

To Get Brilliant Success Databricks Databricks-Generative-AI-Engineer-Associate Questions



P.S. Free 2025 Databricks Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by Itbraindumps: <https://drive.google.com/open?id=1ZKCNx4m8tD4KcJTOCiYUjoMMcPmG84k3>

The candidates can benefit themselves by using our Databricks-Generative-AI-Engineer-Associate test engine and get a lot of test questions like exercises and answers. Our Databricks-Generative-AI-Engineer-Associate exam questions will help them modify the entire syllabus in a short time. And the Software version of our Databricks-Generative-AI-Engineer-Associate Study Materials have the advantage of simulating the real exam, so that the candidates have more experience of the practicing the real exam questions.

Just as an old saying goes, it is better to gain a skill than to be rich. Contemporarily, competence far outweighs family backgrounds and academic degrees. One of the significant factors to judge whether one is competent or not is his or her Databricks-Generative-AI-Engineer-Associate certificates. Generally speaking, Databricks-Generative-AI-Engineer-Associate certificates function as the fundamental requirement when a company needs to increase manpower in its start-up stage. In this respect, our Databricks-Generative-AI-Engineer-Associate practice materials can satisfy your demands if you are now in preparation for a Databricks-Generative-AI-Engineer-Associate certificate.

>> Valid Databricks-Generative-AI-Engineer-Associate Exam Syllabus <<

Study Guide Databricks-Generative-AI-Engineer-Associate Pdf - Databricks-Generative-AI-Engineer-Associate Test Dumps Pdf

The company is preparing for the test candidates to prepare the Databricks-Generative-AI-Engineer-Associate exam guide professional brand, designed to be the most effective and easiest way to help users through their want to get the test Databricks-Generative-AI-Engineer-Associate certification and obtain the relevant certification. In comparison with similar educational products, our Databricks-Generative-AI-Engineer-Associate Training Materials are of superior quality and reasonable price, so our company has become the top enterprise in the international market. Our Databricks-Generative-AI-Engineer-Associate practice materials have been well received mainly for the advantage of high pass rate as 99% to 100%.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q62-Q67):

NEW QUESTION # 62

A Generative AI Engineer is developing a patient-facing healthcare-focused chatbot. If the patient's question is not a medical emergency, the chatbot should solicit more information from the patient to pass to the doctor's office and suggest a few relevant pre-approved medical articles for reading. If the patient's question is urgent, direct the patient to calling their local emergency services.

Given the following user input:

"I have been experiencing severe headaches and dizziness for the past two days." Which response is most appropriate for the chatbot to generate?

- A. Please call your local emergency services.
- B. Headaches can be tough. Hope you feel better soon!
- C. Please provide your age, recent activities, and any other symptoms you have noticed along with your headaches and dizziness.
- D. Here are a few relevant articles for your browsing. Let me know if you have questions after reading them.

Answer: A

Explanation:

* Problem Context: The task is to design responses for a healthcare-focused chatbot that appropriately addresses the urgency of a patient's symptoms.

* Explanation of Options:

* Option A: Suggesting articles might be suitable for less urgent inquiries but is inappropriate for symptoms that could indicate a serious condition.

* Option B: Given the description of severe symptoms like headaches and dizziness, directing the patient to emergency services is prudent. This aligns with medical guidelines that recommend immediate professional attention for such severe symptoms.

* Option C: Offering well-wishes does not address the potential seriousness of the symptoms and lacks appropriate action.

* Option D: While gathering more information is part of a detailed assessment, the immediate need here suggests a more urgent response.

Given the potential severity of the described symptoms, Option B is the most appropriate, ensuring the chatbot directs patients to seek urgent care when needed, potentially saving lives.

NEW QUESTION # 63

A Generative AI Engineer is building a RAG application that will rely on context retrieved from source documents that are currently in PDF format. These PDFs can contain both text and images. They want to develop a solution using the least amount of lines of code.

Which Python package should be used to extract the text from the source documents?

- A. flask
- B. numpy
- C. beautifulsoup
- D. unstructured

Answer: C

Explanation:

* Problem Context: The engineer needs to extract text from PDF documents, which may contain both text and images. The goal is to find a Python package that simplifies this task using the least amount of code.

* Explanation of Options:

* Option A: flask: Flask is a web framework for Python, not suitable for processing or extracting content from PDFs.

* Option B: beautifulsoup: Beautiful Soup is designed for parsing HTML and XML documents, not PDFs.

* Option C: unstructured: This Python package is specifically designed to work with unstructured data, including extracting text from PDFs. It provides functionalities to handle various types of content in documents with minimal coding, making it ideal for the task.

