

Pass-Sure新版Databricks-Generative-AI-Engineer-Associate題庫和資格考試中的領先供應商和奇妙的Databricks Certified Generative AI Engineer Associate



Databricks Databricks-Generative-AI-Engineer-Associate認證考試在競爭激烈的IT行業中越來越受歡迎，報名參加考試的人越來越多。但是它的難度並沒有減小，依然很難通過考試，畢竟這是個權威的檢驗電腦專業知識和資訊技術能力的考試。一般人為了通過Databricks Databricks-Generative-AI-Engineer-Associate 認證考試都需要花費大量的時間和精力來復習備考。

彰顯一個人在某一領域是否成功往往體現在他所獲得的資格證書上，在IT行業也不外如是。所以現在很多人都選擇參加Databricks-Generative-AI-Engineer-Associate資格認證考試來證明自己的實力。但是要想通過Databricks-Generative-AI-Engineer-Associate資格認證卻不是一件簡單的事。不過只要你找對了捷徑，通過考試也就變得容易許多了。這就不得不推薦Fast2test的考試考古題了，它可以讓你少走許多彎路，節省時間幫助你考試合格。

>> 新版Databricks-Generative-AI-Engineer-Associate題庫 <<

最有效的新版Databricks-Generative-AI-Engineer-Associate題庫，免費下載Databricks-Generative-AI-Engineer-Associate考試指南得到妳想要的Databricks證書

拿高薪，是每個人的夢想，但究竟能拿多少錢，得由你的職場身價決定。機會很多時候就在你面前.不管你是否喜歡這樣的機會，只有把握住，迎難而上才能獲得非凡的成就。通過 Databricks 認證考試取得一張“金牌派司”無疑是證明和提升自己身價的一個有效方式。Fast2test Databricks-Generative-AI-Engineer-Associate 題庫覆蓋了真實的Databricks Databricks-Generative-AI-Engineer-Associate 考試指南，適合全球考生適用。

最新的 Generative AI Engineer Databricks-Generative-AI-Engineer-Associate 免費考試真題 (Q49-Q54):

問題 #49

A Generative AI Engineer is building an LLM to generate article summaries in the form of a type of poem, such as a haiku, given the article content. However, the initial output from the LLM does not match the desired tone or style. Which approach will NOT improve the LLM's response to achieve the desired response?

- A. Include few-shot examples in the prompt to the LLM
- B. Provide the LLM with a prompt that explicitly instructs it to generate text in the desired tone and style

- C. Use a neutralizer to normalize the tone and style of the underlying documents
- D. Fine-tune the LLM on a dataset of desired tone and style

答案: C

解題說明:

The task at hand is to improve the LLM's ability to generate poem-like article summaries with the desired tone and style. Using a neutralizer to normalize the tone and style of the underlying documents (option B) will not help improve the LLM's ability to generate the desired poetic style. Here's why:

* Neutralizing Underlying Documents: A neutralizer aims to reduce or standardize the tone of input data. However, this contradicts the goal, which is to generate text with a specific tone and style (like haikus). Neutralizing the source documents will strip away the richness of the content, making it harder for the LLM to generate creative, stylistic outputs like poems.

* Why Other Options Improve Results:

* A (Explicit Instructions in the Prompt): Directly instructing the LLM to generate text in a specific tone and style helps align the output with the desired format (e.g., haikus). This is a common and effective technique in prompt engineering.

* C (Few-shot Examples): Providing examples of the desired output format helps the LLM understand the expected tone and structure, making it easier to generate similar outputs.

* D (Fine-tuning the LLM): Fine-tuning the model on a dataset that contains examples of the desired tone and style is a powerful way to improve the model's ability to generate outputs that match the target format.

Therefore, using a neutralizer (option B) is not an effective method for achieving the goal of generating stylized poetic summaries.

問題 #50

A Generative AI Engineer is working with a retail company that wants to enhance its customer experience by automatically handling common customer inquiries. They are working on an LLM-powered AI solution that should improve response times while maintaining a personalized interaction. They want to define the appropriate input and LLM task to do this.

Which input/output pair will do this?

