

Free PDF Snowflake - DSA-C03 - Professional SnowPro Advanced: Data Scientist Certification Exam Sample Exam



BTW, DOWNLOAD part of Actual4test DSA-C03 dumps from Cloud Storage: https://drive.google.com/open?id=1R24t9g-c98GT8NzD_Xwb0tN1MCIr4k3Z

We strongly recommend using our SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam dumps to prepare for the Snowflake DSA-C03 certification. It is the best way to ensure success. With our Snowflake DSA-C03 practice questions, you can get the most out of your studying and maximize your chances of passing your Snowflake DSA-C03 Exam. Actual4test Snowflake DSA-C03 practice test Actual4test is the answer if you want to score higher in the DSA-C03 exam and achieve your academic goals.

This way you will be able to experience the actual SnowPro Advanced: Data Scientist Certification Exam exam environment and become a more prepared and confident candidate to step into the examination center. You will know where exactly you stand before the actual Snowflake DSA-C03 Certification Exam. The actual Snowflake DSA-C03 exam questions will make you familiar with the inside-out view of the exam pattern and syllabus.

>> DSA-C03 Sample Exam <<

Pass-Sure DSA-C03 Sample Exam & Leading Offer in Qualification Exams & Marvelous Snowflake SnowPro Advanced: Data Scientist Certification Exam

If you want to choose the best DSA-C03 exam bootcamp, you should not miss our DSA-C03 exam materials. We have not only experienced industries elites who compile the high-quality products but also professional IT staff to develop three formats of our DSA-C03 study guide and the fast shopping environment. Buyers can enjoy free-worry shopping experience. Besides we provide one year free updates of our DSA-C03 training braindump and service warranty for buyers. With our DSA-C03 exam questions, your success is guaranteed.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample

Questions (Q110-Q115):

NEW QUESTION # 110

You are tasked with presenting a business case to stakeholders demonstrating the value of a new machine learning model that predicts customer churn. The model has been trained on data within Snowflake, and you have various metrics such as accuracy, precision, recall, and F I-score. You also have feature importance scores generated using a SHAP (SHapley Additive exPlanations) explainer. Which of the following visualization strategies, when combined, would MOST effectively communicate the model's performance and impact to a non-technical audience, while also providing sufficient detail for technical stakeholders?

- A. A scatter plot showing the relationship between two key features identified by SHAP, colored by the model's churn prediction, and a table summarizing the model's performance metrics (accuracy, precision, recall, F I-score). Additionally, include a waterfall plot for a specific customer, illustrating how each feature contributes to the final prediction.
- B. A simple bar chart showing the overall accuracy score of the model alongside a table detailing the precision, recall, and F I-score. Include a word cloud of the most important features from the SHAP values.
- C. A distribution plot (e.g., histogram or KDE) of the predicted churn probabilities, segmented by actual churn status (churned vs. not churned), combined with a SHAP force plot visualizing the feature contributions for a single, randomly selected customer who churned. Add a section on potential cost savings from churn reduction.
- D. A confusion matrix visualizing the true positives, true negatives, false positives, and false negatives, along with a summary plot of the SHAP values showing the impact of each feature on the model's prediction for a representative sample of customers. A line chart showing cumulative churn rate across different customer segments.
- E. A ROC curve (Receiver Operating Characteristic) showing the trade-off between true positive rate and false positive rate, paired with a detailed table of all feature importance scores generated by the SHAP explainer. Present statistical summaries, such as mean and standard deviation, of the top 5 feature values, grouped by predicted churn probability.

Answer: A,D

Explanation:

Options B and D provide a balanced approach for both technical and non-technical audiences- A confusion matrix (Option B) is easily understandable and shows model performance across different prediction outcomes. A summary plot of SHAP values clearly illustrates feature importance and direction of impact. A line chart showing cumulative churn rate across different customer segments highlights the business value-Option D is also highly effective because scatter plots can be easily understood, especially when colored by churn prediction- The table of model metrics provides necessary details. The waterfall plot brings the explanation down to an individual customer level, making the model's behavior more tangible. Options A, C and E have deficits- Option A lacks detailed performance visualization. Option C is technical and might confuse non-technical stakeholders. Option E has too many summary plots.

