

Databricks-Generative-AI-Engineer-Associate 시험유료자료 & Databricks-Generative-AI-Engineer-Associate 높은 통과율 인기덤프



참고: KoreaDumps에서 Google Drive로 공유하는 무료 2026 Databricks Databricks-Generative-AI-Engineer-Associate 시험 문제집이 있습니다: <https://drive.google.com/open?id=1HhPNPEd-cxtE2JKnUcxBpOBsRBmrHmbh>

KoreaDumps의 완벽한 Databricks인증 Databricks-Generative-AI-Engineer-Associate덤프는 고객님의Databricks인증 Databricks-Generative-AI-Engineer-Associate시험을 패스하는 지름길입니다. 시간과 돈을 적게 들이는 반면 효과는 십점만점에 십점입니다. KoreaDumps의 Databricks인증 Databricks-Generative-AI-Engineer-Associate덤프를 선택하시면 고객님의께서 원하시는 시험점수를 받아 자격증을 쉽게 취득할 수 있습니다.

Databricks Databricks-Generative-AI-Engineer-Associate 시험요강:

주제	소개
주제 1	<ul style="list-style-type: none"> Design Applications: The topic focuses on designing a prompt that elicits a specifically formatted response. It also focuses on selecting model tasks to accomplish a given business requirement. Lastly, the topic covers chain components for a desired model input and output.
주제 2	<ul style="list-style-type: none"> Data Preparation: Generative AI Engineers covers a chunking strategy for a given document structure and model constraints. The topic also focuses on filter extraneous content in source documents. Lastly, Generative AI Engineers also learn about extracting document content from provided source data and format.
주제 3	<ul style="list-style-type: none"> Assembling and Deploying Applications: In this topic, Generative AI Engineers get knowledge about coding a chain using a pyfunc mode, coding a simple chain using langchain, and coding a simple chain according to requirements. Additionally, the topic focuses on basic elements needed to create a RAG application. Lastly, the topic addresses sub-topics about registering the model to Unity Catalog using MLflow.

주제 4	<ul style="list-style-type: none"> Application Development: In this topic, Generative AI Engineers learn about tools needed to extract data, Langchain similar tools, and assessing responses to identify common issues. Moreover, the topic includes questions about adjusting an LLM's response, LLM guardrails, and the best LLM based on the attributes of the application.
주제 5	<ul style="list-style-type: none"> Governance: Generative AI Engineers who take the exam get knowledge about masking techniques, guardrail techniques, and legal licensing requirements in this topic.

>> Databricks-Generative-AI-Engineer-Associate 시험유료자료 <<

Databricks-Generative-AI-Engineer-Associate 높은 통과율 인기덤프 - Databricks-Generative-AI-Engineer-Associate 시험대비 덤프 최신문제

KoreaDumps 질문 풀은 실제시험 변화의 기반에서 스케줄에 따라 업데이트 합니다. 만일 Databricks Databricks-Generative-AI-Engineer-Associate 테스트에 어떤 변화가 생긴다면, 적중율이 항상 98% 이상을 유지 할 수 있도록 2일간의 근무일 안에 제품을 업데이트 하도록 합니다. KoreaDumps는 고객들이 테스트에 성공적으로 합격 할 수 있도록 하기 위하여 업데이트 된 버전을 구매후 서비스로 제공해드립니다. 시험에서 불합격받으셨는데 업데이트가 힘든 상황이면 덤프비용을 환불해드립니다.

최신 Generative AI Engineer Databricks-Generative-AI-Engineer-Associate 무료샘플문제 (Q23-Q28):

질문 # 23

A Generative AI Engineer is tasked with developing a RAG application that will help a small internal group of experts at their company answer specific questions, augmented by an internal knowledge base. They want the best possible quality in the answers, and neither latency nor throughput is a huge concern given that the user group is small and they're willing to wait for the best answer. The topics are sensitive in nature and the data is highly confidential and so, due to regulatory requirements, none of the information is allowed to be transmitted to third parties.

Which model meets all the Generative AI Engineer's needs in this situation?

- A. OpenAI GPT-4
- B. BGE-large
- C. Dolly 1.5B
- D. Llama2-70B

정답: B

설명:

Problem Context: The Generative AI Engineer needs a model for a Retrieval-Augmented Generation (RAG) application that provides high-quality answers, where latency and throughput are not major concerns. The key factors are confidentiality and sensitivity of the data, as well as the requirement for all processing to be confined to internal resources without external data transmission.

