



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 1

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following describes concept drift?

- A. Concept drift is when there is a change in the distribution of an input variable
- B. Concept drift is when there is a change in the distribution of a target variable
- C. Concept drift is when there is a change in the relationship between input variables and target variables
- D. Concept drift is when there is a change in the distribution of the predicted target given by the model
- E. None of these describe Concept drift

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 2

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer is monitoring categorical input variables for a production machine learning application. The engineer believes that missing values are becoming more prevalent in more recent data for a particular value in one of the categorical input variables.

Which of the following tools can the machine learning engineer use to assess their theory?

- A. Kolmogorov-Smirnov (KS) test
- B. One-way Chi-squared Test
- C. Two-way Chi-squared Test
- D. Jenson-Shannon distance
- E. None of these

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 3

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist is using MLflow to track their machine learning experiment. As a part of each MLflow run, they are performing hyperparameter tuning. The data scientist would like to have one parent run for the tuning process with a child run for each unique combination of hyperparameter values.

They are using the following code block:

```
with mlflow.start_run(run_name="Parent run") as run:
    print("Start parent run")
with mlflow.start_run(run_name="Child 1", nested=True):
    mlflow.log_param("run_name", "child_1")
with mlflow.start_run(run_name="Child 2", nested=True):
    mlflow.log_param("run_name", "child_2")
```

The code block is not nesting the runs in MLflow as they expected.

Which of the following changes does the data scientist need to make to the above code block so that it successfully nests the child runs under the parent run in MLflow?

- A. Indent the child run blocks within the parent run block
- B. Add the nested=True argument to the parent run
- C. Remove the nested=True argument from the child runs
- D. Provide the same name to the run_name parameter for all three run blocks
- E. Add the nested=True argument to the parent run and remove the nested=True arguments from the child runs

Show Suggested Answer

Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 4

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer wants to log feature importance data from a CSV file at path `importance_path` with an MLflow run for model `model`.

Which of the following code blocks will accomplish this task inside of an existing MLflow run block?

- ```
mlflow.log_model_and_data(
 model,
 importance_path,
 "feature-importance.csv"
)
```
- A.
- ```
mlflow.log_model(  
    model,  
    importance_path,  
    "feature-importance.csv"  
)
```
- B.
- C. `mlflow.log_data(importance_path, "feature-importance.csv")`
- D. `mlflow.log_artifact(importance_path, "feature-importance.csv")`
- E. None of these code blocks can accomplish the task.

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 5

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following is a simple, low-cost method of monitoring numeric feature drift?

- A. Jensen-Shannon test
- B. Summary statistics trends
- C. Chi-squared test
- D. None of these can be used to monitor feature drift
- E. Kolmogorov-Smirnov (KS) test

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 6

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has developed a model to predict ice cream sales using the expected temperature and expected number of hours of sun in the day. However, the expected temperature is dropping beneath the range of the input variable on which the model was trained.

Which of the following types of drift is present in the above scenario?

- A. Label drift
- B. None of these
- C. Concept drift
- D. Prediction drift
- E. Feature drift

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 7

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist wants to remove the `star_rating` column from the Delta table at the location path. To do this, they need to load in data and drop the `star_rating` column. Which of the following code blocks accomplishes this task?

- A. `spark.read.format("delta").load(path).drop("star_rating")`
- B. `spark.read.format("delta").table(path).drop("star_rating")`
- C. Delta tables cannot be modified
- D. `spark.read.table(path).drop("star_rating")`
- E. `spark.sql("SELECT * EXCEPT star_rating FROM path")`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 8

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following operations in Feature Store Client fs can be used to return a Spark DataFrame of a data set associated with a Feature Store table?

- A. fs.create_table
- B. fs.write_table
- C. fs.get_table
- D. There is no way to accomplish this task with fs
- E. fs.read_table

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 9

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer is in the process of implementing a concept drift monitoring solution. They are planning to use the following steps:

1. Deploy a model to production and compute predicted values
2. Obtain the observed (actual) label values
3. _____
4. Run a statistical test to determine if there are changes over time

Which of the following should be completed as Step #3?

- A. Obtain the observed values (actual) feature values
- B. Measure the latency of the prediction time
- C. Retrain the model
- D. None of these should be completed as Step #3
- E. Compute the evaluation metric using the observed and predicted values

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 10

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following is a reason for using Jensen-Shannon (JS) distance over a Kolmogorov-Smirnov (KS) test for numeric feature drift detection?

