

Free PDF Quiz DSA-C03 - The Best New SnowPro Advanced: Data Scientist Certification Exam Test Braindumps



BONUS!!! Download part of BraindumpsIT DSA-C03 dumps for free: <https://drive.google.com/open?id=1KqiQAxpyHYTe2mSGmp0H29zY37y-BvXw>

The clients can consult our online customer service before and after they buy our DSA-C03 useful test guide. We provide considerate customer service to the clients. Before the clients buy our DSA-C03 cram training materials they can consult our online customer service personnel about the products' version and price and then decide whether to buy them or not. After the clients buy the DSA-C03 Study Tool they can consult our online customer service about how to use them and the problems which occur during the process of using. We will help you pass the DSA-C03 exam in the shortest time.

The Snowflake DSA-C03 PDF dumps file is the most convenient way to prepare for the examination. This document is a collection of most probable and realistic SnowPro Advanced: Data Scientist Certification Exam DSA-C03 dumps. With this PDF file, you have SnowPro Advanced: Data Scientist Certification Exam DSA-C03 questions that will appear in the real exam. You can immediately download our DSA-C03 PDF Questions from the BraindumpsIT website after payment. Without place and time limits, you can use the PDF format of SnowPro Advanced: Data Scientist Certification Exam DSA-C03 real exam questions via smartphones, tablets, and laptops.

>> [New DSA-C03 Test Braindumps](#) <<

Customized DSA-C03 Lab Simulation - DSA-C03 Pdf Files

Our qualified team of Snowflake SnowPro Advanced: Data Scientist Certification Exam study material to improve the quality and to match the changes in the syllabus and pattern shared by DSA-C03. Our desktop Snowflake DSA-C03 Practice Exam software is designed for all those candidates who want to learn and practice in the actual Snowflake DSA-C03 exam environment.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q21-Q26):

NEW QUESTION # 21

You have a structured dataset in Snowflake containing customer information and purchase history. You aim to build a multi-class classification model to predict customer churn, categorizing customers into 'Low Risk', 'Medium Risk', and 'High Risk' of churning. After training the model, you want to evaluate its performance. Which of the following metrics and evaluation techniques, when used together, provide the MOST comprehensive understanding of the model's performance across all churn risk categories, especially when dealing with potential class imbalance?

- A. Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R-squared (Coefficient of Determination).
- B. Overall Accuracy, Precision, Recall, F1-Score for each class, and Confusion Matrix.
- C. Only Overall Accuracy and a confusion Matrix.
- D. Log Loss (Cross-Entropy Loss), Gini Coefficient, and Kolmogorov-Smirnov (KS) statistic.
- E. Area Under the ROC Curve (AUC-ROC) for each class (one-vs-rest approach), Precision-Recall Curve for each class, and Cumulative Accuracy Profile (CAP) curve.

Answer: B

Explanation:

Option A offers the most comprehensive evaluation. Overall accuracy provides a general sense of performance, but can be misleading with imbalanced classes. Precision, recall, and F1-score, calculated for each class, give a detailed view of the model's performance on each churn risk category. The confusion matrix provides a visual representation of the model's classification errors, allowing you to identify patterns of misclassification between the different risk levels. Option B, ROC AUC and Precision-Recall curve are also relevant but is better for binary classification (with one-vs-rest extended for multiclass). CAP curves are less common. Option C (Log Loss, Gini, KS) is more suitable for binary classification or ranking problems. Option D (RMSE, MAE, R-squared) are regression metrics, not suitable for classification.

NEW QUESTION # 22

You are a data scientist working for a retail company using Snowflake. You're building a linear regression model to predict sales based on advertising spend across various channels (TV, Radio, Newspaper). After initial EDA, you suspect multicollinearity among the independent variables. Which of the following Snowflake SQL statements or techniques are MOST appropriate for identifying and addressing multicollinearity BEFORE fitting the model? Choose two.

- A. Calculate the Variance Inflation Factor (VIF) for each independent variable using a user-defined function (UDF) in Snowflake that implements the VIF calculation based on R-squared values from auxiliary regressions. This requires fitting a linear regression for each independent variable against all others.
- B. Drop one of the independent variable randomly if they seem highly correlated.
- C. Use t on each independent variable to estimate its uniqueness. If uniqueness is low, multicollinearity is likely.
- D. Generate a correlation matrix of the independent variables using 'CORR aggregate function in Snowflake SQL and examine the correlation coefficients. Values close to +1 or -1 suggest high multicollinearity.
- E. Implement Principal Component Analysis (PCA) using Snowpark Python to transform the independent variables into uncorrelated principal components and then select only the components explaining a certain percentage of the variance.

