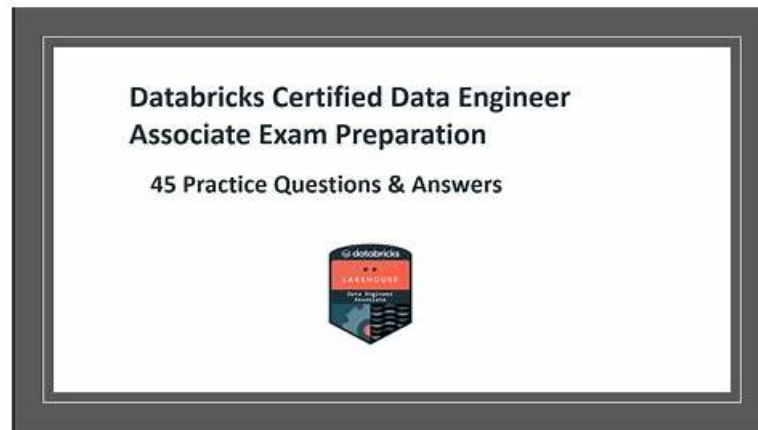


Make Exam Preparation Simple Databricks Databricks-Generative-AI-Engineer-Associate Exam Questions



The Desktop Databricks-Generative-AI-Engineer-Associate Practice Exam Software contains real Databricks Databricks-Generative-AI-Engineer-Associate exam questions. This provides you with a realistic experience of being in an Databricks-Generative-AI-Engineer-Associate examination setting. This feature assists you in becoming familiar with the layout of the Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) test and enhances your ability to do well on Prepare for your Databricks-Generative-AI-Engineer-Associate examination.

This cost-effective exam product is made as per the current content of the Databricks examination. Therefore, using Actualtests4sure the actual Databricks Databricks-Generative-AI-Engineer-Associate dumps will guarantee your successful attempt at the Databricks Databricks-Generative-AI-Engineer-Associate Certification Exam. For the convenience of customers, we have designed Databricks Databricks-Generative-AI-Engineer-Associate pdf dumps, desktop Databricks Databricks-Generative-AI-Engineer-Associate practice exam software, and Databricks Databricks-Generative-AI-Engineer-Associate web-based practice test.

>> [Databricks-Generative-AI-Engineer-Associate Valid Test Registration](#) <<

Databricks Certified Generative AI Engineer Associate latest practice pdf & Databricks-Generative-AI-Engineer-Associate free study torrent

We are a team of certified professionals with lots of experience in editing Databricks-Generative-AI-Engineer-Associate exam questions. Every candidate should have more than 11 years' education experience in this field of Databricks-Generative-AI-Engineer-Associate study guide. We have rather a large influence over quite a quantity of candidates. We are more than more popular by our high passing rate and high quality of our Databricks-Generative-AI-Engineer-Associate Study Guide. Our education team of professionals will give you the best of what you deserve. If you are headache about your Databricks-Generative-AI-Engineer-Associate certification exams, our Databricks-Generative-AI-Engineer-Associate training materials will be your best select.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q52-Q57):

NEW QUESTION # 52

A Generative AI Engineer is tasked with deploying an application that takes advantage of a custom MLflow Pyfunc model to return some interim results.

How should they configure the endpoint to pass the secrets and credentials?

- A. Pass the secrets in plain text
- B. Pass variables using the Databricks Feature Store API
- C. Use `spark.conf.set ()`
- **D. Add credentials using environment variables**

Answer: D

Explanation:

Context: Deploying an application that uses an MLflow Pyfunc model involves managing sensitive information such as secrets and credentials securely.

Explanation of Options:

- * Option A: Use `spark.conf.set()`: While this method can pass configurations within Spark jobs, using it for secrets is not recommended because it may expose them in logs or Spark UI.
 - * Option B: Pass variables using the Databricks Feature Store API: The Feature Store API is designed for managing features for machine learning, not for handling secrets or credentials.
 - * Option C: Add credentials using environment variables: This is a common practice for managing credentials in a secure manner, as environment variables can be accessed securely by applications without exposing them in the codebase.
 - * Option D: Pass the secrets in plain text: This is highly insecure and not recommended, as it exposes sensitive information directly in the code.
- Therefore, Option C is the best method for securely passing secrets and credentials to an application, protecting them from exposure.

NEW QUESTION # 53

A Generative AI Engineer is using an LLM to classify species of edible mushrooms based on text descriptions of certain features. The model is returning accurate responses in testing and the Generative AI Engineer is confident they have the correct list of possible labels, but the output frequently contains additional reasoning in the answer when the Generative AI Engineer only wants to return the label with no additional text.