- A. All of these reasons
- B. JS is not normalized or smoothed
- C. None of these reasons
- D. JS is more robust when working with large datasets
- E. JS does not require any manual threshold or cutoff determinations

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 11

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist is utilizing MLflow to track their machine learning experiments. After completing a series of runs for the experiment with experiment ID `exp_id`, the data scientist wants to programmatically work with the experiment run data in a Spark DataFrame. They have an active MLflow Client `client` and an active Spark session `spark`.

Which of the following lines of code can be used to obtain run-level results for `exp_id` in a Spark DataFrame?

- A. `client.list_run_infos(exp_id)`
- B. `spark.read.format("delta").load(exp_id)`
- C. There is no way to programmatically return row-level results from an MLflow Experiment.
- D. `mlflow.search_runs(exp_id)`
- E. `spark.read.format("mlflow-experiment").load(exp_id)`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 12

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has developed and logged a scikit-learn random forest model model, and then they ended their Spark session and terminated their cluster. After starting a new cluster, they want to review the `feature_importances_` of the original model object.

Which of the following lines of code can be used to restore the model object so that `feature_importances_` is available?

- A. `mlflow.load_model(model_uri)`
- B. `client.list_artifacts(run_id)["feature-importances.csv"]`
- C. `mlflow.sklearn.load_model(model_uri)`
- D. This can only be viewed in the MLflow Experiments UI
- E. `client.pyfunc.load_model(model_uri)`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 13

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following is a simple statistic to monitor for categorical feature drift?

- A. Mode
- B. None of these
- C. Mode, number of unique values, and percentage of missing values
- D. Percentage of missing values
- E. Number of unique values

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 14

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following is a probable response to identifying drift in a machine learning application?

- A. None of these responses
- B. Retraining and deploying a model on more recent data
- C. All of these responses
- D. Rebuilding the machine learning application with a new label variable
- E. Sunsetting the machine learning application

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 15

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has computed updated feature values for all primary key values stored in the Feature Store table features. In addition, feature values for some new primary key values have also been computed. The updated feature values are stored in the DataFrame features_df. They want to replace all data in features with the newly computed data.

Which of the following code blocks can they use to perform this task using the Feature Store Client fs?

A.

```
fs.create_table(  
    name="features",  
    df=features_df,  
    mode="overwrite"  
)
```

B.

```
fs.write_table(  
    name="features",  
    df=features_df,  
)
```

C.

```
fs.write_table(  
    name="features",  
    df=features_df,  
    mode="merge"  
)
```

D.

```
fs.write_table(  
    name="features",  
    df=features_df,  
    mode="overwrite"  
)
```

E.

```
fs.create_table(  
    name="features",  
    df=features_df,  
    mode="merge"  
)
```

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 16

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

After a data scientist noticed that a column was missing from a production feature set stored as a Delta table, the machine learning engineering team has been tasked with determining when the column was dropped from the feature set.

Which of the following SQL commands can be used to accomplish this task?

- A. VERSION
- B. DESCRIBE
- C. HISTORY
- D. DESCRIBE HISTORY
- E. TIMESTAMP

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 17

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following describes label drift?

- A. Label drift is when there is a change in the distribution of the predicted target given by the model
- B. None of these describe label drift
- C. Label drift is when there is a change in the distribution of an input variable
- D. Label drift is when there is a change in the relationship between input variables and target variables
- E. Label drift is when there is a change in the distribution of a target variable

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 18

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following machine learning model deployment paradigms is the most common for machine learning projects?

- A. On-device
- B. Streaming
- C. Real-time
- D. Batch
- E. None of these deployments

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 19

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist would like to enable MLflow Autologging for all machine learning libraries used in a notebook. They want to ensure that MLflow Autologging is used no matter what version of the Databricks Runtime for Machine Learning is used to run the notebook and no matter what workspace-wide configurations are selected in the Admin Console.

Which of the following lines of code can they use to accomplish this task?

- A. `mlflow.sklearn.autolog()`
- B. `mlflow.spark.autolog()`
- C. `spark.conf.set("autologging", True)`
- D. It is not possible to automatically log MLflow runs.
- E. `mlflow.autolog()`

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 20

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has developed a model and computed the RMSE of the model on the test set. They have assigned this value to the variable `rmse`. They now want to manually store the RMSE value with the MLflow run.

They write the following incomplete code block:

```
with mlflow.start_run(experiment_id=exp_id, run_name=run_name) as run:
    # Log rmse
    mlflow.____("rmse", rmse)
```

Which of the following lines of code can be used to fill in the blank so the code block can successfully complete the task?

- A. `log_artifact`
- B. `log_model`
- C. `log_metric`
- D. `log_param`
- E. There is no way to store values like this.

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 21

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following MLflow operations can be used to automatically calculate and log a Shapley feature importance plot?