Answer: A,D

Explanation:

Multicollinearity can be identified by calculating the VIF for each independent variable. VIF is calculated by regressing each independent variable against all other independent variables and calculating $1/(1-RA^2)$, where RA^2 is the R-squared value from the regression. A high VIF suggests high multicollinearity. Correlation matrices generated with 'CORR' can also reveal multicollinearity by showing pairwise correlations between independent variables. PCA using Snowpark is also a viable option, but less direct than VIF and correlation matrix analysis for identifying multicollinearity. APPROX_COUNT_DISTINCT is not directly related to identifying multicollinearity. Randomly dropping variables will also lead to data loss.

NEW QUESTION # 23

You are tasked with performing exploratory data analysis on a table named containing daily sales transactions. The table includes columns like 'transaction_date', 'product_id', 'quantity', and 'price'. Your goal is to identify potential data quality issues and understand the distribution of sales. Which of the following SQL queries using Snowflake's statistical functions and features would be MOST effective for quickly identifying outliers in the 'quantity' column, potential data skewness, and missing values?

<input type="checkbox"/> SELECT AVG(quantity), STDDEV(quantity) FROM sales_data;
<input type="checkbox"/> SELECT COUNT() AS total_records, COUNT(quantity) AS non_null_quantity, APPROX_MEDIAN(quantity), APPROX_COUNT_DISTINCT(product_id) FROM sales_data;
<input type="checkbox"/> SELECT MIN(quantity), MAX(quantity), VARIANCE(quantity), SKEW(quantity) FROM sales_data;
<input type="checkbox"/> SELECT CORR(quantity, price) FROM sales_data;
<input type="checkbox"/> SELECT COUNT() as record_count FROM sales_data;

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

Answer: A,E

Explanation:

Options B and C are the most effective. Option B provides the total record count, the count of non-null quantity values (helping identify missing data), an approximate median, and the approximate distinct count of product IDs. This gives a good overview of data completeness and product diversity. Option C provides the min, max, variance and skew which together help identify possible outliers and understand the data distribution.

NEW QUESTION # 24

A financial institution suspects fraudulent activity based on unusual transaction patterns. They want to use association rule mining to identify relationships between different transaction attributes (e.g., transaction amount, location, time of day, merchant category code) that are indicative of fraud. The data is stored in a Snowflake table called 'TRANSACTIONS'. Which of the following considerations are CRITICAL when applying association rule mining in this fraud detection scenario?

- A. Prioritize rules with high confidence and lift, even if support is relatively low, as rare but highly predictive combinations of attributes can be strong indicators of fraudulent activity.
- B. Ensure that the Apriori algorithm is run directly within Snowflake using SQL to maximize performance and scalability, rather than extracting the data and processing it in an external Python environment.
- C. Carefully discretize continuous variables like 'transaction amount' and 'time of day' into meaningful categories to enable association rule mining, and consider the impact of different discretization strategies on the resulting rules.
- D. Ignore transaction attributes that have a large number of distinct values (e.g., specific location coordinates) as they will likely lead to an explosion of rules and make interpretation difficult.
- E. Focus solely on rules with very high support (e.g., > 0.1) to ensure statistical significance and avoid overfitting to rare fraudulent events.

Answer: A,C

Explanation:

Option B is critical because discretization is essential for handling continuous variables in association rule mining. The way these variables are binned can significantly influence the rules discovered. Option C is also critical because in fraud detection, identifying rare but highly predictive rules is crucial. Low support rules, if they have high confidence and lift, can point to specific patterns indicative of fraud. Option A is incorrect because requiring high support would miss rare fraud patterns. Option D is incorrect because some high cardinality attributes might be important indicators. Option E is incorrect as Apriori algorithm cannot be directly run using SQL, Snowpark and python is a good option.

NEW QUESTION # 25

You are working with a large dataset in Snowflake and need to build a machine learning model using scikit-learn in Python. You want to leverage Snowflake's compute resources for feature engineering to speed up the process. Which of the following approaches correctly combines Snowflake's SQL capabilities with scikit-learn for feature engineering and model training, while minimizing data transfer between Snowflake and the Python environment?

- A. Create Snowflake User-Defined Functions (UDFs) in Python for complex feature engineering calculations. Call these UDFs within a SQL query to apply the feature engineering to the Snowflake data. Load the resulting features into a Pandas DataFrame and train the scikit-learn model.
- B. Use the Snowflake Python Connector to execute individual SQL queries for each feature engineering step. Load the resulting features step-by-step into a Pandas DataFrame and train the scikit-learn model.

