intelligence. An AI system may offer economic value or operational efficiency, but a sustainable assessment also considers its environmental effects, such as energy and resource use, and its social effects, such as accessibility, fairness, employment impacts, and human wellbeing. This balanced view supports decisions that meet present needs without reducing the ability of future generations to meet theirs. The United Nations describes sustainable development as integrating economic, social, and environmental dimensions, while BCS includes sustainability among the broader professional considerations relevant to AI practice. See the [United Nations Sustainable Development Goals](#) and the [BCS Artificial Intelligence certifications](#).

QUESTION NO: 5 - (HOTSPOT)

HOTSPOT

Select the correct AI term from the following list for each statement. Each AI term should be selected one time. (Select THREE

E.


)

- AI
- Deep learning
- ML

Simulates human problem-solving capabilities	Select... Select...AIDeep learningML
Applies data-driven learning techniques to make predictions	Select... Select...AIDeep learningML
Focuses on processing data through intricate neural networks	Select... Select...AIDeep learningML

ANSWER:

Answer:



Answer:

Explanation:

Artificial Intelligence (AI) is the broad field focused on simulating human problem-solving and cognitive abilities, including reasoning, perception, and decision-making. (Reference: AWS Certified AI Practitioner Official Study Guide) Machine Learning (ML) is a subset of AI that uses data-driven algorithms to identify patterns and make predictions without explicit programming for each specific task. (Reference: AWS Machine Learning Overview) Deep learning is a subset of ML that uses neural networks with many layers (deep neural networks) to process complex data and extract high-level features. (Reference: AWS Deep Learning on AWS)

Explanation:

Artificial intelligence is the correct term for simulating human problem-solving capabilities. AI is the broad discipline concerned with building systems that perform tasks associated with human intelligence, such as reasoning, perception, language processing, planning, and decision-making. In this sense, AI describes the overall goal of enabling computer systems to exhibit intelligent behaviour. The BCS AI Foundation syllabus treats AI as the umbrella field under which several approaches and techniques sit.

Machine learning is the correct term for applying data-driven learning techniques to make predictions. Rather than relying entirely on fixed, hand-written rules for every situation, a machine-learning model learns useful relationships and patterns from examples in data. Those learned patterns can then be used to classify new cases, estimate values, or generate predictions. This data-led learning process is central to the definition of machine learning, as described in [IBM's machine-learning overview](#).

Deep learning is the correct term for processing data through intricate neural networks. Deep learning is a specialised area within machine learning that uses neural networks containing multiple layers. These layers progressively transform input data into increasingly useful representations, allowing models to handle complex material such as images, speech, and natural language. The use of multi-layer neural networks is the defining feature that makes “deep learning” the appropriate term here. This relationship is also explained in [IBM's deep-learning overview](#): deep learning is a subset of machine learning, and machine learning is itself within the broader AI field.

QUESTION NO: 6

A documentary filmmaker wants to reach more viewers. The filmmaker wants to automatically add subtitles and voice-overs in multiple languages to their films.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Use Amazon Transcribe and Amazon Translate to generate subtitles in other languages
- B. Use Amazon Textract and Amazon Translate to generate subtitles in other languages
- C. Use Amazon Polly to generate voice-overs in other languages

D. Use Amazon Translate to generate voice-overs in other languages

E. Use Amazon Textract to generate voice-overs in other languages

ANSWER: A C

Explanation:

“Use Amazon Transcribe and Amazon Translate to generate subtitles in other languages” meets the subtitle requirement through a suitable speech-to-text and machine-translation workflow. Amazon Transcribe can create a timestamped transcript from the film’s spoken audio. That transcript supplies the text that can be formatted as captions or subtitles, while Amazon Translate can produce translations of the dialogue into the required target languages. In a production workflow, the translated, time-aligned subtitle content can then be associated with the video for viewers to select.

