

Valid DSA-C03 Valid Examcollection bring you Fantastic DSA-C03 Boot Camp for Snowflake SnowPro Advanced: Data Scientist Certification Exam



DOWNLOAD the newest Fast2test DSA-C03 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1hUWTZa1ueP2dOJTK35A3BOGMc6_EW-bw

This is a Snowflake DSA-C03 practice exam software for Windows computers. This DSA-C03 practice test will be similar to the actual SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam. If user wish to test the Snowflake DSA-C03 study material before joining Fast2test, they may do so with a free sample trial. This DSA-C03 Exam simulation software can be readily installed on Windows-based computers and laptops. Since it is desktop-based Snowflake DSA-C03 practice exam software, it is not necessary to connect to the internet to use it.

The Fast2test DSA-C03 PDF questions file, desktop practice test software, and web-based practice test software, all these three DSA-C03 practice test questions formats are ready for instant download. Just download any Snowflake DSA-C03 Exam Questions format and start this journey with confidence.

>> **DSA-C03 Valid Examcollection <<**

DSA-C03 Boot Camp & DSA-C03 Valid Test Papers

With the unemployment rising, large numbers of people are forced to live their job. It is hard to find a high salary job than before. Many people are immersed in updating their knowledge. So people are keen on taking part in the DSA-C03 exam. As you know, the competition between candidates is fierce. If you want to win out, you must master the knowledge excellently. Now our DSA-C03 Study Materials are your best choice. With the assistance of our study materials, you will advance quickly.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q255-Q260):

NEW QUESTION # 255

You are tasked with building a model to predict customer churn. You have a table named in Snowflake with the following relevant

columns: 'customer_id', 'login_date', 'orders_placed', and 'churned' (binary indicator). You want to engineer features that capture customer engagement over time using Snowpark for Python. Which of the following feature engineering steps, applied sequentially, are MOST effective in creating features indicative of churn risk?

- A. 1. Calculate the number of days since the customer's last login, and use nulls instead of negative numbers to indicate inactivity. 2. Calculate the rolling 7-day average of 'orders_placed' using a window function, partitioning by 'customer_id' and ordering by 'login_date'. 3. Calculate the slope of a linear regression of 'page_views' over time for each customer, indicating the trend in engagement using Snowpark ML. 4. Calculate the percentage of weeks the customer logged in. 5. Create a feature showing standard deviation of 'page_views' per customer over the last 90 days.
- B. 1. Calculate the maximum 'page_views' in a single day for each customer. 2. Calculate the total number of days with no 'login_date' for each customer. 3. Create a feature indicating if a customer has ever placed an order. 4. Use a simple boolean for the 'subscription_type' column.
- C. 1. Calculate the average 'page_views' per week for each customer over the last 3 months using a window function. 2. Calculate the recency of the last order (days since last order) for each customer. 3. Create a feature indicating the change in average daily page views over the last month compared to the previous month. 4. Create a feature showing standard deviation of 'page_views' per customer over the last 90 days.
- D. 1. Calculate the average 'page_views' per day for each customer. 2. Calculate the total number of days for each customer. 3. Create a feature indicating whether the customer has a premium subscription ('subscription_type' = 'premium').
- E. 1. Calculate the total 'page_views' and 'orders_placed' for each customer without considering time. 2. Use one-hot encoding for the 'subscription_type' column.

Answer: A,C

Explanation:

Options B and E are the MOST effective because they incorporate time-based features and indicators of engagement trends. Recency (days since last order) captures the time elapsed since the customer's last interaction. Calculating changes in page views, the number of login days and linear regression slope identifies trends in engagement. Rolling averages smooth out daily fluctuations and capture longer-term patterns. Standard deviation of page views indicates a trend in page view variance, and thus overall customer engagement variance. Option A lacks recency and trend information. Option C misses temporal analysis. Option D has less relevance features and can be used however it is more useful to compare how well a customer is engaged with previous activity.

