

DSA-C03 Latest Test Practice - DSA-C03 Latest Training



BTW, DOWNLOAD part of BraindumpsVCE DSA-C03 dumps from Cloud Storage: <https://drive.google.com/open?id=1PQw0HKlyGZ8pScvAPRRLpnFEWF24ZfjG>

This allows candidates to choose the format that best suits their learning style and preference, ensuring a seamless and effective exam preparation experience. By offering tailored solutions to meet individual needs, BraindumpsVCE has established itself as a trusted provider of top-quality SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam preparation material.

The name of these formats are Snowflake DSA-C03 PDF dumps file, desktop practice test software, and web-based practice test software. All these three Snowflake Cloud DSA-C03 practice test formats are easy to use and perfectly work with all devices, operating systems, and web browsers. The DSA-C03 Pdf Dumps file is a simple collection of Real and Updated SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam questions in PDF format and it is easy to install and use.

>> **DSA-C03 Latest Test Practice** <<

DSA-C03 Latest Training, Reliable DSA-C03 Exam Review

We here guarantee that we will never sell the personal information of our candidates. There is no need for you to worry about the individual privacy under our rigorous privacy DSA-C03 protection system. As regards purchasing, our website and DSA-C03 study materials are absolutely safe and free of virus. For further consideration we will provide professional IT personnel to guide your installation and the use of our DSA-C03 Study Materials remotely. So you can buy our DSA-C03 study materials without any misgivings. If you have any questions, please you contact us online through the email.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q127-Q132):

NEW QUESTION # 127

You are using a Snowflake Notebook to build a churn prediction model. You have engineered several features, and now you want to visualize the relationship between two key features: and , segmented by the target variable 'churned' (boolean). Your goal is to create an interactive scatter plot that allows you to explore the data points and identify any potential patterns.

Which of the following approaches is most appropriate and efficient for creating this visualization within a Snowflake Notebook?

- A. Use the Snowflake Connector for Python to fetch the data, then leverage a Python visualization library like Plotly or Bokeh to generate an interactive plot within the notebook.
- B. Write a stored procedure in Snowflake that generates the visualization data in a specific format (e.g., JSON) and then use a JavaScript library within the notebook to render the visualization.
- C. Use the 'snowflake-connector-python' to pull the data and use 'seaborn' to create static plots.
- D. Create a static scatter plot using Matplotlib directly within the Snowflake Notebook by converting the data to a Pandas DataFrame. This involves pulling all relevant data into the notebook's environment before plotting.
- E. Leverage Snowflake's native support for Streamlit within the notebook to create an interactive application. Query the data directly from Snowflake within the Streamlit app and use Streamlit's plotting capabilities for visualization.

Answer: E

Explanation:

Option D, leveraging Snowflake's native support for Streamlit, is the most appropriate and efficient approach. Streamlit allows you to build interactive web applications directly within the notebook, querying data directly from Snowflake and using Streamlit's built-in plotting capabilities (or integrating with other Python visualization libraries). This avoids pulling large amounts of data into the notebook's environment, which is crucial for large datasets. Option A is inefficient due to the data transfer overhead and limited interactivity. Option B can work but is not as streamlined as using Streamlit within the Snowflake environment. Option C will create static plots only. Option E is overly complex and less efficient than using Streamlit.

NEW QUESTION # 128

You are tasked with building a machine learning pipeline in Snowpark Python to predict customer lifetime value (CLTV). You need to access and manipulate data residing in multiple Snowflake tables and views, including customer demographics, purchase history, and website activity. To improve code readability and maintainability, you decide to encapsulate data access and transformation logic within a Snowpark Stored Procedure. Given the following Python code snippet representing a simplified version of your stored procedure:

- A. The `'session.write_pandas(df, table_name='CLTV PREDICTIONS', auto_create_table=True)` function writes the Pandas DataFrame 'df' containing the CLTV predictions directly to a new Snowflake table named , automatically creating the table if it does not exist.
- B. The `'session.sql('SELECT FROM PURCHASE` line executes a SQL query against the Snowflake database and returns the results as a list of Row objects.
- C. The `'session.table('CUSTOMER DEMOGRAPHICS')` method creates a local Pandas DataFrame containing a copy of the data from the 'CUSTOMER DEMOGRAPHICS' table.
- D. The `'snowflake.snowpark.context.get_active_session()'` function retrieves the active Snowpark session object, enabling interaction with the Snowflake database from within the stored procedure.
- E. The `replace=True, packages=['snowflake-snowpark-python', 'pandas',` decorator registers the Python function as a Snowpark Stored Procedure, allowing it to be called from SQL.

