USAII CAIC Practice Questions

Shared by Stark on 12-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which type of AI mimics the human brain by building networks of artificial neurons?

Options:

A- Fuzzy Logic
B- Genetic Algorithms
C- Neural Networks
D- Expert Systems
E- None of the above

P2P exams

Answer:

C

Explanation:

The correct answer is C. Neural Networks. Neural networks are AI and machine learning models inspired by the structure and function of the human brain. They are built using interconnected artificial neurons, also called nodes, which process input data, apply weights and activation functions, and pass information through layers to produce an output. This structure allows neural networks to learn complex patterns from data.

Fuzzy logic is incorrect because it deals with reasoning under uncertainty using degrees of truth rather than artificial neurons. Genetic algorithms are optimization methods inspired by natural selection and evolution, not the brain's neuron-based structure. Expert systems use rules and knowledge bases to make decisions, but they do not mimic the brain through networks of artificial neurons. Since the question specifically asks about AI that mimics the human brain by building artificial neuron networks, the correct answer is C. Neural Networks.

P2P exams

Question 2

Question Type: MultipleChoice

Which one of the following is a NOT good attribute of solution architecture?

Options:

A- Technology in alignment with business requirements
B- Scalability and flexibility
C- Tightly coupled architecture
D- Risk mitigation
E- Increased ROI

Answer:

C

Explanation:

The correct answer is C. Tightly coupled architecture because a strong solution architecture should promote flexibility, scalability, maintainability, integration readiness, and adaptability. A tightly coupled architecture means system components are highly dependent on one another. This creates problems when teams need to update, scale, replace, test, or modify one part of the system, because changes in one component can easily affect other components. In enterprise AI and software solution design, this increases operational risk, slows innovation, and makes future growth more difficult.

Technology alignment with business requirements is a good attribute because architecture must support business goals and operational needs. Scalability and flexibility are also good attributes because modern solutions must handle growth, changing workloads, and evolving requirements. Risk mitigation is a strong architectural objective because good design reduces security, performance, compliance, and operational risks. Increased ROI is also a desired outcome when architecture improves efficiency and business value. Therefore, the attribute that is NOT good is C. Tightly coupled architecture.

Question 3

Question Type: MultipleChoice

What is a prompt?

Options:

A- Prompt refers to the input the user provides to the generative models.
B- For ChatGPT and OpenAI models, prompts mainly refer to a piece of text in natural language.
C- Prompts are the only way users can control the output generated by those models.
D- a and b only
E- a, b and c only

Answer:

D

Explanation:

The correct answer is D. a and b only because a prompt is the input provided by a user to a generative AI model. In natural language systems such as ChatGPT and other language models, the prompt is usually written as text in natural language. It may be a question, instruction, command, description, context, example, or task requirement that guides the model toward producing a response.

Statement A is correct because prompts are the user-provided input that generative models use to produce outputs. Statement B is also correct because, for ChatGPT and similar models, prompts commonly appear as natural language text. Statement C is not fully correct because prompts are an important way to guide model output, but they are not the only possible control mechanism. Outputs can also be influenced by system instructions, model settings, retrieval context, fine-tuning, guardrails, and application design. Therefore, the best answer is D. a and b only.

Question 4

Question Type: MultipleChoice

Choose the CORRECT example of Reinforcement Learning.

Options:

A- robotics
B- game playing
C- navigation
D- Semi Reinforcement learning
E- Reinforcement learning

Answer:

D

Explanation:

The correct answer is D. All of the above because robotics, game playing, and navigation are all common examples of reinforcement learning. Reinforcement learning is a machine learning approach in which an agent learns by interacting with an environment and receiving rewards or penalties based on its actions. Over time, the agent learns a policy that helps it maximize long-term reward.

Robotics is a strong example because robots can learn movement, object handling, path planning, and control actions through trial and feedback. Game playing is another classic reinforcement learning example because an AI agent can learn winning strategies by trying actions, observing outcomes, and improving decisions over repeated episodes. Navigation is also a valid example because an agent can learn the best route or movement strategy by receiving feedback about distance, obstacles, time, or success in reaching a goal.

Since all three listed options are valid applications of reinforcement learning, the correct answer is D. All of the above.

Question 5

Question Type: MultipleChoice

Which option best is the CORRECT stage of the Data and AI Analytics Business Model Maturity Index?

Options:

A- Business Monitoring
B- Business Insights
C- Business Optimization
D- Cultural Transformation
E- Better performance with large datasets

Answer:

E

Explanation:

The correct answer is E. All of the above because the Data and AI Analytics Business Model Maturity Index describes how organizations progress in their ability to use data, analytics, and AI for business value creation. Business Monitoring is a valid stage because organizations first use data to observe performance, track metrics, and understand what is happening in the business. Business Insights is also a correct stage because analytics then helps organizations explain why things are happening and identify patterns, opportunities, and risks.

Business Optimization is another valid stage because mature organizations use analytics and AI to improve processes, decisions, resources, customer experiences, and operational outcomes. Cultural Transformation is also part of maturity because long-term AI and data success requires a shift in mindset, leadership behavior, decision-making culture, and enterprise-wide adoption of data-driven practices.

Since all listed options represent stages or maturity areas in the Data and AI Analytics Business Model Maturity Index, the correct answer is E. All of the above.