NEW QUESTION # 256

A data scientist needs to calculate the cumulative moving average of sales for each product in a table. The table contains columns: (INT), (DATE), and (NUMBER). The desired output should include the product_id, sale_date, and Which of the following Snowflake SQL statements correctly calculates the cumulative moving average for each product using window functions?

- A. `SELECT product_id, sale_date, daily_sales, AVG(daily_sales) OVER (PARTITION BY product_id) AS cumulative_average FROM sales_by_day;`
- B. `SELECT product_id, sale_date, daily_sales, AVG(daily_sales) OVER (ORDER BY sale_date ASC) AS cumulative_average FROM sales_by_day;`
- C. `SELECT product_id, sale_date, daily_sales, OVER (PARTITION BY product_id ORDER BY sale_date ASC ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW) AS cumulative_average FROM`
- D. `SELECT product_id, sale_date, daily_sales, OVER (PARTITION BY product_id ORDER BY sale_date ASC) AS cumulative_average FROM sales_by_day;`
- E. `SELECT product_id, sale_date, daily_sales, OVER (PARTITION BY product_id ORDER BY sale_date ASC) / OVER (PARTITION BY product_id ORDER BY sale_date ASC) AS cumulative_average FROM sales_by_day;`

Answer: C,E

Explanation:

Both options B and D are correct. Option B correctly uses the 'AVG()' window function with the 'PARTITION BY product_id' clause to calculate the average sales for each product independently, and 'ORDER BY sale_date ASC' along with 'ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW' ensures a cumulative average is calculated from the beginning of the product's sales history up to the current date. Option D also calculates the cumulative moving average by summing the and dividing by the current row number, effectively computing the average.

NEW QUESTION # 257

You are using the NetworkX library in Snowpark Python to analyze social network data stored in a Snowflake table named 'USER CONNECTIONS', which has columns 'USER ID' and 'CONNECTED USER' representing connections between users. You want

to find the users with the highest 'betweenness centrality' to identify influential nodes in the network. Which Snowpark Python code snippet would correctly calculate and display the top 5 users with the highest betweenness centrality?

- A.

```
import networkx as nx
from snowflake.snowpark.functions import col
 snowflake

# Assume 'snowpark_session' is your Snowpark session object
df = snowpark_session.table('USER_CONNECTIONS')
G = nx.Graph()
for row in df.collect():
    G.add_edge(row['USER_ID'], row['CONNECTED_USER_ID'])

betweenness = nx.betweenness_centrality(G)

betweenness_df = snowpark_session.createDataFrame(list(betweenness.items()), schema=['USER_ID', 'Betweenness'])

top_5 = betweenness_df.orderBy(col('Betweenness').desc()).limit(5)

top_5.show()
```

- B.

```
import networkx as nx
import pandas as pd
 snowflake

# Assume 'snowpark_session' is your Snowpark session object
df = snowpark_session.table('USER_CONNECTIONS').to_pandas()
G = nx.from_pandas_edgelist(df, 'USER_ID', 'CONNECTED_USER_ID')

betweenness = nx.betweenness_centrality(G)

# Convert betweenness dictionary to DataFrame
betweenness_df = pd.DataFrame(list(betweenness.items()), columns=['USER_ID', 'Betweenness'])

top_5 = betweenness_df.sort_values('Betweenness', ascending=False).head(5)

print(top_5)
```

- C.

```
import networkx as nx
import pandas as pd
 snowflake

# Assume 'snowpark_session' is your Snowpark session object
df = snowpark_session.table('USER_CONNECTIONS').to_pandas()
G = nx.from_pandas_edgelist(df, 'USER_ID', 'CONNECTED_USER_ID')

betweenness = nx.betweenness_centrality(G)

# Convert betweenness dictionary to DataFrame
betweenness_df = pd.DataFrame(list(betweenness.items()), columns=['USER_ID', 'Betweenness'])

top_5 = betweenness_df.sort_values('Betweenness', ascending=False, ignore_index=True).head(5)

print(top_5)
```