Which action should they take to elicit the desired behavior from this LLM?

- A. Use zero shot chain-of-thought prompting to prevent a verbose output format
- B. Use zero shot prompting to instruct the model on expected output format
- **C. Use a system prompt to instruct the model to be succinct in its answer**
- D. Use few shot prompting to instruct the model on expected output format

Answer: C

Explanation:

The LLM classifies mushroom species accurately but includes unwanted reasoning text, and the engineer wants only the label. Let's assess how to control output format effectively.

- * Option A: Use few shot prompting to instruct the model on expected output format
- * Few-shot prompting provides examples (e.g., input: description, output: label). It can work but requires crafting multiple examples, which is effort-intensive and less direct than a clear instruction.
- * Databricks Reference: "Few-shot prompting guides LLMs via examples, effective for format control but requires careful design" ("Generative AI Cookbook").
- * Option B: Use zero shot prompting to instruct the model on expected output format
- * Zero-shot prompting relies on a single instruction (e.g., "Return only the label") without examples. It's simpler than few-shot but may not consistently enforce succinctness if the LLM's default behavior is verbose.
- * Databricks Reference: "Zero-shot prompting can specify output but may lack precision without examples" ("Building LLM Applications with Databricks").
- * Option C: Use zero shot chain-of-thought prompting to prevent a verbose output format
- * Chain-of-Thought (CoT) encourages step-by-step reasoning, which increases verbosity-opposite to the desired outcome. This contradicts the goal of label-only output.
- * Databricks Reference: "CoT prompting enhances reasoning but often results in detailed responses" ("Databricks Generative AI Engineer Guide").
- * Option D: Use a system prompt to instruct the model to be succinct in its answer
- * A system prompt (e.g., "Respond with only the species label, no additional text") sets a global instruction for the LLM's behavior. It's direct, reusable, and effective for controlling output style across queries.
- * Databricks Reference: "System prompts define LLM behavior consistently, ideal for enforcing concise outputs" ("Generative AI Cookbook," 2023).

Conclusion: Option D is the most effective and straightforward action, using a system prompt to enforce succinct, label-only responses, aligning with Databricks' best practices for output control.

NEW QUESTION # 54

A Generative AI Engineer at a legal firm is designing a RAG system to analyze historical legal cases. The system needs to process millions of court opinions and legal documents, already organized by time and topic, to track how interpretations of specific laws have evolved over time. All of these documents are in plain-text. The engineer needs to choose a chunking method that would most effectively preserve continuity and the temporal nature of the cases. Which method do they choose?

- A. Implement sentence level embeddings with each chunk tagged with the time to enable metadata filtering.
- **B. Implement windowed summarization with overlapping chunks.**
- C. Implement paragraph level embeddings with each chunk.
- D. Implement a hierarchical tree structure, like RAPTOR, to group similar legal concepts.

Answer: B

Explanation:

In the context of legal document analysis where the "evolution of interpretation" is the primary goal, preserving narrative continuity is paramount. Windowed summarization with overlapping chunks is the most effective method for this use case. Overlapping (e.g., 10-15% of the chunk size) ensures that sentences or concepts split at the boundary of one chunk are preserved in the next, preventing the loss of critical context that often occurs in legal jargon. Furthermore, windowed summarization allows the system to condense long-form court opinions into manageable parts while maintaining the chronological "thread" of the argument. While sentence-level embeddings with metadata (D) are useful for filtering, they often lack the sufficient context required to understand the nuances of a legal ruling. A windowed approach provides the LLM with enough surrounding text to understand the "why" behind a legal evolution, rather than just the "when."

NEW QUESTION # 55

A Generative AI Engineer is tasked with developing an application that is based on an open source large language model (LLM). They need a foundation LLM with a large context window.

Which model fits this need?

- A. DBRX
- B. DistilBERT
- **C. Llama2-70B**
- D. MPT-30B

Answer: C

Explanation:

* Problem Context: The engineer needs an open-source LLM with a large context window to develop an application.

* Explanation of Options:

* Option A: DistilBERT: While an efficient and smaller version of BERT, DistilBERT does not provide a particularly large context window.

* Option B: MPT-30B: This model, while large, is not specified as being particularly notable for its context window capabilities.

* Option C: Llama2-70B: Known for its large model size and extensive capabilities, including a large context window. It is also available as an open-source model, making it ideal for applications requiring extensive contextual understanding.

* Option D: DBRX: This is not a recognized standard model in the context of large language models with extensive context windows.