* Option D: numpy: Numpy is a powerful library for numerical computing in Python and does not provide any tools for text extraction from PDFs.

Given the requirement, Option C (unstructured) is the most appropriate as it directly addresses the need to efficiently extract text from PDF documents with minimal code.

NEW QUESTION # 64

A Generative AI Engineer is deciding between using LSH (Locality Sensitive Hashing) and HNSW (Hierarchical Navigable Small World) for indexing their vector database. Their top priority is semantic accuracy. Which approach should the Generative AI Engineer use to evaluate these two techniques?

- A. Compare the cosine similarities of the embeddings of returned results against those of a representative sample of test inputs
- B. Compare the Bilingual Evaluation Understudy (BLEU) scores of returned results for a representative sample of test inputs

- C. Compare the Levenshtein distances of returned results against a representative sample of test inputs
- D. Compare the Recall-Oriented-Understudy for Gisting Evaluation (ROUGE) scores of returned results for a representative sample of test inputs

Answer: A

Explanation:

The task is to choose between LSH and HNSW for a vector database index, prioritizing semantic accuracy.

The evaluation must assess how well each method retrieves semantically relevant results. Let's evaluate the options.

* Option A: Compare the cosine similarities of the embeddings of returned results against those of a representative sample of test inputs

* Cosine similarity measures semantic closeness between vectors, directly assessing retrieval accuracy in a vector database. Comparing returned results' embeddings to test inputs' embeddings evaluates how well LSH or HNSW preserves semantic relationships, aligning with the priority.

* Databricks Reference: "Cosine similarity is a standard metric for evaluating vector search accuracy" ("Databricks Vector Search Documentation," 2023).

* Option B: Compare the Bilingual Evaluation Understudy (BLEU) scores of returned results for a representative sample of test inputs

* BLEU evaluates text generation (e.g., translations), not vector retrieval accuracy. It's irrelevant for indexing performance.

* Databricks Reference: "BLEU applies to generative tasks, not retrieval" ("Generative AI Cookbook").

* Option C: Compare the Recall-Oriented-Understudy for Gisting Evaluation (ROUGE) scores of returned results for a representative sample of test inputs

* ROUGE is for summarization evaluation, not vector search. It doesn't measure semantic accuracy in retrieval.

* Databricks Reference: "ROUGE is unsuited for vector database evaluation" ("Building LLM Applications with Databricks").

* Option D: Compare the Levenshtein distances of returned results against a representative sample of test inputs

* Levenshtein distance measures string edit distance, not semantic similarity in embeddings. It's inappropriate for vector-based retrieval.

* Databricks Reference: No specific support for Levenshtein in vector search contexts.

Conclusion: Option A (cosine similarity) is the correct approach, directly evaluating semantic accuracy in vector retrieval, as recommended by Databricks for Vector Search assessments.

NEW QUESTION # 65

A Generative AI Engineer has a provisioned throughput model serving endpoint as part of a RAG application and would like to monitor the serving endpoint's incoming requests and outgoing responses. The current approach is to include a micro-service in between the endpoint and the user interface to write logs to a remote server.

Which Databricks feature should they use instead which will perform the same task?

- A. Inference Tables
- B. DBSQL
- C. Vector Search
- D. Lakeview

Answer: A

Explanation:

Problem Context: The goal is to monitor the serving endpoint for incoming requests and outgoing responses in a provisioned throughput model serving endpoint within a Retrieval-Augmented Generation (RAG) application. The current approach involves using a microservice to log requests and responses to a remote server, but the Generative AI Engineer is looking for a more streamlined solution within Databricks.

Explanation of Options:

* Option A: Vector Search: This feature is used to perform similarity searches within vector databases.

It doesn't provide functionality for logging or monitoring requests and responses in a serving endpoint, so it's not applicable here.

* Option B: Lakeview: Lakeview is not a feature relevant to monitoring or logging request-response cycles for serving endpoints. It might be more related to viewing data in Databricks Lakehouse but doesn't fulfill the specific monitoring requirement.

* Option C: DBSQL: Databricks SQL (DBSQL) is used for running SQL queries on data stored in Databricks, primarily for analytics purposes. It doesn't provide the direct functionality needed to monitor requests and responses in real-time for an inference endpoint.

* Option D: Inference Tables: This is the correct answer. Inference Tables in Databricks are designed to store the results and metadata of inference runs. This allows the system to log incoming requests and outgoing responses directly within Databricks, making it an ideal choice for monitoring the behavior of a provisioned serving endpoint. Inference Tables can be queried and analyzed,

enabling easier monitoring and debugging compared to a custom microservice.