Answer: A,B,D,E

Explanation:

Option A is correct because is the standard method for accessing the active Snowpark session within a stored procedure. Option C is correct as the `gsproc` decorator is required to register the function as a Snowpark Stored Procedure, specifying necessary packages. Option D correctly explains how to execute SQL queries using the session object and retrieve results. Option E accurately describes the function's ability to write a Pandas DataFrame to a Snowflake table and create it if it doesn't exist. Option B is incorrect because returns a Snowpark DataFrame, not a Pandas DataFrame. A Snowpark DataFrame is a lazily evaluated representation of the data, while a Pandas DataFrame is an in-memory copy.

NEW QUESTION # 129

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'?

```

○ PUT file:///path/to/model.joblib @model_stage; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON
RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib') HANDLER = 'main.predict' AS $$ import joblib import pandas as pd from
snowflake.types import PandasDataFrame class PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__))
self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred =
self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON
RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' TARGET_PATH = '@model_stage/model.joblib' AS $$
import joblib import pandas as pd import os class PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model =
joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main =
PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES
= ('scikit-learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os class
PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df =
pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-
learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os import snowflake.snowpark class PythonPredictor: def
__init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred =
self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage OVERWRITE=TRUE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib','pandas','snowflake-snowpark-
python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os import snowflake.snowpark class PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__))
self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

```

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

Answer: D

Explanation:

Option E is correct. `AUTO_COMPRESS=FALSE` isn't strictly needed for joblib files, and adding `OVERWRITE=TRUE` ensures the model is updated in case of re-deployment. Importantly, the `IMPORTS` parameter correctly specifies the location of the model within the stage, and includes necessary packages (pandas and snowflake-snowpark-python) and the code correctly loads the model from the stage within the Python handler. Also snowflake.snowpark need to be imported to avoid errors with snowpark library. Option A lacks pandas, snowflake-snowpark-python, imports section. Option B has wrong syntax TARGET_PATH which is incorrect syntax. Option C doesn't import snowpark, so might give error. Option D doesn't OVERWRITE flag, and can give error if trying to overwrite files.

NEW QUESTION # 130

You have trained a classification model in Snowflake using Snowpark ML to predict customer churn. After deploying the model, you observe that the model performs well on the training data but poorly on new, unseen data'. You suspect overfitting. Which of the following strategies can be applied within Snowflake to detect and mitigate overfitting during model validation, considering the model is already deployed and receiving inference requests through a Snowflake UDF?

- A. Create shadow UDFs that score data using alternative models. Compare the performance metrics (such as accuracy, precision, recall) between the production UDF and shadow UDFs using Snowflake's query capabilities. If shadow models consistently outperform the production model on certain data segments, retrain the production model incorporating those data segments with higher weights.
- B. Calculate the Area Under the Precision-Recall Curve (AUPRC) using Snowflake SQL on both the training and validation datasets. A significant difference indicates overfitting. Then, retrain the model in Snowpark ML with added L1 or L2 regularization, adjusting the regularization strength based on validation set performance, and redeploy the UDF.
- C. Implement k-fold cross-validation within the Snowpark ML training pipeline using Snowflake's distributed compute. Track the mean and standard deviation of the performance metrics (e.g., accuracy, F1-score) across folds. A high variance suggests overfitting. Use this information to tune hyperparameters or select a simpler model architecture before deployment.
- D. Monitor the UDF execution time in Snowflake. A sudden increase in execution time indicates overfitting. Use the 'EXPLAIN' command on the UDF's underlying SQL query to identify performance bottlenecks and rewrite the query for optimization.
- E. Since the model is already deployed, the only option is to collect inference requests and compare the distributions of predicted values in each batch with the predicted values on the training set. A large difference indicates overfitting; model must be retrained outside of the validation process.

