

Valid Exam Snowflake DSA-C03 Preparation, Latest DSA-C03 Learning Material



2026 Latest ExamsLabs DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: https://drive.google.com/open?id=17kD6agegdVx4KqkPkQUWJ_f-X62tgmbR

Our DSA-C03 study guide boosts both the high passing rate which is about 98%-100% and the high hit rate to have few difficulties to pass the test. Our DSA-C03 exam simulation is compiled based on the resources from the authorized experts' diligent working and the Real DSA-C03 Exam and confer to the past years' exam papers thus they are very practical. The content of the questions and answers of DSA-C03 exam quiz is refined and focuses on the most important information.

You can free download part of ExamsLabs's practice questions and answers about Snowflake certification DSA-C03 exam online, as an attempt to test our quality. As long as you choose to purchase ExamsLabs's products, we will do our best to help you pass Snowflake Certification DSA-C03 Exam disposably.

>> Valid Exam Snowflake DSA-C03 Preparation <<

Latest Snowflake DSA-C03 Learning Material | Latest DSA-C03 Exam Simulator

We believe that you can buy our DSA-C03 demo PDF torrent without any misgivings, Firstly, we have a strong experts team who are devoted themselves to research of the technology, which ensure the high-quality of our DSA-C03 Dump guide, ExamsLabs offers SnowPro Advanced: Data Scientist Certification Exam DSA-C03 free Updates. It is no exaggeration to say that the value of the certification training materials is equivalent to all exam related reference books.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q40-Q45):

NEW QUESTION # 40

A healthcare provider has a Snowflake table 'MEDICAL RECORDS' containing patient notes stored as unstructured text in a column called 'NOTE TEXT'. They want to identify different patient groups based on the topics discussed in these notes. They aim

to use a combination of unsupervised and supervised learning. Which of the following represents a robust workflow to achieve this goal?

- A. MultiOutputClassifier wrapped around a Logistic Regression model) within Snowflake (using Snowpark), using the original 'NOTE TEXT' as input features (TF-IDF or word embeddings) and the manually assigned topic labels as target variables. Use the trained model to classify the remaining patient notes into relevant patient groups.
- B. Perform topic modeling (e.g., LDA) directly on the 'NOTE_TEXT' column using a Python UDF in Snowflake. Manually label a subset of the resulting topics. Then, train a supervised classifier (e.g., Naive Bayes) to predict the identified topics for new patient notes.
- C. **Perform topic modeling on a sample of the 'NOTE TEXT' data using a Snowflake Python UDF. Manually review the top documents for each identified topic, and assign labels describing the patient group represented by each topic. Train a supervised multi-label classification model (e.g., using scikit-learn's**
- D. Export all 'NOTE TEXT' data to an external system, use an existing NLP pipeline for topic modeling and manual labeling, then create a Snowflake UDF that replicates this entire pipeline internally.
- E. Use a Snowflake external function to call a pre-trained topic modeling model (e.g., BERTopic) hosted on Google Cloud AI Platform. Assign topic probabilities to each patient note. Then, perform K-Means clustering on the topic probabilities to identify patient segments. No manual labeling is performed.

Answer: C

Explanation:

Option D is the most comprehensive and practical. First, it uses unsupervised topic modeling to discover potential patient groups. Second, it uses manual labeling to create a supervised training dataset. Third, it trains a supervised multi-label classification model within Snowflake (using Snowpark), allowing for automated patient group assignment based on the text of their notes, leveraging TF-IDF or word embeddings for feature representation. This balances the efficiency of unsupervised learning with the accuracy of supervised learning. It also highlights Snowflake's ability to directly train and deploy models using Snowpark.

NEW QUESTION # 41

You are building a model to predict loan defaults using a dataset stored in Snowflake. After training your model and calculating residuals, you create a scatter plot of the residuals against the predicted values. The plot shows a cone-shaped pattern, with residuals spreading out more as the predicted values increase. Which of the following SQL queries, run within a Snowpark Python session, could be used to address the underlying issue indicated by this residual pattern, assuming the predicted values are stored in a column named and the residuals in a column named 'loan_default_residuar' in a Snowflake table named 'loan_predictions'?

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

Answer: A

Explanation:

A cone-shaped pattern in the residuals plot (heteroscedasticity) indicates that the variance of the errors is not constant. Applying a transformation like Box-Cox to the target variable before retraining the model (Option D) is the most appropriate way to address this. Option A attempts to filter outliers based on the residuals, but does not address the heteroscedasticity itself and requires statistical functions unavailable within standard SQL. Option B attempts to take the natural log of the residuals, which is nonsensical as residuals can be negative. Option C attempts to filter based on the rank of residuals, which is similarly unhelpful, does not fix the problem, and uses inappropriate outlier removal with SQL QUALIFY clause. Option E scaling the features might sometimes improve model performance, but it does not directly address heteroscedasticity.

NEW QUESTION # 42

You are a data scientist working with a Snowflake table named 'CUSTOMER DATA' that contains a 'PHONE NUMBER' column stored as VARCHAR. The 'PHONE NUMBER' column sometimes contains non-numeric characters like hyphens and parentheses, and in some rows the data is missing. You need to create a new table 'CLEANED CUSTOMER DATA' with a column named 'CLEANED PHONE NUMBER' that contains only the numeric part of the phone number (as VARCHAR) and replaces missing or invalid phone numbers with NULL. Which of the following Snowpark Python code snippets achieves this most efficiently, ensuring no errors occur during the data transformation, and considers Snowflake's performance best practices?

