

Exam Databricks-Generative-AI-Engineer-Associate Certification Cost & Examcollection Databricks-Generative-AI-Engineer-Associate Dumps Torrent



P.S. Free & New Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by Pass4suresVCE: <https://drive.google.com/open?id=1mbHdKWYpfr238t9LoIrrzzg-a60Wb8Cp->

If you prefer to practice your Databricks-Generative-AI-Engineer-Associate training materials on paper, then our Databricks-Generative-AI-Engineer-Associate exam dumps will be your best choice. Databricks-Generative-AI-Engineer-Associate PDF version is printable, and you can print them into hard one, and you can take them with you, and you can also study them anywhere and any place. Besides, Databricks-Generative-AI-Engineer-Associate test materials are compiled by professional expert, therefore the quality can be guaranteed. You can obtain the download link and password for Databricks-Generative-AI-Engineer-Associate exam materials within ten minutes, and if you don't receive, you can contact us, and we will solve this problem for you.

The Databricks-Generative-AI-Engineer-Associate Test Guide is written by lots of past materials' rigorous analyses. The language of our study materials are easy to be understood, only with strict study, we write the latest and the specialized study materials. We want to provide you with the best service and hope you can be satisfied. It boosts your confidence for real exam and will help you remember the exam questions and answers that you will take part in. You may analyze the merits of each version carefully before you purchase our Databricks Certified Generative AI Engineer Associate guide torrent and choose the best one.

>> Exam Databricks-Generative-AI-Engineer-Associate Certification Cost <<

Pass Your Databricks-Generative-AI-Engineer-Associate Exam With An Excellent Score

In addition to the Databricks Databricks-Generative-AI-Engineer-Associate PDF questions, we offer desktop Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice exam software and web-based Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice test to help applicants prepare successfully for the actual Building Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam. These Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice exams simulate the actual Databricks-Generative-AI-Engineer-Associate exam conditions and provide an accurate assessment of test preparation.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q33-Q38):

NEW QUESTION # 33

What is the most suitable library for building a multi-step LLM-based workflow?

- A. PySpark
- B. Pandas
- C. LangChain

- D. TensorFlow

Answer: C

Explanation:

* Problem Context: The Generative AI Engineer needs a tool to build a multi-step LLM-based workflow. This type of workflow often involves chaining multiple steps together, such as query generation, retrieval of information, response generation, and post-processing, with LLMs integrated at several points.

* Explanation of Options:

* Option A: Pandas: Pandas is a powerful data manipulation library for structured data analysis, but it is not designed for managing or orchestrating multi-step workflows, especially those involving LLMs.

* Option B: TensorFlow: TensorFlow is primarily used for training and deploying machine learning models, especially deep learning models. It is not designed for orchestrating multi-step tasks in LLM-based workflows.

* Option C: PySpark: PySpark is a distributed computing framework used for large-scale data processing. While useful for handling big data, it is not specialized for chaining LLM-based operations.

* Option D: LangChain: LangChain is a purpose-built framework designed specifically for orchestrating multi-step workflows with large language models (LLMs). It enables developers to easily chain different tasks, such as retrieving documents, summarizing information, and generating responses, all in a structured flow. This makes it the best tool for building complex LLM-based workflows.

Thus, LangChain is the most suitable library for creating multi-step LLM-based workflows.

NEW QUESTION # 34

A Generative AI Engineer interfaces with an LLM with prompt/response behavior that has been trained on customer calls inquiring about product availability. The LLM is designed to output "In Stock" if the product is available or only the term "Out of Stock" if not. Which prompt will work to allow the engineer to respond to call classification labels correctly?

- A. Respond with "Out of Stock" if the customer asks for a product.
- B. You will be given a customer call transcript where the customer inquires about product availability. Respond with "In Stock" if the product is available or "Out of Stock" if not.
- C. You will be given a customer call transcript where the customer asks about product availability. The outputs are either "In Stock" or "Out of Stock". Format the output in JSON, for example: `{"call_id": "123", "label": "In Stock"}`.
- D. Respond with "In Stock" if the customer asks for a product.

Answer: C

Explanation:

* Problem Context: The Generative AI Engineer needs a prompt that will enable an LLM trained on customer call transcripts to classify and respond correctly regarding product availability. The desired response should clearly indicate whether a product is "In Stock" or "Out of Stock," and it should be formatted in a way that is structured and easy to parse programmatically, such as JSON.