NEW QUESTION # 111

You are analyzing website traffic data stored in a Snowflake table named 'WEB_EVENTS'. This table contains a 'TIMESTAMP' column representing when the event occurred and a 'PAGE_VIEWS' column indicating the number of page views for that event. You need to identify the day with the highest number of page views and also the day with lowest number of page views along with average number of page views. How can you accomplish this using Snowflake SQL?

- `SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews DESC LIMIT 1; SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews ASC LIMIT 1; SELECT AVG(PAGE_VIEWS) FROM WEB_EVENTS`
- `SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews DESC; SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews ASC;SELECT AVG(PAGE_VIEWS) FROM WEB_EVENTS`
- `SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews, RANK() OVER (ORDER BY TotalViews DESC) AS ViewRank FROM WEB_EVENTS GROUP BY EventDate ORDER BY ViewRank; SELECT AVG(PAGE_VIEWS) FROM WEB_EVENTS`
- `SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews DESC LIMIT 1; SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews ASC LIMIT 1; SELECT AVG(PAGE_VIEWS) OVER() FROM WEB_EVENTS LIMIT 1;`
- `SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews DESC LIMIT 1; SELECT DATE(TIMESTAMP) AS EventDate, SUM(PAGE_VIEWS) AS TotalViews FROM WEB_EVENTS GROUP BY EventDate ORDER BY TotalViews ASC LIMIT 1; SELECT APPROX_AVG(PAGE_VIEWS) FROM WEB_EVENTS`

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

Answer: E

Explanation:

Option D provides the correct answer. The first two queries correctly identify the day with the highest and lowest total views, respectively, using 'DATE(TIMESTAMP)' to extract the date, to aggregate page views, 'GROUP BY' to group by date, 'ORDER BY' to sort, and 'LIMIT 1' to select only the top/bottom day. It also has the correct query to identify average page views 'SELECT OVER() FROM WEB_EVENTS LIMIT 1'. Other Options A and E are quite close but they don't identify the same, in option A, 'SELECT AVG(PAGE_VIEWS) FROM WEB_EVENTS', the AVG page views won't tell us the dates of min max and Avg views. Similar is the problem with option E, 'SELECT FROM WEB_EVENTS The APPROX_AVG won't tell us which day has highest or lowest.

NEW QUESTION # 112

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?

- `'''sql SELECT cortex.complete('Answer the following question about the product features: ' || question, 'snowflake_cortex.llama2_70b_chat') AS response;'''`
- `'''sql SELECT cortex.complete('As a helpful customer support agent, answer the following question about our product features: ' || question || 'Keep your response concise and friendly.', 'snowflake_cortex.llama2_70b_chat') AS response;'''`
- `'''sql SELECT cortex.complete('You are a product expert. Provide only the description of the feature. Do not include anything else. ' || question, 'snowflake_cortex.llama2_70b_chat') AS response;'''`
- `'''sql SELECT cortex.complete('Given the following product details: ' || (SELECT array_agg(feature_name || ':' || feature_description) FROM PRODUCT_DETAILS) || '. Answer the question: ' || question, 'snowflake_cortex.llama2_70b_chat') AS response;'''`
- `'''sql SELECT cortex.complete('Product Feature description from Product detail table is: ' || (SELECT feature_description FROM PRODUCT_DETAILS where feature_name = feature_name) || '. Answer the question: ' || question, 'snowflake_cortex.llama2_70b_chat') AS response;'''`

- A. Option A
- B. Option E
- C. Option B
- 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 # 113

You are building a predictive model on customer churn using Snowflake data. You observe that the distribution of 'TIME SINCE LAST PURCHASE' is heavily left-skewed. Which of the following strategies would be MOST appropriate to handle this skewness before feeding the data into a linear regression model to improve its performance? (Select TWO)

- **A. Apply a square root transformation to the 'TIME SINCE LAST PURCHASE' column.**
- B. Standardize the 'TIME SINCE LAST PURCHASE' column using Z-score normalization.
- **C. Use a winsorization technique to cap extreme values in the 'TIME SINCE LAST PURCHASE' column at a predefined percentile (e.g., 99th percentile).**
- D. Remove all records with 'TIME SINCE LAST PURCHASE' values below the mean.
- E. Apply a logarithmic transformation to the 'TIME SINCE LAST PURCHASE' column.