Question 6

Question Type: MultipleChoice

Which of the following is a regulatory requirement?

Options:

A- HIPAA
B- GDPR
C- CCPA
D- a and b only
E- a, b and c only

Answer:

E

Explanation:

The correct answer is E. a, b and c only because HIPAA, GDPR, and CCPA are all regulatory requirements that organizations may need to follow when collecting, processing, storing, sharing, or protecting data. In AI and data-driven business environments, regulatory compliance is important because AI systems often depend on large volumes of personal, sensitive, customer, health, or behavioral data.

HIPAA is a United States healthcare regulation that protects sensitive patient health information. GDPR is a European Union data protection regulation that governs personal data processing, privacy rights, consent, data minimization, and accountability. CCPA is a California privacy law that gives consumers rights related to personal information, including access, deletion, and opt-out rights. These regulations affect how organizations design AI systems, manage data governance, ensure privacy, reduce risk, and maintain responsible AI practices.

Since all three listed items are valid regulatory requirements, the correct answer is E. a, b and c

only.

Question 7

Question Type: MultipleChoice

Select the BEST choice for ML solutions architecture coverage.

Options:

A- Business understanding
B- Identification and verification of ML techniques
C- System architecture of the ML technology platform
D- a and b only
E- a, b and c only

Answer:

E

Explanation:

The correct answer is E. a, b and c only because ML solution architecture must cover the complete path from business need to technical implementation. Business understanding is essential because an ML solution should begin with a clear problem statement, business objective, success criteria, expected value, and operational impact. Without business understanding, the model may solve the wrong problem or fail to create measurable value.

Identification and verification of ML techniques are also part of ML solution architecture because teams must choose suitable algorithms, validate model approaches, compare methods, and confirm that the selected technique fits the data, use case, performance expectations, and business constraints. System architecture of the ML technology platform is equally important because ML solutions require data pipelines, infrastructure, compute resources, model deployment environments, monitoring, security, scalability, and integration with enterprise systems.

Since all three areas are important parts of ML solution architecture coverage, the best answer is E.

Question 8

Question Type: MultipleChoice

If humans are unlabeling the data and the machine is correctly labeling current or future data points, it's _____.

Options:

A- supervised learning
B- unsupervised learning
C- reinforcement learning
D- semi-reinforcement learning
E- semi-supervised learning

P2P exams

Answer:

E

Explanation:

Semi-supervised learning is the correct answer because it combines a small amount of labeled data with a larger amount of unlabeled data. In this scenario, humans are not fully labeling the data, but the machine is still able to correctly label current or future data points by learning patterns from the available data. That matches the concept of semi-supervised learning, where the model uses limited human-provided labels and extends learning to unlabeled examples.

Supervised learning is not the best answer because supervised learning depends on clearly labeled training data supplied by humans. Unsupervised learning is also incorrect because it identifies hidden patterns or clusters without using labels, rather than predicting correct labels for future data points. Reinforcement learning is based on rewards, penalties, actions, and an environment, which is not described here. ''Semi-reinforcement learning'' is not a standard main category in machine learning.

Therefore, the most accurate answer is **E. Semi-supervised learning**.

Question 9

Question Type: MultipleChoice

Unsupervised learning is a type of machine learning where the algorithm learns from a _____.

Options:

A- labeled dataset
B- Unlabeled dataset
C- explained dataset
D- a and b only
E- a and c only

Answer:

B

Explanation:

The correct answer is B. Unlabeled dataset. Unsupervised learning is a machine learning approach where the algorithm works with data that does not contain predefined labels, target outputs, or correct answers. Instead of being told what each data point represents, the model analyzes the structure of the data and identifies hidden patterns, groupings, similarities, or relationships on its own.

This type of learning is commonly used for clustering, association rule mining, anomaly detection, dimensionality reduction, and customer segmentation. For example, an unsupervised learning model may group customers based on buying behavior without being given category labels in advance.

A labeled dataset is used in supervised learning, where the model learns from input-output pairs. An ''explained dataset'' is not a standard machine learning category. Since unsupervised learning specifically depends on unlabeled data, the correct answer is B. Unlabeled dataset.

Question 10

Question Type: MultipleChoice

Which of the following is the CORRECT key areas as ethical principles?

Options:

A- Respect for human autonomy
B- Prevention of harm
C- Explicability
D- a and c only

E- a, b and c only

Answer:

E

Explanation:

The correct answer is E. a, b and c only because respect for human autonomy, prevention of harm, and explicability are all recognized ethical principles in responsible AI. Respect for human autonomy means AI systems should support human decision-making rather than unfairly manipulate, replace, or override people in ways that remove meaningful human control. This is especially important in business, healthcare, finance, hiring, and other high-impact AI use cases.

Prevention of harm is also a core ethical principle because AI systems should be designed and deployed to reduce physical, psychological, financial, social, operational, and reputational risks. Organizations must consider safety, reliability, misuse prevention, bias reduction, and risk controls.

Explicability is correct because AI decisions should be understandable, explainable, and auditable where appropriate. Stakeholders should be able to understand how and why an AI system produces important outputs. Since all three listed items are valid ethical principles, the correct answer is E. a, b and c only.

To Get Premium Files for CAIC Visit

https://www.p2pexams.com/products/caic

For More Free Questions Visit

https://www.p2pexams.com/usaii/pdf/caic