Explanation of Options:

* Option A: Dolly 1.5B: This model does not typically support RAG applications as it's more focused on image generation tasks.

* Option B: OpenAI GPT-4: While GPT-4 is powerful for generating responses, its standard deployment involves cloud-based processing, which could violate the confidentiality requirements due to external data transmission.

* Option C: BGE-large: The BGE (Big Green Engine) large model is a suitable choice if it is configured to operate on-premises or within a secure internal environment that meets regulatory requirements.

Assuming this setup, BGE-large can provide high-quality answers while ensuring that data is not transmitted to third parties, thus aligning with the project's sensitivity and confidentiality needs.

* Option D: Llama2-70B: Similar to GPT-4, unless specifically set up for on-premises use, it generally relies on cloud-based services, which might risk confidential data exposure.

Given the sensitivity and confidentiality concerns, BGE-large is assumed to be configurable for secure internal use, making it the optimal choice for this scenario.

질문 # 24

A Generative AI Engineer is creating an LLM-powered application that will need access to up-to-date news articles and stock prices.

The design requires the use of stock prices which are stored in Delta tables and finding the latest relevant news articles by searching the internet.

How should the Generative AI Engineer architect their LLM system?

- A. Create an agent with tools for SQL querying of Delta tables and web searching, provide retrieved values to an LLM for generation of response.
- B. Use an LLM to summarize the latest news articles and lookup stock tickers from the summaries to find stock prices.
- C. Download and store news articles and stock price information in a vector store. Use a RAG architecture to retrieve and generate at runtime.
- D. Query the Delta table for volatile stock prices and use an LLM to generate a search query to investigate potential causes of the stock volatility.

정답: A

설명:

To build an LLM-powered system that accesses up-to-date news articles and stock prices, the best approach is to create an agent that has access to specific tools (option D).

Agent with SQL and Web Search Capabilities:

By using an agent-based architecture, the LLM can interact with external tools. The agent can query Delta tables (for up-to-date stock prices) via SQL and perform web searches to retrieve the latest news articles. This modular approach ensures the system can access both structured (stock prices) and unstructured (news) data sources dynamically.

Why This Approach Works:

SQL Queries for Stock Prices: Delta tables store stock prices, which the agent can query directly for the latest data.

Web Search for News: For news articles, the agent can generate search queries and retrieve the most relevant and recent articles, then pass them to the LLM for processing.

Why Other Options Are Less Suitable:

A (Summarizing News for Stock Prices): This convoluted approach would not ensure accuracy when retrieving stock prices, which are already structured and stored in Delta tables.

B (Stock Price Volatility Queries): While this could retrieve relevant information, it doesn't address how to obtain the most up-to-date news articles.

C (Vector Store): Storing news articles and stock prices in a vector store might not capture the real-time nature of stock data and news updates, as it relies on pre-existing data rather than dynamic querying.

Thus, using an agent with access to both SQL for querying stock prices and web search for retrieving news articles is the best approach for ensuring up-to-date and accurate responses.

질문 # 25

A Generative AI Engineer is designing a RAG application for answering user questions on technical regulations as they learn a new sport.

What are the steps needed to build this RAG application and deploy it?

- A. Ingest documents from a source -> Index the documents and saves to Vector Search -> User submits queries against an LLM -> LLM retrieves relevant documents -> Evaluate model -> LLM generates a response -> Deploy it using Model Serving
- B. Ingest documents from a source -> Index the documents and save to Vector Search -> User submits queries against an LLM -> LLM retrieves relevant documents -> LLM generates a response -> Evaluate model -> Deploy it using Model Serving
- C. Ingest documents from a source -> Index the documents and save to Vector Search -> Evaluate model -> Deploy it using Model Serving
- D. User submits queries against an LLM -> Ingest documents from a source -> Index the documents and save to Vector Search -> LLM retrieves relevant documents -> LLM generates a response -> Evaluate model -> Deploy it using Model Serving

정답: B

설명:

The Generative AI Engineer needs to follow a methodical pipeline to build and deploy a Retrieval-Augmented Generation (RAG)

application. The steps outlined in option B accurately reflect this process:

Ingest documents from a source: This is the first step, where the engineer collects documents (e.g., technical regulations) that will be used for retrieval when the application answers user questions.