Answer: B,C

Explanation:

Options A and C are correct because they describe strategies for detecting and mitigating overfitting during the model validation process using Snowflake's capabilities. AUPRC is a good performance metric to compare the training vs validation set results to catch overfitting, and regularization can be used to avoid overfitting. Option C directly incorporates cross-validation into the model training workflow within Snowflake, allowing for early detection and mitigation of overfitting through hyperparameter tuning and

model selection. Option B is incorrect because it focuses on performance optimization, not overfitting. Option D describes an AIB testing or champion-challenger setup which could be a strategy to use to detect data drift over time, but not overfitting. E is only partially correct as it describes one way to detect data drift, but not overfitting.

NEW QUESTION # 131

A pharmaceutical company is testing a new drug to lower blood pressure. They conduct a clinical trial with 200 patients. After treatment, the sample mean reduction in systolic blood pressure is 10 mmHg, with a sample standard deviation of 15 mmHg. You want to construct a 99% confidence interval for the true mean reduction in systolic blood pressure. Which of the following statements is most accurate concerning the appropriate distribution and critical value to use?

- A. Use a z-distribution because the sample size is large ($n > 30$), and the critical value is approximately 2.576.
- B. Use a t-distribution with 200 degrees of freedom, and the critical value is close to 2.576.
- C. Use a t-distribution with 199 degrees of freedom, and the critical value is slightly larger than 2.576.
- D. Use a z-distribution because we are estimating mean, and use a critical value of 1.96.
- E. Use a chi-squared distribution with 199 degrees of freedom.

Answer: C

Explanation:

The correct answer is B. While the sample size is considered 'large' ($n > 30$), it's more accurate to use a t-distribution when the population standard deviation is unknown and estimated by the sample standard deviation. The t-distribution accounts for the added uncertainty from estimating the standard deviation. The degrees of freedom are $n-1 = 199$. The critical value for a 99% confidence interval with a t-distribution and 199 degrees of freedom will be slightly larger than the z-score of 2.576. Option A is incorrect because using t-distribution is slightly better. Option C is incorrect because chi-squared distribution is for variance/standard deviation. Option D is incorrect since 1.96 is z score for 95%. Option E is incorrect as the degrees of freedom should be $n-1$.

NEW QUESTION # 132

.....

Created on the exact pattern of the Actual DSA-C03 Tests, BraindumpsVCE's dumps comprise questions and answers and provide all important information in easy to grasp and simplified content. The easy language does not pose any barrier for any learner. The complex portions of the certification syllabus have been explained with the help of simulations and real-life based instances. The best part of BraindumpsVCE's dumps is their relevance, comprehensiveness and precision. You need not to try any other source for exam preparation. The innovatively crafted dumps will serve you the best; imparting you information in fewer number of questions and answers.

DSA-C03 Latest Training: https://www.braindumpsvce.com/DSA-C03_exam-dumps-torrent.html

Use these tools for your help and guidance and then you will definitely be happy with the results of BraindumpsVCE's DSA-C03 updated audio study guide and DSA-C03 from BraindumpsVCE cbt f To have enjoyable and fantastic journey for the DSA-C03 latest computer based training you must trust the helping tools of BraindumpsVCE which are made to deal with this purpose, In order to provide all customers with the suitable study materials, a lot of experts from our company designed the DSA-C03 training materials.

Reasoning: Using Arithmetic to Solve Your Problems, With a bunch of courteous employees DSA-C03 and staff dedicated to the aftersales stuff enthusiastically, Use these tools for your help and guidance and then you will definitely be happy with the results of BraindumpsVCE's DSA-C03 updated audio study guide and DSA-C03 from BraindumpsVCE cbt f To have enjoyable and fantastic journey for the DSA-C03 latest computer based training you must trust the helping tools of BraindumpsVCE which are made to deal with this purpose.

IT-Tests DSA-C03 Test Study Guide, Answer Snowflake DSA-C03 Practice Exam Questions

In order to provide all customers with the suitable study materials, a lot of experts from our company designed the DSA-C03 training materials, As long as you never give up yourself, you are bound to become successful.

When you choose our DSA-C03 valid training dumps, you will enjoy one year free update for DSA-C03 pdf torrent without any additional cost, Contact live chat to know more.

