

DSA-C03 Practice Materials & DSA-C03 Best Questions & DSA-C03 Exam Guide



P.S. Free 2026 Snowflake DSA-C03 dumps are available on Google Drive shared by DumpsTorrent:
https://drive.google.com/open?id=1DkKl_N2W0gs5bqeVxkFtU73libjw2Egb

Our web-based practice test is accessible from anywhere with an internet connection, which means you can take it at your convenience. This Snowflake DSA-C03 Practice Test is designed to simulate the actual exam and help you become familiar with the test format. You can access the web-based practice exam from anywhere with an internet connection to study on the go or from the comfort of your own home. You can receive your mock exam result instantly.

There are a lot of materials for Snowflake DSA-C03 practice test. DumpsTorrent is the only site providing with the finest Snowflake DSA-C03 dumps torrent. All DumpsTorrent test questions are the latest and we guarantee you can pass your exam at first time. DSA-C03 Questions and answers DumpsTorrent provide are rewritten by the modern information technology experts, which is good for you.

>> **DSA-C03 Test Engine** <<

Latest DSA-C03 Test Fee & DSA-C03 Instant Download

If you want to clear the exam for Snowflake DSA-C03 certification along with your job, there is no need to worry about it. You can choose flexible timings for the learning session and get all the SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam questions online and practice with Snowflake DSA-C03 exam dumps any time you want. There is no strict schedule for it.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q283-Q288):

NEW QUESTION # 283

Which of the following statements about Z-tests and T-tests are generally true? Select all that apply.

- A. Both Z-tests and T-tests assume that the data is non-normally distributed.
- B. A T-test has fewer degrees of freedom compared to the Z-test, making it more robust to outliers.

- C. A Z-test requires knowing the population standard deviation, while a T-test estimates it from the sample data.
- D. As the sample size increases, the T-distribution approaches the standard normal (Z) distribution.
- E. A T-test is generally used when the sample size is large ($n > 30$) and the population standard deviation is known.

Answer: C,D

Explanation:

The correct answers are A and C. A Z-test requires knowing the population standard deviation, while a T-test estimates it from the sample data. As the sample size increases, the T-distribution approaches the standard normal (Z) distribution, which is a core concept in statistical inference. B is incorrect because a T-test is generally used for small sample sizes ($n < 30$) or when the population standard deviation is unknown. D is incorrect because both tests assume the underlying population distribution is approximately normal, especially for smaller sample sizes (though the Central Limit Theorem allows us to relax this assumption somewhat for large samples). E is incorrect because fewer degrees of freedom make the t-test less robust to outliers. Also the robustness is provided by the population distribution being approximately normal.

NEW QUESTION # 284

You are building a customer support chatbot using Snowflake Cortex and a large language model (LLM). You want to use prompt engineering to improve the chatbot's ability to answer complex questions about product features. You have a table PRODUCT DETAILS with columns 'feature_name'. Which of the following prompts, when used with the COMPLETE function in Snowflake Cortex, is MOST likely to yield the best results for answering user questions about specific product features, assuming you are aiming for concise and accurate responses focused solely on providing the requested feature description and avoiding extraneous chatbot-like conversation?

- A. Option B
- B. Option E
- C. Option A
- D. Option D
- E. Option C

Answer: E

Explanation:

Option C is the best prompt because it directly instructs the LLM to act as a product expert and provide only the feature description, minimizing extraneous conversation or information. Options A and B lack specific instructions, potentially leading to verbose responses. Option D includes all product details in the prompt, which might overwhelm the LLM. Option E tries to fetch a specific feature description, but the SQL is incorrect. Correct SQL will increase token usage and may not lead to a concise response.

NEW QUESTION # 285

You are tasked with predicting sales (SALES AMOUNT) for a retail company using linear regression in Snowflake. The dataset includes features like 'ADVERTISING SPEND', 'PROMOTIONS', 'SEASONALITY INDEX', and 'COMPETITOR PRICE'. After training a linear regression model named 'sales model', you observe that the model performs poorly on new data, indicating potential issues with multicollinearity or overfitting. Which of the following strategies, applied directly within Snowflake, would be MOST effective in addressing these issues and improving the model's generalization performance? Choose ALL that apply.

