

Oracle

1Z0-1110-26

Oracle Cloud Infrastructure Data Science Professional

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/1z0-1110-26>

Latest Version: 6.0

Question: 1

A bike sharing platform has collected user commute data for the past 3 years. For increasing profitability and making useful inferences, a machine learning model needs to be built from the accumulated data. Which of the following options has the correct order of the required machine learning tasks for building a model?

- A. Data Access, Feature Exploration, Data Exploration, Feature Engineering, Modeling
- B. Data Access, Data Exploration, Feature Exploration, Feature Engineering, Modeling
- C. Data Access, Data Exploration, Feature Engineering, Feature Exploration, Modeling
- D. Data Access, Feature Exploration, Feature Engineering, Data Exploration, Modeling

Answer: C

Explanation:

Detailed Answer in Step-by-Step Solution:

Data Access: The first step in any machine learning workflow is accessing the raw data. This involves retrieving the user commute data collected over the past 3 years from the bike-sharing platform's storage system.

Data Exploration: Once data is accessed, it's explored to understand its structure, quality, and patterns (e.g., missing values, distributions). This step helps identify what preprocessing is needed.

Feature Engineering: After understanding the data, features are created or transformed (e.g., commute duration, time of day) to improve model performance. This step precedes feature exploration because you need engineered features to analyze further.

Feature Exploration: This involves analyzing the engineered features (e.g., correlation analysis, importance ranking) to refine them or select the most relevant ones for modeling.

Modeling: Finally, the prepared data and features are used to train and evaluate a machine learning model.

Option C (Data Access, Data Exploration, Feature Engineering, Feature Exploration, Modeling) follows this logical sequence, aligning with standard ML workflows.

The correct order reflects the machine learning lifecycle as outlined in Oracle's OCI Data Science documentation. Data Access is the initial step to retrieve data, followed by Data Exploration to assess it (e.g., using OCI Data Science Notebook Sessions with tools like pandas). Feature Engineering transforms raw data into meaningful inputs, followed by Feature Exploration to analyze feature importance (e.g., using ADS SDK's correlation tools). Modeling is the final step where the model is built and trained. This sequence is consistent with Oracle's recommended practices for building ML models in OCI Data Science (Oracle Cloud Infrastructure Data Science Service Documentation, "Machine Learning Lifecycle").

Question: 2

You have been given a collection of digital files required for a business audit. They consist of several different formats that you would like to annotate using Oracle Cloud Infrastructure (OCI) Data Labeling. Which THREE types of files could this tool annotate?

- A. Video footage of a conversation in a conference room
- B. Images of computer server racks
- C. A typewritten document that details an annual budget
- D. A collection of purchase orders for office supplies
- E. An audio recording of a phone conversation

Answer: A, B, C

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI Data Labeling Capabilities: OCI Data Labeling is designed to annotate data for machine learning, supporting specific file types like images, text documents, and videos.

Evaluate Options:

- A . Video footage: Supported for tasks like object detection or action recognition.
- B . Images: Supported for image classification, object detection, etc.
- C . Typewritten document: Supported as text data for tasks like entity extraction or classification.
- D . Purchase orders: While potentially text-based, this is ambiguous without format clarification (e.g., PDF, image). OCI supports text annotation, but “purchase orders” isn’t a specific file type—it’s assumed as text here.
- E . Audio recording: Not supported, as OCI Data Labeling focuses on visual and textual data, not audio.

Select Three: A (video), B (images), and C (text documents) are explicitly supported file types. OCI Data Labeling supports annotating datasets of images, text, and videos, as per the official documentation. Video footage (A) can be annotated for tasks like object tracking, images (B) for classification or detection, and typewritten documents (C) for text-based annotations (e.g., named entity recognition). Audio files (E) are not supported, and while purchase orders (D) could be text, the question specifies “typewritten document” as a clearer match. (Oracle Cloud Infrastructure Data Labeling Service Documentation, "Supported Data Types").

Question: 3

Which TWO statements about Oracle Cloud Infrastructure (OCI) Open Data service are true?

- A. Open Data includes text and image data repositories for AI and ML.
- B. Audio and video formats are not available.

- C. Each dataset in Open Data consists of code and tooling usage examples for consumption and reproducibility.
- D. Open Data is a dataset repository made for the people that create, use, and manipulate datasets.
- E. A primary goal of Open Data is for users to contribute to the data repositories to expand the content offered.
- F. Subscribers can pay and log into Open Data to view curated datasets that are otherwise not available to the public.

Answer: A, D

Explanation:

Detailed Answer in Step-by-Step Solution:

Analyze OCI Open Data: OCI Open Data is a free service providing access to public datasets for AI/ML use cases.

Evaluate Statements:

A: True—Open Data includes text and image datasets (e.g., geospatial images).

B: False—Video and other formats may be available depending on the dataset; no strict exclusion exists.

C: False—Datasets may include metadata, but code/tooling examples aren't guaranteed.

D: True—It's designed for data scientists and analysts who work with datasets.

E: False—It's not a user-contributed repository; it's curated by Oracle.

F: False—Open Data is free and public, not subscription-based.

Select Two: A and D align with the service's purpose and offerings.

OCI Open Data provides access to datasets like text and images (A) for AI/ML, aimed at data professionals (D). It's a free, curated service, not user-contributed (E) or paid (F), and while it focuses on certain formats, it doesn't explicitly exclude audio/video (B). (Oracle Cloud Infrastructure Open Data Documentation, "Overview of Open Data").

Question: 4

You are running a pipeline in the OCI Data Science service and want to override some of the pipeline's default settings. Which of the following statements about overriding pipeline defaults is true?