- A. Input: Customer service chat logs; Output Group the chat logs by users, followed by summarizing each user's interactions, then respond
- B. Input: Customer service chat logs; Output: Find the answers to similar questions and respond with a summary
- C. Input: Customer reviews; Output Classify review sentiment
- D. Input: Customer reviews; Output Group the reviews by users and aggregate per-user average rating, then respond

答案: B

解題說明:

The task described in the question involves enhancing customer experience by automatically handling common customer inquiries using an LLM-powered AI solution. This requires the system to process input data (customer inquiries) and generate personalized, relevant responses efficiently. Let's evaluate the options step-by-step in the context of Databricks Generative AI Engineer principles, which emphasize leveraging LLMs for tasks like question answering, summarization, and retrieval-augmented generation (RAG).

* Option A: Input: Customer reviews; Output: Group the reviews by users and aggregate per-user average rating, then respond

* This option focuses on analyzing customer reviews to compute average ratings per user. While this might be useful for sentiment analysis or user profiling, it does not directly address the goal of handling common customer inquiries or improving response times for personalized interactions. Customer reviews are typically feedback data, not real-time inquiries requiring immediate responses.

* Databricks Reference: Databricks documentation on LLMs (e.g., "Building LLM Applications with Databricks") emphasizes that LLMs excel at tasks like question answering and conversational responses, not just aggregation or statistical analysis of reviews.

* Option B: Input: Customer service chat logs; Output: Group the chat logs by users, followed by summarizing each user's interactions, then respond

* This option uses chat logs as input, which aligns with customer service scenarios. However, the output-grouping by users and summarizing interactions-focusses on user-specific summaries rather than directly addressing inquiries. While summarization is an LLM capability, this approach lacks the specificity of finding answers to common questions, which is central to the problem.

* Databricks Reference: Per Databricks' "Generative AI Cookbook," LLMs can summarize text, but for customer service, the emphasis is on retrieval and response generation (e.g., RAG workflows) rather than user interaction summaries alone.

* Option C: Input: Customer service chat logs; Output: Find the answers to similar questions and respond with a summary

* This option uses chat logs (real customer inquiries) as input and tasks the LLM with identifying answers to similar questions, then providing a summarized response. This directly aligns with the goal of handling common inquiries efficiently while maintaining personalization (by referencing past interactions or similar cases). It leverages LLM capabilities like semantic search, retrieval, and response generation, which are core to Databricks' LLM workflows.

* Databricks Reference: From Databricks documentation ("Building LLM-Powered Applications," 2023), an exact extract states: "For customer support use cases, LLMs can be used to retrieve relevant answers from historical data like chat logs and

generate concise, contextually appropriate responses."This matches Option C's approach of finding answers and summarizing them.

* Option D: Input: Customer reviews; Output: Classify review sentiment

* This option focuses on sentiment classification of reviews, which is a valid LLM task but unrelated to handling customer inquiries or improving response times in a conversational context.

It's more suited for feedback analysis than real-time customer service.

* Databricks Reference: Databricks' "Generative AI Engineer Guide" notes that sentiment analysis is a common LLM task, but it's not highlighted for real-time conversational applications like customer support.

Conclusion: Option C is the best fit because it uses relevant input (chat logs) and defines an LLM task (finding answers and summarizing) that meets the requirements of improving response times and maintaining personalized interaction. This aligns with Databricks' recommended practices for LLM-powered customer service solutions, such as retrieval-augmented generation (RAG) workflows.

問題 #51

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. 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"}.
- 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. Respond with "Out of Stock" if the customer asks for a product.
- D. Respond with "In Stock" if the customer asks for a product.

答案: A

解題說明:

* 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.

問題 #52

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

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

答案: B

解題說明:

* 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.

問題 #53

A Generative AI Engineer is creating an LLM-based application. The documents for its retriever have been chunked to a maximum of 512 tokens each. The Generative AI Engineer knows that cost and latency are more important than quality for this application. They have several context length levels to choose from.

Which will fulfill their need?