Answer: A,C

Explanation:

For left-skewed data, a square root transformation (Option B) can help reduce the impact of smaller values and bring the distribution

closer to normal. Winsorization (Option C) can mitigate the influence of extreme values on the left tail of the distribution, making the data more suitable for linear regression. Logarithmic transformation is more suitable for right-skewed data (Option A). Z-score normalization (Option D) centers the data around zero but doesn't change the skewness. Removing records below the mean (Option E) is generally not a good practice as it can introduce bias and lose valuable information.

NEW QUESTION # 114

You are analyzing customer transaction data in Snowflake to identify fraudulent activities. The 'TRANSACTION_AMOUNT' column exhibits a right-skewed distribution. Which of the following Snowflake queries is MOST effective in identifying outliers based on the Interquartile Range (IQR) method, specifically targeting unusually large transaction amounts? Assume IQR is already calculated as variable and Q1 as and Q3 as in snowflake session.

- A. SELECT TRANSACTION ID FROM TRANSACTIONS WHERE TRANSACTION_AMOUNT < q1 - (1.5 iqr);
- **B. SELECT TRANSACTION ID FROM TRANSACTIONS WHERE TRANSACTION_AMOUNT > q3 + (1.5 iqr);**
- C. SELECT TRANSACTION ID FROM TRANSACTIONS WHERE TRANSACTION_AMOUNT > (SELECT MEDIAN(TRANSACTION_AMOUNT) FROM TRANSACTIONS);
- D. SELECT TRANSACTION ID FROM TRANSACTIONS WHERE TRANSACTION_AMOUNT > (SELECT + 3 FROM TRANSACTIONS);
- E. SELECT TRANSACTION ID FROM TRANSACTIONS WHERE TRANSACTION_AMOUNT > (SELECT WITHIN GROUP (ORDER BY TRANSACTION_AMOUNT) FROM TRANSACTIONS);

Answer: B

Explanation:

Option B correctly implements the IQR method for identifying outliers. The formula 'Q3 + (1.5 identifies values significantly higher than the third quartile, which is appropriate for detecting outliers in a right-skewed distribution. Options A uses standard deviation, which is less robust to outliers. Option C finds the 95th percentile but might not isolate extreme outliers. Option D looks for lower bound outliers, which is not the case. Option E uses the median which not effective to find out outlier values.

NEW QUESTION # 115

.....

The proper answer to your questions is Actual4test. When studying for the SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) certification exam, Actual4test is one of the most helpful resources. Actual4test guarantees success on the first try by providing you with actual SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam questions in PDF, desktop practice exam software, and a web-based practice exam.

DSA-C03 Training Kit: https://www.actual4test.com/DSA-C03_examcollection.html

Snowflake DSA-C03 Sample Exam We ensure that all objectives of the exam are covered in depth so you'll be ready for any question on the exam, Snowflake DSA-C03 Sample Exam You can decide whether to buy it or not until you think our products are truly helpful, In order to pass Snowflake DSA-C03 exam easily, many candidates are eager to find the most helpful DSA-C03 dumps guide: SnowPro Advanced: Data Scientist Certification Exam anxiously as the best shortcut, Choosing right study materials is a smart way for most office workers who have enough time and energy to attending classes about DSA-C03 reliable study questions.

It takes time for a signal to move down a cable and it takes even more DSA-C03 Premium Files time for a repeater to rebuild, retime, and retransmit a signal, One of the most common forms of payment is the personal check.

Professional DSA-C03 Sample Exam & Leading Offer in Qualification Exams & Free Download Snowflake SnowPro Advanced: Data Scientist Certification Exam

We ensure that all objectives of the exam are covered in depth so you'll DSA-C03 Actual Dumps be ready for any question on the exam, You can decide whether to buy it or not until you think our products are truly helpful.

In order to pass Snowflake DSA-C03 Exam easily, many candidates are eager to find the most helpful DSA-C03 dumps guide: SnowPro Advanced: Data Scientist Certification Exam anxiously as the best shortcut.