- A. `mlflow.shap.log_explanation`
- B. None of these operations can accomplish the task.
- C. `mlflow.shap`
- D. `mlflow.log_figure`
- E. `client.log_artifact`

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 22

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has developed a scikit-learn random forest model model, but they have not yet logged model with MLflow. They want to obtain the input schema and the output schema of the model so they can document what type of data is expected as input.

Which of the following MLflow operations can be used to perform this task?

- A. `mlflow.models.schema.infer_schema`
- B. `mlflow.models.signature.infer_signature`
- C. `mlflow.models.Model.get_input_schema`
- D. `mlflow.models.Model.signature`
- E. There is no way to obtain the input schema and the output schema of an unlogged model.

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 23

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer and data scientist are working together to convert a batch deployment to an always-on streaming deployment. The machine learning engineer has expressed that rigorous data tests must be put in place as a part of their conversion to account for potential changes in data formats.

Which of the following describes why these types of data type tests and checks are particularly important for streaming deployments?

- A. Because the streaming deployment is always on, all types of data must be handled without producing an error
- B. All of these statements
- C. Because the streaming deployment is always on, there is no practitioner to debug poor model performance
- D. Because the streaming deployment is always on, there is a need to confirm that the deployment can autoscale
- E. None of these statements

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 24

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following deployment paradigms can centrally compute predictions for a single record with exceedingly fast results?

- A. Streaming
- B. Batch
- C. Edge/on-device
- D. None of these strategies will accomplish the task.
- E. Real-time

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 25

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineering team wants to build a continuous pipeline for data preparation of a machine learning application. The team would like the data to be fully processed and made ready for inference in a series of equal-sized batches.

Which of the following tools can be used to provide this type of continuous processing?

- A. Spark UDFs
- B. Structured Streaming
- C. MLflow
- D. Delta Lake
- E. AutoML

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 26

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer wants to deploy a model for real-time serving using MLflow Model Serving. For the model, the machine learning engineer currently has one model version in each of the stages in the MLflow Model Registry. The engineer wants to know which model versions can be queried once Model Serving is enabled for the model.

Which of the following lists all of the MLflow Model Registry stages whose model versions are automatically deployed with Model Serving?

- A. Staging, Production, Archived
- B. Production
- C. None, Staging, Production, Archived
- D. Staging, Production
- E. None, Staging, Production

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 27

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has written a function to track the runs of their random forest model. The data scientist is changing the number of trees in the forest across each run. Which of the following MLflow operations is designed to log single values like the number of trees in a random forest?

- A. `mlflow.log_artifact`
- B. `mlflow.log_model`
- C. `mlflow.log_metric`
- D. `mlflow.log_param`
- E. There is no way to store values like this.

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 28

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer is converting a Hyperopt-based hyperparameter tuning process from manual MLflow logging to MLflow Autologging. They are trying to determine how to manage nested Hyperopt runs with MLflow Autologging.

Which of the following approaches will create a single parent run for the process and a child run for each unique combination of hyperparameter values when using Hyperopt and MLflow Autologging?

- A. Starting a manual parent run before calling fmin
- B. Ensuring that a built-in model flavor is used for the model logging
- C. Starting a manual child run within the objective_function
- D. There is no way to accomplish nested runs with MLflow Autologging and Hyperopt
- E. MLflow Autologging will automatically accomplish this task with Hyperopt

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 29

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has created a Python function `compute_features` that returns a Spark DataFrame with the following schema:

```
customer_id STRING,  
spend DOUBLE,  
units INT,  
loyal INT,  
region STRING
```

The resulting DataFrame is assigned to the `features_df` variable. The data scientist wants to create a Feature Store table using `features_df`.

Which of the following code blocks can they use to create and populate the Feature Store table using the Feature Store Client `fs`?

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    df=features_df,  
    description="Customer features"  
)
```

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    description="Customer features"  
)
```

C. `features_df.write.mode("fs").path("new_table")`

```
fs.create_table(  
    name="new_table",  
    primary_keys="customer_id",  
    function=compute_features,  
    description="Customer features"  
)
```

E. `features_df.write.mode("feature").path("new_table")`

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 30

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following is a benefit of logging a model signature with an MLflow model?

- A. The model will have a unique identifier in the MLflow experiment
- B. The schema of input data can be validated when serving models
- C. The model can be deployed using real-time serving tools
- D. The model will be secured by the user that developed it
- E. The schema of input data will be converted to match the signature

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 31

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following statements describes streaming with Spark as a model deployment strategy?