Index the documents and save to Vector Search: Once the documents are ingested, they need to be embedded using a technique like embeddings (e.g., with a pre-trained model like BERT) and stored in a vector database (such as Pinecone or FAISS). This enables fast retrieval based on user queries.

User submits queries against an LLM: Users interact with the application by submitting their queries. These queries will be passed to the LLM.

LLM retrieves relevant documents: The LLM works with the vector store to retrieve the most relevant documents based on their vector representations.

LLM generates a response: Using the retrieved documents, the LLM generates a response that is tailored to the user's question.

Evaluate model: After generating responses, the system must be evaluated to ensure the retrieved documents are relevant and the generated response is accurate. Metrics such as accuracy, relevance, and user satisfaction can be used for evaluation.

Deploy it using Model Serving: Once the RAG pipeline is ready and evaluated, it is deployed using a model-serving platform such as Databricks Model Serving. This enables real-time inference and response generation for users.

By following these steps, the Generative AI Engineer ensures that the RAG application is both efficient and effective for the task of answering technical regulation questions.

질문 # 26

A Generative AI Engineer just deployed an LLM application at a digital marketing company that assists with answering customer service inquiries.

Which metric should they monitor for their customer service LLM application in production?

- A. HuggingFace Leaderboard values for the base LLM
- B. Energy usage per query
- C. Number of customer inquiries processed per unit of time
- D. Final perplexity scores for the training of the model

정답: C

설명:

When deploying an LLM application for customer service inquiries, the primary focus is on measuring the operational efficiency and quality of the responses. Here's why A is the correct metric:

* Number of customer inquiries processed per unit of time: This metric tracks the throughput of the customer service system, reflecting how many customer inquiries the LLM application can handle in a given time period (e.g., per minute or hour). High throughput is crucial in customer service applications where quick response times are essential to user satisfaction and business efficiency.

* Real-time performance monitoring: Monitoring the number of queries processed is an important part of ensuring that the model is performing well under load, especially during peak traffic times. It also helps ensure the system scales properly to meet demand.

Why other options are not ideal:

* B. Energy usage per query: While energy efficiency is a consideration, it is not the primary concern for a customer-facing application where user experience (i.e., fast and accurate responses) is critical.

* C. Final perplexity scores for the training of the model: Perplexity is a metric for model training, but it doesn't reflect the real-time operational performance of an LLM in production.

* D. HuggingFace Leaderboard values for the base LLM: The HuggingFace Leaderboard is more relevant during model selection and benchmarking. However, it is not a direct measure of the model's performance in a specific customer service application in production.

Focusing on throughput (inquiries processed per unit time) ensures that the LLM application is meeting business needs for fast and efficient customer service responses.

질문 # 27

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

정답: A

설명:

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.

질문 # 28

.....

Databricks-Generative-AI-Engineer-Associate인증시험은 IT업계에 종사하고 계신 분이시라면 최근 많은 인기를 누리고 있다는 것을 알고 계실것입니다. Databricks-Generative-AI-Engineer-Associate인증시험을 패스하여 자격증을 취득하는데 가장 쉬운 방법은 KoreaDumps에서 제공해드리는 Databricks-Generative-AI-Engineer-Associate덤프를 공부하는 것입니다. Databricks Databricks-Generative-AI-Engineer-Associate덤프에 있는 문제와 답만 기억하시면 Databricks-Generative-AI-Engineer-Associate시험을 패스하는데 많은 도움이 됩니다.덤프구매후 최신버전으로 업데이트되면 업데이트버전을 시스템 자동으로 구매시 사용한 메일주소로 발송해드려 덤프유효기간을 최대한 길게 연장해드립니다.

Databricks-Generative-AI-Engineer-Associate높은 통과율 인기덤프 : https://www.koreadumps.com/Databricks-Generative-AI-Engineer-Associate_exam-braindumps.html