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

Answer: B

Explanation:

Option E is the most efficient because it leverages Snowpark's built-in functions for string manipulation and conditional logic directly. It first removes all non-numeric characters using 'regexp_replace' and then uses 'iff(if and only if)' to replace empty strings (resulting from cleaning) with NULL. This approach avoids using UDFs (User-Defined Functions), which can introduce overhead. Option B, although using 'regexp_replace', requires an additional 'with_column' to handle empty strings after cleaning. Option A introduces UDF that decreases performance. Option C calls UDF with undefined 'call_udf' function and 'snowflake-snowpark-python' library. Option D is missing dataframe and its transformation is not happening on top of Dataframe. Option E is preferable over Option B, as it uses the single transformation.

NEW QUESTION # 43

You have trained a complex Random Forest model in Snowflake to predict loan default risk. You wish to understand the individual and combined effects of 'credit_score' and 'debt_to_income_ratio' on the predicted probability of default. Which approach is MOST suitable for visualizing and interpreting these relationships?

- A. Fit a simpler linear model (e.g., Logistic Regression) to the data and interpret its coefficients.
- B. Calculate feature importance using SNOWFLAKE.ML.FEATURE IMPORTANCE and focus on the features with the highest scores.
- C. Examine the model's overall accuracy (e.g., AUC) and assume the relationships are well-represented.
- D. Create a two-way Partial Dependence Plot (PDP) showing the interaction between 'credit_score' and 'debt_to_income_ratio'.
- E. Generate individual Partial Dependence Plots (PDPs) for 'credit_score' and 'debt_to_income_ratio'.

Answer: D

Explanation:

The correct answer is C. While individual PDPs (option B) provide insights into the individual effects of each feature, a two-way PDP specifically visualizes and helps interpret the interaction between 'credit_score' and 'debt_to_income_ratio'. This is crucial for understanding how the combined effect of these features influences the predicted probability of default. Feature importance (option A) indicates feature relevance but doesn't show the nature of the relationship. Simplifying the model (option D) sacrifices the complexity captured by the Random Forest. Overall accuracy (option E) doesn't provide specific insights into feature relationships.

NEW QUESTION # 44

A financial services company wants to predict loan defaults. They have a table 'LOAN APPLICATIONS' with columns 'application_id', 'applicant_income', 'applicant_age', and 'loan_amount'. You need to create several derived features to improve model performance.

Which of the following derived features, when used in combination, would provide the MOST comprehensive view of an applicant's financial stability and ability to repay the loan? Select all that apply

- A. Requires external data from a credit bureau to determine total debt, then calculated as 'total_debt / applicant_income' (Assume credit bureau integration is already in place)
- B. Calculated as 'applicant_income / loan_amount'.
- C. Calculated as 'loan_amount / applicant_age'.
- D. Calculated as 'applicant_age * applicant_age'.
- E. Calculated as 'applicant_age / applicant_income'.

Answer: A,B,C

Explanation:

The best combination provides diverse perspectives on financial stability. 'total_debt / applicant_income' directly reflects the applicant's ability to cover the loan with their income. 'loan_amount / applicant_age' represents the loan burden relative to the applicant's age and can expose risk in younger, less established applicants. 'applicant_age * applicant_age' provides the most comprehensive view, including existing debt obligations from external data. 'age_squared' and are less directly

informative about repayment ability. They could potentially capture non-linear relationships, but 'age_squared' is more likely to introduce overfitting, relies on an external data source, making it a powerful, but potentially more complex, feature to implement.

NEW QUESTION # 45

ExamsLabs is a very wonderful and effective platform to give chances to our worthy clients who want to achieve their expected scores and gain their DSA-C03 certifications. With our professional experts' tireless efforts, our DSA-C03 exam torrent is equipped with a simulated examination system with timing function, allowing you to examine your learning results at any time, keep checking for defects, and improve your strength. Whether you are a student or an office worker, you can be satisfied here, and you will never regret if you choose our DSA-C03 Exam Torrent.

Latest DSA-C03 Learning Material: <https://www.examslabs.com/Snowflake/SnowPro-Advanced/best-DSA-C03-exam-dumps.html>

Snowflake Valid Exam DSA-C03 Preparation Furthermore, they can be downloaded to all electronic devices so that you can have a rather modern study experience conveniently, Snowflake Valid Exam DSA-C03 Preparation When others are fighting to improve their vocational ability, if you still making no progress and take things as they are, then you will be eliminated, Snowflake Valid Exam DSA-C03 Preparation How can I renew my products after the expiry date?

Reducing the Cost or Price, I think it is a good time for you to achieve a salary increase and change your life style when you find DSA-C03 training dumps, Furthermore, they can be downloaded DSA-C03 to all electronic devices so that you can have a rather modern study experience conveniently.

100% Pass Snowflake Valid Exam DSA-C03 Preparation - Unparalleled SnowPro Advanced: Data Scientist Certification Exam

When others are fighting to improve their vocational ability, if you Associate DSA-C03 Level Exam still making no progress and take things as they are, then you will be eliminated, How can I renew my products after the expiry date?

ExamsLabs makes your investment secure with its money back guarantee policy. And there is no exaggeration that with our DSA-C03 training guide, you can get 100% pass guarantee.

P.S. Free & New DSA-C03 dumps are available on Google Drive shared by ExamsLabs: https://drive.google.com/open?id=17kD6agegdVx4KqkPkQUWJ_f-X62tgmbR