SISA CSPAI Practice Questions

Shared by Burris on 09-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

In a time-series prediction task, how does an RNN effectively model sequential data?

Options:

A- By focusing on the overall sequence structure rather than individual time steps for a more holistic approach.
B- By processing each time step independently, optimizing the model's performance over time.
C- By storing only the most recent time step, ensuring efficient memory usage for real-time predictions
D- By using hidden states to retain context from prior time steps, allowing it to capture dependencies across the sequence.

Answer:

D

Question 2

Question Type: MultipleChoice

How does AI enhance customer experience in retail environments?

Options:

A- By integrating personalized interactions with AI-driven analytics for a more customized shopping experience.
B- By optimizing customer service through automated systems and tailored recommendations.
C- By ensuring every customer receives the same generic response from automated systems.
D- By automating repetitive tasks and providing consistent data driven insights to improve customer service.

Answer:

A

Question 3

Question Type: MultipleChoice

In a Transformer model processing a sequence of text for a translation task, how does incorporating positional encoding impact the model's ability to generate accurate translations?

Options:

A- It ensures that the model treats all words as equally important, regardless of their position in the sequence.
B- It simplifies the model's computations by merging all words into a single representation, regardless of their order
C- It speeds up processing by reducing the number of tokens the model needs to handle.
D- It helps the model distinguish the order of words in the sentence, leading to more accurate translation by maintaining the context of each word's position.

Answer:

D

Question 4

Question Type: MultipleChoice

What is a potential risk associated with hallucinations in LLMs, and how should it be addressed to ensure Responsible AI?

Options:

A- Hallucinations can lead to creative outputs, which are beneficial for all applications; hence, no measures are necessary.
B- Hallucinations cause models to slow down; optimizing hardware performance is necessary to mitigate this issue.
C- Hallucinations can produce inaccurate or misleading information; it should be addressed by incorporating external knowledge bases and retrieval systems.
D- Hallucinations are primarily due to overfitting; regularization techniques should be applied during training.

Answer:

C

Question 5

Question Type: MultipleChoice

What is a key benefit of using GenAI for security analytics?

Options:

A- Increasing data silos to protect information.
B- Predicting future threats through pattern recognition in large datasets.
C- Limiting analysis to historical data only.
D- Reducing the use of analytics tools to save costs.

Answer:

B

Question 6

Question Type: MultipleChoice

When integrating LLMs using a Prompting Technique, what is a significant challenge in achieving consistent performance across diverse applications?

Options:

A- Handling the security concerns that arise from dynamically generated prompts
B- Overcoming the lack of transparency in understanding how the LLM interprets varying prompt structures.
C- The need for optimizing prompt templates to ensure generalization across different contexts.
D- Reducing latency in generating responses to meet real-time application requirements.

Answer:

C

Question 7

Question Type: MultipleChoice

In the Retrieval-Augmented Generation (RAG) framework, Which option best is the most critical factor for improving factual consistency in generated outputs?

Options:

A- Fine-tuning the generative model with synthetic datasets generated from the retrieved documents
B- Utilising an ensemble of multiple LLMs to cross-check the generated outputs.
C- Implementing a redundancy check by comparing the outputs from different retrieval modules.
D- Tuning the retrieval model to prioritize documents with the highest semantic similarity

Answer:

D

Question 8

Question Type: MultipleChoice

Which option best is a primary goal of enforcing Responsible AI standards and regulations in the development and deployment of LLMs?

Options:

A- Maximizing model performance while minimizing computational costs.
B- Developing AI systems with the highest accuracy regardless of data privacy concerns
C- Focusing solely on improving the speed and scalability of AI systems
D- Ensuring that AI systems operate safely, ethically, and without causing harm.

Answer:

D

Question 9

Question Type: MultipleChoice

In line with the US Executive Order on AI, a company's AI application has encountered a

security vulnerability. What should be prioritized to align with the order's expectations?

Options:

A- Implementing a rapid response to address and remediate the vulnerability, followed by a review of security practices.
B- Immediate public disclosure of the vulnerability.
C- Halting all AI projects until a full investigation is complete.
D- Ignoring the vulnerability if it does not affect core functionalities.

Answer:

A

P2P exams

Question 10

Question Type: MultipleChoice

How does the multi-head self-attention mechanism improve the model's ability to learn complex relationships in data?

Options:

A- By forcing the model to focus on a single aspect of the input at a time.
B- By ensuring that the attention mechanism looks only at local context within the input
C- By simplifying the network by removing redundancy in attention layers.
D- By allowing the model to focus on different parts of the input through multiple attention heads

P2P exams

Answer:

D

To Get Premium Files for CSPAI Visit

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

For More Free Questions Visit

https://www.p2pexams.com/sisa/pdf/cspai