“Use Amazon Polly to generate voice-overs in other languages” meets the audio requirement once the dialogue has been translated. Amazon Polly is a text-to-speech service that converts text into synthetic speech and provides voices for multiple languages and locales. The translated scripts can therefore be submitted to Polly to create speech audio for each target language, which can be synchronised and mixed with the film as a voice-over track. Together, these services support the required automated pipeline: transcribe the original speech, translate its text for multilingual subtitles and scripts, then synthesize translated speech for multilingual voice-overs. AWS documents Transcribe’s media transcription capabilities at [Amazon Transcribe Developer Guide](#) and Polly’s multilingual text-to-speech capability at [Amazon Polly Developer Guide](#).

QUESTION NO: 7

A company is developing an ML application. The application must automatically group similar customers and products based on their characteristics.

Which ML strategy should the company use to meet these requirements?

A. Unsupervised learning

B. Supervised learning

C. Reinforcement learning

D. Semi-supervised learning

ANSWER: A

Explanation:

Unsupervised learning is correct because the requirement is to discover natural groupings among customers and products according to their shared characteristics, without being given predefined group labels. This task is commonly called clustering. A clustering algorithm examines feature data—such as purchasing behaviour, demographics, product attributes, or usage patterns—and assigns records to groups whose members are more similar to one another than to members of other groups. These groups can then support practical activities such as customer segmentation, product recommendation, targeted marketing, catalogue organisation, and the identification of related products.

The key feature of this scenario is that the application must automatically find the structure within the data. There is no stated target value or known category for each customer or product from which a model could learn. Unsupervised learning is designed for precisely this exploratory purpose: finding patterns, relationships, and clusters in unlabelled data. The BCS AI Foundation syllabus identifies clustering as an unsupervised machine-learning approach, and the general definition and uses of clustering are also described by [scikit-learn’s clustering documentation](#). For a broader overview of how unsupervised learning derives insights from unlabelled datasets, see [Google’s Machine Learning clustering overview](#).

QUESTION NO: 8

Who was the pioneer of computer programming?

A. Dame Wendy Hall.

- B. Karen Spark Jones.
- C. Ada Lovelace.
- D. Sophie Wilson

ANSWER: C

Explanation:

Ada Lovelace is widely recognised as the pioneer of computer programming because her notes on Charles Babbage's proposed Analytical Engine went beyond describing the machine's mechanisms. In her 1843 translation of Luigi Menabrea's article about the Analytical Engine, Lovelace added extensive original notes that explained how a general-purpose machine could manipulate symbols according to specified operations. Most notably, her Note G set out a step-by-step method for calculating Bernoulli numbers. This is commonly regarded as the first published algorithm specifically intended to be processed by a machine, and is why she is often described as the first computer programmer.

Her contribution was also conceptually important to artificial intelligence and computing history: she recognised that such a machine could operate on more than numerical quantities when those quantities could be represented symbolically. This anticipates the general-purpose, programmable nature of modern computers, rather than viewing the Analytical Engine merely as an advanced calculator. The Computer History Museum discusses Lovelace's algorithm and its historical significance at [Computer History Museum: Ada Lovelace](#). The Science Museum Group also provides context on the Analytical Engine and Lovelace's notes at [Ada Lovelace and the Analytical Engine](#).

QUESTION NO: 9

A company is using custom models in Amazon Bedrock for a generative AI application. The company wants to use a company-managed encryption key to encrypt the model artifacts that the model customization jobs create. Which AWS service meets these requirements?

- A. AWS Key Management Service (AWS KMS)
- B. Amazon Inspector
- C. Amazon Macie
- D. AWS Secrets Manager

ANSWER: A

Explanation:

AWS Key Management Service (AWS KMS) is correct because it provides the customer managed keys that Amazon Bedrock can use to encrypt artifacts produced by model-customization jobs. A customer managed KMS key is created and controlled in the company's AWS account, allowing the company to define key policies, control which IAM principals and AWS services may use the key, audit key usage through AWS CloudTrail, and manage key rotation in accordance with its security and governance requirements. When configuring a Bedrock customization job, the company can supply the ARN of an appropriate KMS key for output-data encryption. Bedrock then uses that key to protect the resulting customized-model artifacts at rest. This separates the management of encryption keys from the generative AI service while retaining AWS's managed encryption integration. The key policy must grant Amazon Bedrock the required permissions to use the key on the company's behalf, and the relevant job role also needs suitable KMS permissions. AWS documents KMS as its central service for creating and managing encryption keys: [AWS KMS concepts](#). The Amazon Bedrock documentation also describes using a KMS key to encrypt model customization job output: [Create a model customization job](#).

