Databricks
Machine-Learning-Associate
Databricks Certified Machine Learning Associate

Questions And Answers PDF Format:

For More Information – Visit link below:
https://www.certsgrade.com/

Version = **Product**

Latest Version: 6.0

Question: 1

A senior machine learning engineer is developing a machine learning pipeline. They set up the pipeline to automatically transition a new version of a registered model to the Production stage in the Model Registry once it passes all tests using the MLflow Client API client.
Which operation was used to transition the model to the Production stage?
Response:

A. Client.update_model_stage
B. client.transition_model_version_stage
C. client.transition_model_version
D. client.update_model_version

Answer: B

Question: 2

A data scientist has computed updated rows that contain new feature values for primary keys already stored in the Feature Store table features. The updated feature values are stored in the DataFrame features_df.
They want to update the rows in features if the associated primary key is in features_df. If a row's primary key is not in features_df, they want the row to remain unchanged in features.
Which code block using the Feature Store Client fs can be used to accomplish this task?
Response:

A. fs.write_table(
name="features",
df=features_df,
mode="merge"
)
B. fs.write_table(
name="features",
df=features_df,
mode="overwrite"
)
C. fs.write_table(
name="features",
df=features_df,
)
D. fs.create_table(
name="features",

```
df=features_df,
mode="append"
)
E. fs.refresh_table(
name="features",
df=features_df,
mode="overwrite"
)
```

Answer: A

Question: 3

When AutoML explores the key attributes of a dataset, which of the following elements does it typically not assess?
Response:

A. The dataset's memory footprint.
B. The potential impact of outliers on model performance.
C. The balance or imbalance of classes in classification tasks.
D. The encryption level of the dataset.

Answer: D

Question: 4

Where can you find the code that was executed with a run in the MLflow UI?
Response:

A. In the run's metadata section.
B. Inside the associated Git repository.
C. Under the "Code" tab in the run's details page.
D. It is not possible to view the executed code in the MLflow UI.

Answer: C

Question: 5

A data scientist has developed a two-class decision tree classifier using Spark ML and computed the predictions in a Spark DataFrame preds_df with the following schema:
- prediction DOUBLE
- actual DOUBLE

Which of the following code blocks can be used to compute the accuracy of the model according to the data in preds_df and assign it to the accuracy variable?
Response:

A. accuracy = RegressionEvaluator(
predictionCol="prediction",
labelCol="actual",
metricName="accuracy"
)
B. accuracy = MulticlassClassificationEvaluator(
predictionCol="prediction",
labelCol="actual",
metricName="accuracy"
)
accuracy = classification_evaluator.evaluate(preds_df)
C. classification_evaluator = BinaryClassificationEvaluator(
predictionCol="prediction",
labelCol="actual",
metricName="accuracy"
)
D. accuracy = Summarizer(
predictionCol="prediction",
labelCol="actual",
metricName="accuracy"
)
E. classification_evaluator = BinaryClassificationEvaluator(
predictionCol="prediction",
labelCol="actual",
metricName="accuracy"
)
accuracy = classification_evaluator.evaluate(preds_df)

Answer: E

Question: 6

How can you identify the best run using the MLflow Client API?
Response:

A. By manually reviewing each run's metrics.
B. Utilizing the search_runs function with a specific metric sort order.
C. Comparing run IDs manually for performance metrics.
D. Using a custom Python script outside of MLflow.

Answer: B

Question: 7

Which of the following are key components of ML workflows in Databricks?
Response:

A. Data ingestion
B. Model serving
C. Feature extraction
D. Manual model tuning

Answer: A,B,C

Question: 8

A machine learning team wants to use the Python library newpackage on all of their projects. They share a cluster for all of their projects. Which approach makes the Python library newpackage available to all notebooks run on a cluster?
Response:

A. Edit the cluster to use the Databricks Runtime for Machine Learning
B. Set the runtime-version variable in their Spark session to "ml"
C. Running %pip install newpackage once on any notebook attached to the cluster
D. Adding /databricks/python/bin/pip install newpackage to the cluster's bash init script
E. There is no way to make the newpackage library available on a cluster

Answer: D

Question: 9

A data scientist is developing a machine learning model. They made changes to their code in a text editor on their local machine, committed them to the project's Git repository, and pushed the changes to an online Git provider. Now, they want to load those changes into Databricks. The Databricks workspace contains an out-of-date version of the Git repository.
How can the data scientist complete this task?
Response:

A. Open the Repo Git dialog and enable automatic syncing.
B. Open the Repo Git dialog and click the "Sync" button.
C. Open the Repo Git dialog and click the "Merge" button.
D. Open the Repo Git dialog and enable automatic pulling.
E. Open the Repo Git dialog and click the "Pull" button.

Question: 10

Which of the following steps are necessary to commit changes from a Databricks Repo to an external Git provider?
(Select two)
Response:

A. Merge changes to the master branch in the external Git provider
B. Use Databricks notebooks to push changes
C. Stage and commit changes in the Databricks workspace
D. Pull requests from the Databricks workspace to the Git provider

Answer: B,C

PRODUCT FEATURES

- **100% Money Back Guarantee**
- **90 Days Free updates**
- **Special Discounts on Bulk Orders**
- **Guaranteed Success**
- **50,000 Satisfied Customers**
- **100% Secure Shopping**
- **Privacy Policy**
- **Refund Policy**

16 USD Discount Coupon Code: NB4XKTMZ