- Databricks-Generative-AI-Engineer-Associate인증덤프문제 □ Databricks-Generative-AI-Engineer-Associate시험패스 가능한 공부자료 □ Databricks-Generative-AI-Engineer-Associate최신 기출문제 □ ▶ www.exampassdump.com <을(를) 열고 > Databricks-Generative-AI-Engineer-Associate □를 입력하고 무료 다운로드를 받으십시오 Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부
- 시험패스 가능한 Databricks-Generative-AI-Engineer-Associate시험유료자료 덤프데모문제 □ ⇒ www.itdumpskr.com <에서 검색만 하면> Databricks-Generative-AI-Engineer-Associate □<를 무료로 다운로드 할 수 있습니다 Databricks-Generative-AI-Engineer-Associate완벽한 덤프문제자료
- 시험패스에 유효한 Databricks-Generative-AI-Engineer-Associate시험유료자료 최신버전 덤프데모 문제 □ ⇒ kr.fast2test.com □을 통해 쉽게 【 Databricks-Generative-AI-Engineer-Associate 】 무료 다운로드 받기 Databricks-Generative-AI-Engineer-Associate시험패스 가능한 공부자료
- Databricks-Generative-AI-Engineer-Associate시험유료자료 최신버전 인증덤프문제 □ 지금 [www.itdumpskr.com] 에서 “ Databricks-Generative-AI-Engineer-Associate ”를 검색하고 무료로 다운로드 하세요 Databricks-Generative-AI-Engineer-Associate최신 기출문제
- Databricks-Generative-AI-Engineer-Associate최신 업데이트 덤프 □ Databricks-Generative-AI-Engineer-Associate최신 시험 예상문제모음 □ Databricks-Generative-AI-Engineer-Associate시험대비 최신버전 덤프샘플 □ 시험 자료를 무료로 다운로드 하려면 ⇒ www.passtip.net □□□을 통해 【 Databricks-Generative-AI-Engineer-Associate 】 를

검색하십시오Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부

- 퍼펙트한 Databricks-Generative-AI-Engineer-Associate시험유료자료 덤프데모문제 보기 □ { www.itdumpskr.com } 웹사이트에서 { Databricks-Generative-AI-Engineer-Associate }를 열고 검색하여 무료 다운로드Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부
- 시험준비에 가장 좋은 Databricks-Generative-AI-Engineer-Associate시험유료자료 덤프데모문제 다운 □ ✓ www.passtip.net □ ✓ □의 무료 다운로드 ⇒ Databricks-Generative-AI-Engineer-Associate ◀페이지가 지금 열립니다 Databricks-Generative-AI-Engineer-Associate시험대비 덤프데모문제
- Databricks-Generative-AI-Engineer-Associate시험유료자료 최신 시험대비 덤프공부자료 □ ⇒ www.itdumpskr.com □ □ □에서 ⇒ Databricks-Generative-AI-Engineer-Associate □ □ □를 검색하고 무료로 다운로드 하세요Databricks-Generative-AI-Engineer-Associate높은 통과율 시험공부
- Databricks-Generative-AI-Engineer-Associate시험대비 덤프데모문제 □ Databricks-Generative-AI-Engineer-Associate시험패스 가능한 공부자료 □ Databricks-Generative-AI-Engineer-Associate완벽한 덤프자료 □ ⇒ www.koreadumps.com □ □ □ 웹사이트를 열고 ⇒ Databricks-Generative-AI-Engineer-Associate □를 검색하여 무료 다운로드Databricks-Generative-AI-Engineer-Associate최신시험후기
- Databricks-Generative-AI-Engineer-Associate완벽한 덤프자료 □ Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부 □ Databricks-Generative-AI-Engineer-Associate최신시험후기 □ [www.itdumpskr.com] 을(를) 열고 【 Databricks-Generative-AI-Engineer-Associate 】를 검색하여 시험 자료를 무료로 다운로드하십시오Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부
- 시험패스 가능한 Databricks-Generative-AI-Engineer-Associate시험유료자료 덤프데모문제 □ 무료로 쉽게 다운로드 하려면“ kr.fast2test.com ”에서 ⇒ Databricks-Generative-AI-Engineer-Associate □를 검색하세요Databricks-Generative-AI-Engineer-Associate최신 업데이트버전 덤프공부
- bookmarkcitizen.com, lucxlnb380313.azzablog.com, gretafvbf214390.blogdeazar.com, joyceuxdu383878.qodsblog.com, umairvwem721318.westxwiki.com, safiyalrxb224116.liberty-blog.com, aoiferbv933238.fliplife-wiki.com, albieddlo137654.blogitright.com, my-social-box.com, bookmarkinglife.com, Disposable vapes

2026 KoreaDumps 최신 Databricks-Generative-AI-Engineer-Associate PDF 버전 시험 문제집과 Databricks-Generative-AI-Engineer-Associate 시험 문제 및 답변 무료 공유: <https://drive.google.com/open?id=1HhPNPEd-cxtE2JKnUcxBpOBsRBmrHmbh>