QUESTION NO: 10

A company acquires International Organization for Standardization (ISO) accreditation to manage AI risks and to use AI responsibly. What does this accreditation certify?

- A. All members of the company are ISO certified.

- B. All AI systems that the company uses are ISO certified.
- C. All AI application team members are ISO certified.
- D. The company's AI management system is ISO certified.

ANSWER: D

Explanation:

The company's AI management system is ISO certified. In this context, the relevant standard is ISO/IEC 42001, which specifies requirements for establishing, implementing, maintaining and continually improving an Artificial Intelligence Management System (AIMS). Certification provides independent assurance that the organisation has a governed framework for managing AI across its lifecycle, including policies, roles, risk assessment, impact assessment, controls, monitoring and continual improvement. It supports responsible AI by requiring the organisation to address AI-related risks and opportunities in a systematic, auditable way that is appropriate to its objectives and context. The certification therefore applies to the organisation's management framework and its processes for governing AI, rather than representing a personal qualification for every employee. ISO describes ISO/IEC 42001 as the world's first AI management-system standard and explains that it provides a structured way to manage risks and opportunities associated with AI. Further detail is available from [ISO/IEC 42001:2023](#) and ISO's overview of [what ISO/IEC 42001 means for AI governance](#).

QUESTION NO: 11 - (SIMULATION)

A company wants to develop a solution that uses generative AI to create content for product advertisements, including sample images and slogans.

Select the correct model type from the following list for each action. Each model type should be selected one time. (Select THREE)


E.
)

- Diffusion model
- Object detection model
- Transformer-based model

Create high-quality images that are influenced by the generated slogans and product

Create contextually relevant slogans based on the advertisement product

Ensure that company brand elements are properly placed in the images



ANSWER: See the explanation for the answer

Explanation:

Select the following model types:

Create high-quality images influenced by generated slogans and product: Diffusion model

Create contextually relevant slogans based on the advertisement product: Transformer-based model
Ensure company brand elements are properly placed in the images: Object detection model

Diffusion models generate images, transformer-based models generate and understand contextual text such as slogans, and object-detection models identify visual objects or brand elements to verify their placement.

QUESTION NO: 12

A company wants to use Amazon Q Business for its data. The company needs to ensure the security and privacy of the data.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Enable AWS Key Management Service (AWS KMS) keys for the Amazon Q Business enterprise index.
- B. Set up cross-account access to the Amazon Q index.
- C. Configure Amazon Inspector for authentication.
- D. Allow public access to the Amazon Q index.
- E. Configure AWS Identity and Access Management (IAM) for authentication.

ANSWER: A E

Explanation:

Enable AWS Key Management Service (AWS KMS) keys for the Amazon Q Business enterprise index because encryption protects the confidentiality of enterprise data while it is stored. Amazon Q Business supports AWS KMS keys for encrypting data at rest, giving an organization control over its customer managed key, including key policies, rotation choices, and auditing of key use through AWS services. This is an important privacy control when the index contains internal documents or data retrieved from connected sources.

Configure AWS Identity and Access Management (IAM) for authentication because identity and access controls ensure that only approved users, groups, and services can administer Amazon Q Business resources or access relevant AWS capabilities. IAM policies apply least-privilege permissions, while authentication establishes the identity making a request. Used together with Amazon Q Business user and application access configuration, IAM provides a strong administrative security boundary and helps prevent unauthorized access to data and configuration.

These controls address complementary aspects of security: AWS KMS safeguards stored data through encryption, and IAM governs authenticated, authorized access to AWS resources. AWS describes Amazon Q Business encryption and access-management capabilities in its [Amazon Q Business security documentation](#) and explains permissions management in the [AWS IAM User Guide](#).