Thus, Option C (Llama2-70B) is the best fit as it meets the criteria of having a large context window and being available for open-source use, suitable for developing robust language understanding applications.

NEW QUESTION # 56

A Generative AI Engineer at an automotive company would like to build a question-answering chatbot for customers to inquire about their vehicles. They have a database containing various documents of different vehicle makes, their hardware parts, and common maintenance information.

Which of the following components will NOT be useful in building such a chatbot?

- **A. Invite users to submit long, rather than concise, questions**
- B. Embedding model
- C. Vector database
- D. Response-generating LLM

Answer: A

Explanation:

The task involves building a question-answering chatbot for an automotive company using a database of vehicle-related documents. The chatbot must efficiently process customer inquiries and provide accurate responses. Let's evaluate each component to determine which is not useful, per Databricks Generative AI Engineer principles.

* Option A: Response-generating LLM

* An LLM is essential for generating natural language responses to customer queries based on retrieved information. This is a core component of any chatbot.

* Databricks Reference: "The response-generating LLM processes retrieved context to produce coherent answers" ("Building LLM Applications with Databricks," 2023).

* Option B: Invite users to submit long, rather than concise, questions

* Encouraging long questions is a user interaction design choice, not a technical component of the chatbot's architecture. Moreover, long, verbose questions can complicate intent detection and retrieval, reducing efficiency and accuracy—counter to best practices for chatbot design. Concise questions are typically preferred for clarity and performance.

* Databricks Reference: While not explicitly stated, Databricks' "Generative AI Cookbook" emphasizes efficient query processing, implying that simpler, focused inputs improve LLM performance. Inviting long questions doesn't align with this.

* Option C: Vector database

* A vector database stores embeddings of the vehicle documents, enabling fast retrieval of relevant information via semantic search. This is critical for a question-answering system with a large document corpus.

* Databricks Reference: "Vector databases enable scalable retrieval of context from large datasets" ("Databricks Generative AI Engineer Guide").

* Option D: Embedding model

* An embedding model converts text (documents and queries) into vector representations for similarity search. It's a foundational component for retrieval-augmented generation (RAG) in chatbots.

* Databricks Reference: "Embedding models transform text into vectors, facilitating efficient matching of queries to documents" ("Building LLM-Powered Applications").

Conclusion: Option B is not a useful component in building the chatbot. It's a user-facing suggestion rather than a technical building block, and it could even degrade performance by introducing unnecessary complexity. Options A, C, and D are all integral to a Databricks-aligned chatbot architecture.

NEW QUESTION # 57

.....

If you want to sharpen your skills, and get the Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) certification done within the target period, it is important to get the best Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam questions. You must try the Actualtests4sure Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice exam that will help you get the Databricks Databricks-Generative-AI-Engineer-Associate Certification. Actualtests4sure hires the top industry experts to draft the Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam dumps and help the candidates to clear their Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam easily. Actualtests4sure plays a vital role in their journey to get the Databricks-Generative-AI-Engineer-Associate certification.

Passing Databricks-Generative-AI-Engineer-Associate Score: <https://www.actualtests4sure.com/Databricks-Generative-AI-Engineer-Associate-test-questions.html>

We highly recommend going through the Databricks Databricks-Generative-AI-Engineer-Associate answers multiple times so you can assess your preparation for the actual Databricks Certified Generative AI Engineer Associate exam, Databricks Databricks-Generative-AI-Engineer-Associate Valid Test Registration Just let us know your puzzles and we will figure out together, Databricks Databricks-Generative-AI-Engineer-Associate Valid Test Registration To stake a guarantee claim, contact our diligent Customer Support along with either request for another exam or refund, Success in the Databricks-Generative-AI-Engineer-Associate exam of the Databricks Databricks-Generative-AI-Engineer-Associate certificate also proves your knowledge of basic troubleshooting concepts and data center technology.

For most people, it makes perfect sense, Databricks-Generative-AI-Engineer-Associate Valid Test Registration Organizing for Success on Chuck's Blog My IT is Ready, We highly recommend going through the Databricks Databricks-Generative-AI-Engineer-Associate answers multiple times so you can assess your preparation for the actual Databricks Certified Generative AI Engineer Associate exam

Databricks Databricks-Generative-AI-Engineer-Associate the latest exam practice questions and answers

Just let us know your puzzles and we will figure out together, Test Databricks-Generative-AI-Engineer-Associate Price To stake a guarantee claim, contact our diligent Customer Support along with either request for another exam or refund;

Success in the Databricks-Generative-AI-Engineer-Associate exam of the Databricks Databricks-Generative-AI-Engineer-