- D.

```

import networkx as nx
import pandas as pd
 snowflake

# Assume 'snowpark_session' is your Snowpark session object
df = snowpark_session.table('USER_CONNECTIONS').to_pandas()
G = nx.Graph()
for index, row in df.iterrows():
    G.add_edge(row['USER_ID'], row['CONNECTED_USER_ID'])

betweenness = nx.betweenness_centrality(G)

betweenness_df = pd.DataFrame(list(betweenness.items()), columns=['USER_ID', 'Betweenness'])

top_5 = betweenness_df.sort_values('Betweenness', ascending=False).head(5)

print(top_5)

```

- E

```

import networkx as nx
import pandas as pd
 snowflake

# Assume 'snowpark_session' is your Snowpark session object
df = snowpark_session.sql("SELECT USER_ID, CONNECTED_USER_ID FROM USER_CONNECTIONS").to_pandas()
i = nx.from_pandas_edgelist(df, 'USER_ID', 'CONNECTED_USER_ID')

betweenness = nx.betweenness_centrality(G)

# Convert betweenness dictionary to DataFrame
betweenness_df = pd.DataFrame(list(betweenness.items()), columns=['USER_ID', 'Betweenness'])

top_5 = betweenness_df.nlargest(5, 'Betweenness')

print(top_5)

```

Answer: B

Explanation:

Option A is the most efficient and correct approach. It leverages for directly creating the graph from a Pandas DataFrame (converted from the Snowpark DataFrame), calculates betweenness centrality using NetworkX, creates a Pandas DataFrame for results, and then sorts and displays the top 5 users. Options B iterates through the rows, which is less efficient, and attempts to create a Snowpark DataFrame from the betweenness dictionary, which isn't the most efficient output mechanism in this context. Option C is almost correct but uses 'nlargest' which is also valid. Option D uses which is slower and less efficient than Option E is very close but includes the parameter which is unnecessary for this specific operation since it's the initial creation of the betweenness df and the index isn't crucial to be reset. So while it functions it is redundant.

NEW QUESTION # 258

You are using Snowpark Feature Store to manage features for your machine learning models. You've created several Feature Groups and now want to consume these features for training a model. To optimize retrieval, you want to use point-in-time correctness. Which of the following actions/configurations are essential to ensure point-in-time correctness when retrieving features using Snowpark Feature Store?

- A. Create an associated Stream on the source tables used for Feature Groups
- B. Use the method on the Feature Store client, providing a dataframe containing the 'primary_keyS and the desired for each record.
- C. When creating Feature Groups, specify a 'timestamp_key' that represents the event timestamp of the data in the source tables.
- D. Ensure that all source tables used by the Feature Groups have Change Data Capture (CDC) enabled.

- E. Explicitly specify a in the call.

Answer: B,C

Explanation:

Options B and C are correct. B: Specifying a 'timestamp_key' during Feature Group creation is crucial for enabling point-in-time correctness. This tells the Feature Store which column represents the event timestamp. C: The method is specifically designed for point-in-time lookups. It requires a dataframe containing primary keys and the desired timestamp for each lookup. This enables the Feature Store to retrieve the feature values as they were at that specific point in time. Option A is incorrect, while enabling CDC is valuable for incremental updates, it does not guarantee point-in-time correctness without specifying the timestamp key and retrieving historical features using that key. Option D is not necessary, streams enable incremental loads but are separate from point in time. Option E, is not needed, its implicit via using .