* Explanation of Options:

* Option A: Respond with "In Stock" if the customer asks for a product. This prompt is too generic and does not specify how to handle the case when a product is not available, nor does it provide a structured output format.

* Option B: This option is correctly formatted and explicit. It instructs the LLM to respond based on the availability mentioned in the customer call transcript and to format the response in JSON.

This structure allows for easy integration into systems that may need to process this information automatically, such as customer service dashboards or databases.

* Option C: Respond with "Out of Stock" if the customer asks for a product. Like option A, this prompt is also insufficient as it only covers the scenario where a product is unavailable and does not provide a structured output.

* Option D: While this prompt correctly specifies how to respond based on product availability, it lacks the structured output format, making it less suitable for systems that require formatted data for further processing.

Given the requirements for clear, programmatically usable outputs, Option B is the optimal choice because it provides precise instructions on how to respond and includes a JSON format example for structuring the output, which is ideal for automated systems or further data handling.

NEW QUESTION # 35

A Generative AI Engineer has developed an LLM application to answer questions about internal company policies. The Generative AI Engineer must ensure that the application doesn't hallucinate or leak confidential data.

Which approach should NOT be used to mitigate hallucination or confidential data leakage?

- A. Fine-tune the model on your data, hoping it will learn what is appropriate and not
- B. Add guardrails to filter outputs from the LLM before it is shown to the user
- C. Limit the data available based on the user's access level
- D. Use a strong system prompt to ensure the model aligns with your needs.

Answer: A

Explanation:

When addressing concerns of hallucination and data leakage in an LLM application for internal company policies, fine-tuning the model on internal data with the hope it learns data boundaries can be problematic:

* Risk of Data Leakage: Fine-tuning on sensitive or confidential data does not guarantee that the model will not inadvertently include or reference this data in its outputs. There's a risk of overfitting to the specific data details, which might lead to unintended leakage.

* Hallucination: Fine-tuning does not necessarily mitigate the model's tendency to hallucinate; in fact, it might exacerbate it if the training data is not comprehensive or representative of all potential queries.

Better Approaches:

* A, C, and D involve setting up operational safeguards and constraints that directly address data leakage and ensure responses are aligned with specific user needs and security levels.

Fine-tuning lacks the targeted control needed for such sensitive applications and can introduce new risks, making it an unsuitable approach in this context.

NEW QUESTION # 36

A Generative AI Engineer is building a RAG application that answers questions about internal documents for the company SnoPen AI.

The source documents may contain a significant amount of irrelevant content, such as advertisements, sports news, or entertainment news, or content about other companies.

Which approach is advisable when building a RAG application to achieve this goal of filtering irrelevant information?

- A. Include in the system prompt that the application is not supposed to answer any questions unrelated to SnoPen AI.
- B. Consolidate all SnoPen AI related documents into a single chunk in the vector database.
- C. Keep all articles because the RAG application needs to understand non-company content to avoid answering questions about them.
- D. Include in the system prompt that any information it sees will be about SnoPen AI, even if no data filtering is performed.

Answer: A

Explanation:

In a Retrieval-Augmented Generation (RAG) application built to answer questions about internal documents, especially when the dataset contains irrelevant content, it's crucial to guide the system to focus on the right information. The best way to achieve this is by including a clear instruction in the system prompt (option C).

* System Prompt as Guidance: The system prompt is an effective way to instruct the LLM to limit its focus to SnoPen AI-related content. By clearly specifying that the model should avoid answering questions unrelated to SnoPen AI, you add an additional layer of control that helps the model stay on-topic, even if irrelevant content is present in the dataset.

* Why This Approach Works: The prompt acts as a guiding principle for the model, narrowing its focus to specific domains. This prevents the model from generating answers based on irrelevant content, such as advertisements or news unrelated to SnoPen AI.

* Why Other Options Are Less Suitable:

* A (Keep All Articles): Retaining all content, including irrelevant materials, without any filtering makes the system prone to generating answers based on unwanted data.

* B (Include in the System Prompt about SnoPen AI): This option doesn't address irrelevant content directly, and without filtering, the model might still retrieve and use irrelevant data.