- A. context length 512: smallest model is 0.13GB and embedding dimension 384
- B. context length 32768: smallest model is 14GB and embedding dimension 4096
- C. context length 2048: smallest model is 11GB and embedding dimension 2560
- D. context length 514; smallest model is 0.44GB and embedding dimension 768

答案: A

解題說明:

When prioritizing cost and latency over quality in a Large Language Model (LLM)-based application, it is crucial to select a configuration that minimizes both computational resources and latency while still providing reasonable performance. Here's why this is the best choice:

* Context length: The context length of 512 tokens aligns with the chunk size used for the documents (maximum of 512 tokens per chunk). This is sufficient for capturing the needed information and generating responses without unnecessary overhead.

* Smallest model size: The model with a size of 0.13GB is significantly smaller than the other options.

This small footprint ensures faster inference times and lower memory usage, which directly reduces both latency and cost.

* Embedding dimension: While the embedding dimension of 384 is smaller than the other options, it is still adequate for tasks where cost and speed are more important than precision and depth of understanding.

This setup achieves the desired balance between cost-efficiency and reasonable performance in a latency-sensitive, cost-conscious application.

問題 #54

.....

通過Databricks Databricks-Generative-AI-Engineer-Associate的考試是不簡單的，選擇合適的培訓是你成功的第一步，選擇好的資訊來源是你成功的保障，而Fast2test的產品是有很好的資訊來源保障。如果你選擇了Fast2test的產品不僅可以100%保證你通過Databricks Databricks-Generative-AI-Engineer-Associate認證考試，還可以為你提供長達一年的免費更新。

Databricks-Generative-AI-Engineer-Associate最新題庫: <https://tw.fast2test.com/Databricks-Generative-AI-Engineer-Associate-premium-file.html>

如果你想知道你是不是充分準備好了考試，那麼你可以利用軟體版的 Databricks-Generative-AI-Engineer-Associate 考試題庫來測試一下自己的水準，但是如果你選擇了我們的Fast2test，你會覺得拿到Databricks Databricks-Generative-AI-Engineer-Associate認證考試的證書不是那麼難了，免費測試：在您決定購買Databricks-Generative-AI-Engineer-Associate題庫之前，您可以先下載我們為您提供的免費樣品，其中有PDF版本和軟體版本，如需要軟體版本請與我們的客服人員及時索取，Databricks 新版Databricks-Generative-AI-Engineer-Associate題庫再沒有比這個資料更好的工具了，但是，這樣的做法真的對Databricks-Generative-AI-Engineer-Associate問題集練習有好處嗎，本站Databricks-Generative-AI-Engineer-Associate認證題庫學習資料根據Databricks Certified Generative AI Engineer Associate Databricks-Generative-AI-Engineer-Associate考試的變化動態更新，我們會在第壹時間更新Databricks-Generative-AI-Engineer-

Associate題庫學習資料，確保Databricks Certified Generative AI Engineer Associate Databricks-Generative-AI-Engineer-Associate考試學習資料是最新的，助您通過Databricks Certified Generative AI Engineer Associate Databricks-Generative-AI-Engineer-Associate認證考試！Fast2test是到目前為止國內最早期的微軟思科認證專業網站，從2000年建站至今已經10年的時間，一直致力於報道微軟思科認證原創文章是我們一貫保持的特點，Databricks Databricks-Generative-AI-Engineer-Associate 新版題庫那是領導對自己工作能力的認可，是事業飛黃騰達的跳板。

白龍心提了起來，賈執事很容易見到，每月外門例事自然能夠見到，如果你想知道你是不是充分準備好了考試，那麼你可以利用軟體版的 Databricks-Generative-AI-Engineer-Associate 考試題庫來測試一下自己的水準，但是如果你選擇了我們的Fast2test，你會覺得拿到Databricks Databricks-Generative-AI-Engineer-Associate認證考試的證書不是那麼難了。

準備充分的新版Databricks-Generative-AI-Engineer-Associate題庫和資格考試領導者和可信賴的Databricks-Generative-AI-Engineer-Associate最新題庫

免費測試：在您決定購買Databricks-Generative-AI-Engineer-Associate題庫之前，您可以先下載我們為您提供的免費樣品，其中有PDF版本和軟體版本，如需要軟體版本請與我們的客服人員及時索取，再沒有比這個資料更好的工具了，但是，這樣的做法真的對Databricks-Generative-AI-Engineer-Associate問題集練習有好處嗎？