- A. Decrease the 'MAX_ITERATIONS' parameter in the 'CREATE MODEL' statement to prevent the model from overfitting to the training data.
- B. Increase the size of the training dataset significantly by querying data from external sources.
- C. Apply Ridge Regression by adding an L2 regularization term during model training. This can be achieved by setting the 'REGULARIZATION' parameter of the 'CREATE MODEL' statement to 'L2'.
- D. Perform feature scaling (e.g., standardization or min-max scaling) on the input features before training the model, using Snowflake's built-in functions or user-defined functions (UDFs) for scaling.
- E. Manually remove highly correlated features (e.g., if 'ADVERTISING SPEND' and 'PROMOTIONS' have a correlation coefficient above 0.8) based on a correlation matrix calculated using 'CORR' function and feature selection techniques.

Answer: C,D,E

Explanation:

Options A, B, and D are the most effective strategies for addressing multicollinearity and overfitting in this scenario. Ridge

Regression (A) adds an L2 regularization term, which penalizes large coefficients and reduces overfitting. Manually removing highly correlated features (B) addresses multicollinearity directly. Performing feature scaling (D) ensures that features with different scales do not disproportionately influence the model. Increasing training data (C) is generally helpful, but doesn't directly solve multicollinearity. Decreasing MAX ITERATIONS (E) might prevent the model from fully converging, but is a less targeted approach than regularization or feature selection.

NEW QUESTION # 286

You've trained a machine learning model using Scikit-learn and saved it as 'model.joblib'. You need to deploy this model to Snowflake. Which sequence of commands will correctly stage the model and create a Snowflake external function to use it for inference, assuming you already have a Snowflake stage named 'model_stage'?

- A. Option B
- B. Option C
- **C. Option E**
- D. Option A
- E. Option D

Answer: C

Explanation:

NEW QUESTION # 287

You are developing a regression model in Snowflake to predict housing prices. You've trained a model using Snowflake ML functions and now need to rigorously validate its performance. You have a separate validation dataset stored in a table named 'HOUSING VALIDATION'. Which of the following SQL statements, when executed in Snowflake, would accurately calculate the Root Mean Squared Error (RMSE) of your model's predictions against the actual prices in the validation dataset, assuming your model is named 'HOUSING PRICE MODEL' and the prediction function generated by CREATE SNOWFLAKE.ML.FORECAST is called PREDICT?

- A. Option B
- B. Option C
- **C. Option E**
- D. Option A
- E. Option D

Answer: C

Explanation:

Option E is the correct answer because it correctly calculates the RMSE using the Snowflake ML PREDICT function in conjunction with the POWER and AVG functions within a SQL query. It constructs an object for input to PREDICT, excluding the actual price to prevent data leakage. Options A, B, and C have syntax errors or incorrect function usage for calculating RMSE in Snowflake and assume a PREDICT function that is generated by CREATE SNOWFLAKE.ML.FORECAST, they don't use SNOWFLAKE.ML.PREDICT directly. Option D assumes a function named ROOT MEAN SQUARED ERROR which is not a native Snowflake function.

NEW QUESTION # 288

.....

We believe that if you can learn about several advantages of DSA-C03 preparation questions, I believe you have more understanding of the real questions and answers. You can download the trial versions of the DSA-C03 Exam Questions for free. After using the trial version of our DSA-C03 study materials, I believe you will have a deeper understanding of the advantages of our DSA-C03 training engine.

Latest DSA-C03 Test Fee: <https://www.dumpstorrent.com/DSA-C03-exam-dumps-torrent.html>

The DSA-C03 exam bootcamp is quite necessary for the passing of the exam, Snowflake DSA-C03 Test Engine This kind of access could cost you over \$100,000 if you were to purchase all materials individually, Snowflake DSA-C03 Test Engine The