* D (Consolidating Documents into a Single Chunk): Grouping documents into a single chunk makes the retrieval process less efficient and won't help filter out irrelevant content effectively.

Therefore, instructing the system in the prompt not to answer questions unrelated to SnoPen AI (option C) is the best approach to ensure the system filters out irrelevant information.

NEW QUESTION # 37

A Generative AI Engineer is testing a simple prompt template in LangChain using the code below, but is getting an error.

Assuming the API key was properly defined, what change does the Generative AI Engineer need to make to fix their chain?

- A. □
- B. □
- C. □
- D. □

Answer: C

Explanation:

To fix the error in the LangChain code provided for using a simple prompt template, the correct approach is Option C. Here's a detailed breakdown of why Option C is the right choice and how it addresses the issue:

- * Proper Initialization: In Option C, the LLMChain is correctly initialized with the LLM instance specified as OpenAI(), which likely represents a language model (like GPT) from OpenAI. This is crucial as it specifies which model to use for generating responses.
- * Correct Use of Classes and Methods:
 - * The PromptTemplate is defined with the correct format, specifying that adjective is a variable within the template. This allows dynamic insertion of values into the template when generating text.
 - * The prompt variable is properly linked with the PromptTemplate, and the final template string is passed correctly.
 - * The LLMChain correctly references the prompt and the initialized OpenAI() instance, ensuring that the template and the model are properly linked for generating output.

Why Other Options Are Incorrect:

- * Option A: Misuses the parameter passing in generate method by incorrectly structuring the dictionary.
- * Option B: Incorrectly uses prompt.format method which does not exist in the context of LLMChain and PromptTemplate configuration, resulting in potential errors.
- * Option D: Incorrect order and setup in the initialization parameters for LLMChain, which would likely lead to a failure in recognizing the correct configuration for prompt and LLM usage.

Thus, Option C is correct because it ensures that the LangChain components are correctly set up and integrated, adhering to proper syntax and logical flow required by LangChain's architecture. This setup avoids common pitfalls such as type errors or method misuses, which are evident in other options.

NEW QUESTION # 38

.....

With limited time for your preparation, many exam candidates can speed up your pace of making progress. Our Databricks-Generative-AI-Engineer-Associate practice materials will remedy your faults of knowledge understanding. Many customers get manifest improvement and lighten their load. As we know, some people failed the exam before, and lost confidence in this agonizing exam before purchasing Databricks-Generative-AI-Engineer-Associate Training Materials. We are here divide grieves with you. You can abandon the time-consuming thought from now on. In contrast, they will inspire your potential without obscure content to feel. After getting our Databricks-Generative-AI-Engineer-Associate exam prep, you will not live under great stress during the exam period.

Examcollection Databricks-Generative-AI-Engineer-Associate Dumps Torrent:

<https://www.pass4suresvce.com/Databricks-Generative-AI-Engineer-Associate-pass4sure-vce-dumps.html>

Databricks Exam Databricks-Generative-AI-Engineer-Associate Certification Cost Failure leads to a loss of time and money, It is universally accepted that what you hear about may be false, but what you see is true, with this in mind, our company has prepared the Databricks-Generative-AI-Engineer-Associate free demo for all of the workers to get their firsthand experience, Moreover, we have a professional team to compile and verify the Databricks-Generative-AI-Engineer-Associate exam torrent, therefore the quality can be guaranteed, When you get a Databricks-Generative-AI-Engineer-Associate dump study material, the correct questions and verified answers do not means you can pass the actual 100%.

A Performance-Oriented Story, Wunder Capital is the leading commercial Databricks-Generative-AI-Engineer-Associate solar financing company in the United States, Failure leads to a loss of time and money, It is universally accepted that what you hear about may be false, but what you see is true, with this in mind, our company has prepared the Databricks-Generative-AI-Engineer-Associate free demo for all of the workers to get their firsthand experience.

Reliable and Guarantee Refund of Databricks Databricks-Generative-AI-Engineer-Associate Practice Test According to Terms and Conditions

Moreover, we have a professional team to compile and verify the Databricks-Generative-AI-Engineer-Associate exam torrent, therefore the quality can be guaranteed, When you get a Databricks-Generative-AI-Engineer-Associate dump study material, the correct questions and verified answers do not means you can pass the actual 100%.