- A. The inference of batch processed records as soon as a trigger is hit
- B. The inference of all types of records in real-time
- C. The inference of batch processed records as soon as a Spark job is run
- D. The inference of incrementally processed records as soon as trigger is hit
- E. The inference of incrementally processed records as soon as a Spark job is run

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 32

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer has deployed a model recommender using MLflow Model Serving. They now want to query the version of that model that is in the Production stage of the MLflow Model Registry.

Which of the following model URIs can be used to query the described model version?

- A. `https://<databricks-instance>/model-serving/recommender/Production/invocations`
- B. The version number of the model version in Production is necessary to complete this task.
- C. `https://<databricks-instance>/model/recommender/stage-production/invocations`
- D. `https://<databricks-instance>/model-serving/recommender/stage-production/invocations`
- E. `https://<databricks-instance>/model/recommender/Production/invocations`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 33

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following tools can assist in real-time deployments by packaging software with its own application, tools, and libraries?

- A. Cloud-based compute
- B. None of these tools
- C. REST APIs
- D. Containers
- E. Autoscaling clusters

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 34

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer has registered a sklearn model in the MLflow Model Registry using the sklearn model flavor with UI model_uri. Which of the following operations can be used to load the model as an sklearn object for batch deployment?

- A. `mlflow.spark.load_model(model_uri)`
- B. `mlflow.pyfunc.read_model(model_uri)`
- C. `mlflow.sklearn.read_model(model_uri)`
- D. `mlflow.pyfunc.load_model(model_uri)`
- E. `mlflow.sklearn.load_model(model_uri)`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 35

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist set up a machine learning pipeline to automatically log a data visualization with each run. They now want to view the visualizations in Databricks. Which of the following locations in Databricks will show these data visualizations?

- A. The MLflow Model Registry Model page
- B. The Artifacts section of the MLflow Experiment page
- C. Logged data visualizations cannot be viewed in Databricks
- D. The Artifacts section of the MLflow Run page
- E. The Figures section of the MLflow Run page

Show Suggested Answer



Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 36

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A data scientist has developed a scikit-learn model `sklearn_model` and they want to log the model using MLflow.

They write the following incomplete code block:

```
with mlflow.start_run(experiment_id=exp_id, run_name=run_name) as run:  
    # Log model  
    _____
```

Which of the following lines of code can be used to fill in the blank so the code block can successfully complete the task?

- A. `mlflow.spark.track_model(sklearn_model, "model")`
- B. `mlflow.sklearn.log_model(sklearn_model, "model")`
- C. `mlflow.spark.log_model(sklearn_model, "model")`
- D. `mlflow.sklearn.load_model("model")`
- E. `mlflow.sklearn.track_model(sklearn_model, "model")`

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 37

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

Which of the following describes the concept of MLflow Model flavors?

- A. A convention that deployment tools can use to wrap preprocessing logic into a Model
- B. A convention that MLflow Model Registry can use to version models
- C. A convention that MLflow Experiments can use to organize their Runs by project
- D. A convention that deployment tools can use to understand the model
- E. A convention that MLflow Model Registry can use to organize its Models by project

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 38

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

In a continuous integration, continuous deployment (CI/CD) process for machine learning pipelines, which of the following events commonly triggers the execution of automated testing?

- A. The launch of a new cost-efficient SQL endpoint
- B. CI/CD pipelines are not needed for machine learning pipelines
- C. The arrival of a new feature table in the Feature Store
- D. The launch of a new cost-efficient job cluster
- E. The arrival of a new model version in the MLflow Model Registry

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 39

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineering team has written predictions computed in a batch job to a Delta table for querying. However, the team has noticed that the querying is running slowly. The team has already tuned the size of the data files. Upon investigating, the team has concluded that the rows meeting the query condition are sparsely located throughout each of the data files.

Based on the scenario, which of the following optimization techniques could speed up the query by colocating similar records while considering values in multiple columns?

- A. Z-Ordering
- B. Bin-packing
- C. Write as a Parquet file
- D. Data skipping
- E. Tuning the file size

Show Suggested Answer





Actual exam question from Databricks's Certified Machine Learning Professional

Question #: 40

Topic #: 1

[\[All Certified Machine Learning Professional Questions\]](#)

A machine learning engineer needs to deliver predictions of a machine learning model in real-time. However, the feature values needed for computing the predictions are available one week before the query time.

Which of the following is a benefit of using a batch serving deployment in this scenario rather than a real-time serving deployment where predictions are computed at query time?

- A. Batch serving has built-in capabilities in Databricks Machine Learning
- B. There is no advantage to using batch serving deployments over real-time serving deployments
- C. Computing predictions in real-time provides more up-to-date results
- D. Testing is not possible in real-time serving deployments
- E. Querying stored predictions can be faster than computing predictions in real-time

Show Suggested Answer