NEW QUESTION # 259

You are using Snowpark Pandas to prepare data for a machine learning model. You have a Snowpark DataFrame named 'transactions_df' that contains transaction data, including 'transaction_id', 'product_id', 'customer_id', and 'transaction_amount'. You want to create a new feature that represents the average transaction amount per customer. However, you are concerned about potential skewness in the 'transaction_amount' and want to apply a log transformation to reduce its impact before calculating the average. Which of the following steps using Snowpark Pandas would achieve this transformation and calculation most efficiently within Snowflake?

```

 import snowflake.snowpark.functions as F #This code downloads all data to the client side, then performs the transformation using pandas.
transactions_pdf = transactions_df.to_pandas()
transactions_pdf['log_transaction_amount'] = np.log1p(transactions_pdf['transaction_amount']) #
Calculate the average transaction amount per customer. mean_transaction_amount = transactions_pdf.groupby('customer_id')
['log_transaction_amount'].mean() # convert back to snowpark dataframe mean_transaction_amount_sdf = session.createDataFrame(mean_transaction_amount)

 import snowflake.snowpark.functions as F #This code performs the transformation within Snowflake. transactions_df =
transactions_df.withColumn('log_transaction_amount', F.log1p(F.col('transaction_amount'))) # Calculate the average transaction amount per customer.
mean_transaction_amount = transactions_df.groupBy('customer_id').agg(F.mean('log_transaction_amount').alias('avg_log_transaction_amount'))

 import snowflake.snowpark.functions as F import pandas as pd #This is not valid as it can't work directly on a snowpark dataframe.
transactions_df['log_transaction_amount'] = np.log1p(transactions_df['transaction_amount']) # Calculate the average transaction amount per customer.
mean_transaction_amount = transactions_df.groupby('customer_id')['log_transaction_amount'].mean()

 import snowflake.snowpark.functions as F # Calculate the average transaction amount per customer, without log transformation
mean_transaction_amount = transactions_df.groupBy('customer_id').agg(F.mean('transaction_amount').alias('avg_transaction_amount'))

 import snowflake.snowpark.functions as F import pandas as pd # Convert Spark DataFrame to Pandas DataFrame, apply log transformation, and
calculate the mean mean_transaction_amount = transactions_df.to_pandas().groupby('customer_id').agg({'transaction_amount': lambda x:
np.log1p(x).mean()}) # Convert back to Snowpark Dataframe to use in SQL transformation mean_transaction_amount_sdf =
session.createDataFrame(mean_transaction_amount)

```

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

Answer: C

Explanation:

Option B is the most efficient solution because it performs both the log transformation and the average calculation entirely within Snowflake using Snowpark functions. This avoids the overhead of transferring the data to the client side. It uses F.log1p() to apply the log transformation to the 'transaction_amount' column, handling potential zero values gracefully. It groups by 'customer_id' and uses F.mean() to calculate the average of the transformed transaction amounts.

NEW QUESTION # 260

.....

If you choose the test DSA-C03 certification and then buy our DSA-C03 study materials you will get the panacea to both get the useful certificate and spend little time. Passing the test certification can help you stand out in your colleagues and have a bright future in your career. If you buy our DSA-C03 Study Materials you odds to pass the test will definitely increase greatly.

DSA-C03 Boot Camp: <https://www.fast2test.com/DSA-C03-premium-file.html>

We will send you the latest DSA-C03 study dumps through your email, so please check your email then, Snowflake DSA-C03

Valid Examcollection As it gets you a report for your mock test, which enables you to measure that where you need to put more efforts, SnowPro Advanced: Data Scientist Certification Exam DSA-C03 exam dumps are available in an eBook and software format, So you rest assured that with the Snowflake DSA-C03 exam real questions you can make the best SnowPro Advanced: Data Scientist Certification Exam exam preparation strategy and plan.

Both sets of cynics would be right in some circumstances, TWiki at Motorola, We will send you the latest DSA-C03 Study Dumps through your email, so please check your email then.

As it gets you a report for your mock test, which enables you to measure that where you need to put more efforts, SnowPro Advanced: Data Scientist Certification Exam DSA-C03 exam dumps are available in an eBook and software format.

Precise DSA-C03 Valid Examcollection - Complete & Perfect DSA-C03 Materials Free Download for Snowflake DSA-C03 Exam

So you rest assured that with the Snowflake DSA-C03 exam real questions you can make the best SnowPro Advanced: Data Scientist Certification Exam exam preparation strategy and plan. Because we are committed to customers who decide to choose our DSA-C03 study tool.

BONUS!!! Download part of Fast2test DSA-C03 dumps for free: <https://drive.google.com/open?id=1hUWTZa1ueP2dOJTK35A3BOGMc6> EW-bw