- C. Write a complex SQL query in Snowmake to perform all feature engineering, then load the resulting features into a Pandas DataFrame and train the scikit-learn model.
- D. Implement the feature engineering steps directly in Python using Pandas and scikit-learn, then load the raw data into a Pandas DataFrame and apply the transformations. Finally, train the scikit-learn model.
- E. Use Snowflake external functions to invoke a remote service (e.g., AWS Lambda) for feature engineering. Pass data from Snowflake to the remote service, receive the engineered features back, and load them into a Pandas DataFrame for model training.

Answer: A

Explanation:

Option D is the most efficient approach. Using Snowflake UDFs in Python allows you to perform complex feature engineering directly within Snowflake's compute environment, minimizing the amount of data that needs to be transferred to the Python environment. This reduces network latency and improves performance. Option A may be workable but it would need writing complex SQL queries. Option B will involve a lot of individual interactions between Snowflake and python making this a slower and more complex process. Option C would bring the data out to python before processing it with Pandas and scikit-learn, meaning you'd lose out on the compute of Snowflake. Option E is a viable solution to offload compute to a different compute environment than the python environment and load into a Pandas DataFrame.

NEW QUESTION # 26

.....

As long as you face problems with the DSA-C03 exam, our company is confident to help you solve. Give our DSA-C03 practice quiz a choice is to give you a chance to succeed. We are very willing to go hand in hand with you on the way to preparing for DSA-C03 Exam. And we have three different versions of our DSA-C03 learning materials, you will find that it is so interesting and funny to study with our study guide.

Customized DSA-C03 Lab Simulation: https://www.braindumpsit.com/DSA-C03_real-exam.html

For example, in order to make every customer can purchase at ease, our DSA-C03 preparation quiz will provide users with three different versions for free trial, corresponding to the three official versions, One-year free update right will enable you get the latest DSA-C03 vce dumps anytime and you just need to check your mailbox, BraindumpsIT Customized DSA-C03 Lab Simulation is most reliable platform for your ultimate success, we are offering services for last 10 years and have gathered almost 70,000+ satisfied customer around the world.

Understand how to... Integrate Windows Vista into an Active DSA-C03 Directory domain, and the Government, a.k.a, For example, in order to make every customer can purchase at ease, our DSA-C03 Preparation quiz will provide users with three different versions for free trial, corresponding to the three official versions.

Magnificent DSA-C03 Preparation Dumps: SnowPro Advanced: Data Scientist Certification Exam Represent the Most Popular Simulating Exam - BraindumpsIT

One-year free update right will enable you get the latest DSA-C03 vce dumps anytime and you just need to check your mailbox, BraindumpsIT is most reliable platform for your ultimate success, we are offering Reliable DSA-C03 Exam Sample services for last 10 years and have gathered almost 70,000+ satisfied customer around the world.

To pass the SnowPro Advanced: Data Scientist Certification Exam test in a short time, you must prepare with DSA-C03 exam questions that are real and updated, Of course, we will humbly accept your opinions on our DSA-C03 quiz guide.

- Hot New DSA-C03 Test Braindumps 100% Pass | Valid Customized DSA-C03 Lab Simulation: SnowPro Advanced: Data Scientist Certification Exam □ ➔ www.practicevce.com □ is best website to obtain ★ DSA-C03 ★ for free download □ Practice DSA-C03 Test Online
- 100% Pass Quiz Snowflake - DSA-C03 Useful New Test Braindumps □ Enter ➔ www.pdfvce.com □□□ and search for 【 DSA-C03 】 to download for free □ DSA-C03 Study Guide Pdf
- Reliable DSA-C03 Practice Materials □ Valid Braindumps DSA-C03 Questions □ DSA-C03 Valid Braindumps Ebook □ Easily obtain free download of 【 DSA-C03 】 by searching on [www.prep4away.com] □ Test DSA-C03 Guide Online
- 100% Pass 2026 DSA-C03: Efficient New SnowPro Advanced: Data Scientist Certification Exam Test Braindumps □ Enter ✓ www.pdfvce.com □✓□ and search for □ DSA-C03 □ to download for free □ DSA-C03 Valid Braindumps

Ebook

2026 Latest BraindumpsIT DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: <https://drive.google.com/open?id=1KqiQAxpyHYTe2mSGmp0H29zY37y-BvXw>