Thus, Inference Tables are the optimal feature for monitoring request and response logs within the Databricks infrastructure for a model serving endpoint.

NEW QUESTION # 66

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

```
from langchain.chains import LLMChain
from langchain_community.llms import OpenAI
from langchain_core.prompts import PromptTemplate

prompt_template = "Tell me a {adjective} joke"

prompt = PromptTemplate(
    input_variables=["adjective"],
    template=prompt_template
)

llm = LLMChain(prompt=prompt)
llm.generate([{"adjective": "funny"}])
```

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

```
prompt_template = "Tell me a {adjective} joke"
```

```
prompt = PromptTemplate(
    input_variables=["adjective"],
    template=prompt_template
```

```
llm = LLMChain(llm=OpenAI(), prompt=prompt)
```

- A. llm.generate([{"adjective": "funny"}])

```
prompt_template = "Tell me a {adjective} joke"

prompt = PromptTemplate(
    input_variables=["adjective"],
    template=prompt_template
    llm=OpenAI()
)
```

```
llm = LLMChain(prompt=prompt)
llm.generate([{"adjective": "funny"}])
```

- B. llm.generate([{"adjective": "funny"}])

```
prompt_template = "Tell me a {adjective} joke"

prompt = PromptTemplate(
    input_variables=["adjective"],
    template=prompt_template
```

```
llm = LLMChain(prompt=prompt.format("funny"))
```

- C. llm.generate()

```
prompt_template = "Tell me a {adjective} joke"

prompt = PromptTemplate(
    input_variables=["adjective"],
    template=prompt_template
)
```

```
llm = LLMChain(prompt=prompt)
llm.generate("funny")
```

- D.

Answer: B

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 # 67

.....

It is our mission to help you pass the exam. Databricks-Generative-AI-Engineer-Associate guide torrent will provide you with 100% assurance of passing the professional qualification exam. We are very confident in the quality of Databricks-Generative-AI-Engineer-Associate study guide. And we believe that all students who have purchased our study materials will be able to successfully pass the professional qualification exam as long as they follow the content provided by Databricks-Generative-AI-Engineer-Associate study guide, study it on a daily basis, and conduct regular self-examination through mock exams. Once you unfortunately fail the exam, Databricks-Generative-AI-Engineer-Associate Guide Torrent will provide you with a full refund and the refund process is very simple. As long as you provide your staff with your transcripts, you will receive a refund soon. Of course, before you buy, Databricks-Generative-AI-Engineer-Associate certification training offers you a free trial service, as long as you log on our website, you can download our trial questions bank for free. I believe that after you try Databricks-Generative-AI-Engineer-Associate certification training, you will love them.

Study Guide Databricks-Generative-AI-Engineer-Associate Pdf: https://www.itbraindumps.com/Databricks-Generative-AI-Engineer-Associate_exam.html

Databricks Valid Databricks-Generative-AI-Engineer-Associate Exam Syllabus The procedure is very easy and time-saving, If you have some other questions about Databricks Databricks-Generative-AI-Engineer-Associate dumps torrent, ask for our customer service staff, they will contact you 24/7 online for you soon, so you can place your order assured and trusted, In addition, except Databricks-Generative-AI-Engineer-Associate, many other certification exams are also useful, So you can choose the version of Databricks-Generative-AI-Engineer-Associate training guide according to your interests and habits.

When presented with a design challenge, most of us want Databricks-Generative-AI-Engineer-Associate to roll up our sleeves and start right in with the fun part: coming up with ideas to bring the project to life.

We provide 24/7 (24 hours 7 days) online customers service, Databricks-Generative-AI-Engineer-Associate Test Dumps Pdf The procedure is very easy and time-saving, If you have some other questions about Databricks Databricks-Generative-AI-Engineer-Associate Dumps Torrent, ask for our customer service Databricks-Generative-AI-Engineer-Associate Test Dumps Pdf staff, they will contact you 24/7 online for you soon, so you can place your order assured and trusted.

100% Pass Quiz Databricks - Databricks-Generative-AI-Engineer-Associate Fantastic Valid Exam Syllabus

In addition, except Databricks-Generative-AI-Engineer-Associate, many other certification exams are also useful, So you can choose the version of Databricks-Generative-AI-Engineer-Associate training guide according to your interests and habits.

Your personal data and website is always being Valid Databricks-Generative-AI-Engineer-Associate Exam Syllabus monitored by our team, we never disclose your personal information with third party.

BTW, DOWNLOAD part of Itbraindump Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage: <https://drive.google.com/open?id=1ZKCNx4m8tD4KcJTOCiYUjoMMcPmG84k3>