- A. Pipeline defaults can be overridden only during pipeline creation.
- B. Pipeline defaults can be overridden only by the Administrator.
- C. Pipeline defaults can be overridden before starting the pipeline execution.
- D. Pipeline defaults cannot be overridden once the pipeline has been created.

Answer: C

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI Data Science Pipelines: Pipelines automate ML workflows with configurable steps.

Check Override Mechanism: Defaults (e.g., compute shape, storage) can be modified before execution via the OCI Console, SDK, or CLI.

Evaluate Options:

A: False—Overrides can occur post-creation, before running.

B: False—Any authorized user, not just admins, can override defaults.

C: True—Settings can be adjusted before execution starts.

D: False—Defaults can be changed post-creation, pre-execution.

Conclusion: C is correct as it reflects the flexibility of pipeline configuration.

OCI Data Science Pipelines allow users to override default settings (e.g., compute resources, environment variables) before execution, as noted in the official documentation. This can be done via the UI or programmatically, offering flexibility beyond creation time (A) and without admin-only restrictions (B). (Oracle Cloud Infrastructure Data Science Pipelines Documentation, "Configuring Pipelines").

Question: 5

How are datasets exported in the OCI Data Labeling service?

A. As a binary file

B. As an XML file

C. As a line-delimited JSON file

D. As a CSV file

Answer: C

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI Data Labeling Export: After annotation, datasets are exported for ML use.

Check Supported Formats: OCI Data Labeling exports annotations in a structured, machine-readable format.

Evaluate Options:

A: Binary isn't a standard export format for annotations.

B: XML isn't used; JSON is preferred for flexibility.

C: Line-delimited JSON is the correct format, aligning with ML workflows.

D: CSV is common but not the default for OCI Data Labeling.

Conclusion: C matches the official export format.

OCI Data Labeling exports annotated datasets as line-delimited JSON files, which store each annotation as a separate JSON object per line, suitable for ML pipelines. This is explicitly stated in the documentation. (Oracle Cloud Infrastructure Data Labeling Service Documentation, "Exporting Datasets").

Question: 6

As a data scientist for a hardware company, you have been asked to predict the revenue demand for the upcoming quarter. You develop a time series forecasting model to analyze the data.

a. Select the correct sequence of steps to predict the revenue demand values for the upcoming quarter.

- A. Verify, prepare model, deploy, save, predict
- B. Predict, deploy, save, verify, prepare model
- C. Prepare model, deploy, verify, save, predict
- D. Prepare model, verify, save, deploy, predict

Answer: D

Explanation:

Detailed Answer in Step-by-Step Solution:

Prepare Model: Build and train the time series model using historical data.

Verify: Validate the model's accuracy (e.g., using metrics like MAE or RMSE).

Save: Store the trained model (e.g., in the OCI Model Catalog).

Deploy: Make the model available for predictions (e.g., via OCI Model Deployment).

Predict: Generate revenue forecasts for the upcoming quarter.

Evaluate Options: D follows this logical flow; others (e.g., A starts with "verify" before preparation) don't.

In OCI Data Science, the workflow for time series forecasting involves preparing the model (training), verifying its performance, saving it to the catalog, deploying it, and then predicting. This sequence is standard for ML deployment in OCI, as per the documentation. (Oracle Cloud Infrastructure Data Science Documentation, "Time Series Forecasting Workflow").

Question: 7

You have trained a binary classifier for a loan application and saved this model into the model catalog. A colleague wants to examine the model, and you need to share the model with your colleague. From the model catalog, which model artifacts can be shared?

- A. Metadata, hyperparameters, metrics only
- B. Model metadata and hyperparameters only
- C. Models and metrics only
- D. Models, model metadata, hyperparameters, metrics

Answer: D

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand Model Catalog: The OCI Model Catalog stores trained models and associated details.

Identify Shareable Artifacts: When sharing, all components—model file, metadata (e.g., name, description), hyperparameters (e.g., learning rate), and metrics (e.g., accuracy)—are accessible.

Evaluate Options:

A: Excludes the model itself—incorrect.

B: Excludes metrics—incorrect.

C: Excludes metadata and hyperparameters—incorrect.

D: Includes all components—correct.

Conclusion: D is comprehensive and accurate.

The OCI Model Catalog allows sharing of the model artifact (the trained model), metadata, hyperparameters, and performance metrics, enabling full examination by colleagues. This is detailed in the official documentation. (Oracle Cloud Infrastructure Data Science Documentation, "Model Catalog Overview").

Question: 8

You want to create a user group for a team of external data science consultants. The consultants should only have the ability to see Data Science resource details but not have the ability to create, delete, or update Data Science resources. What verb should you write in the policy?

- A. Use
- B. Inspect
- C. Manage
- D. Read

Answer: B

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI IAM Policies: Policies control access using verbs like “inspect,” “read,” “use,” and “manage.”

Define Requirement: Consultants need view-only access without modification rights.

Evaluate Verbs:

A . Use: Allows interaction (e.g., running jobs)—too permissive.

B . Inspect: Grants view-only access to resource details—matches requirement.

C . Manage: Full control (create, update, delete)—too permissive.

D . Read: Includes viewing content (e.g., data), not just metadata—slightly broader than needed.

Conclusion: “Inspect” (B) is the precise verb for view-only access to resource details.

In OCI Identity and Access Management (IAM), the “inspect” verb allows listing and viewing resource metadata without granting modification or data access rights, ideal for this scenario.

This is confirmed in the IAM policy reference. (Oracle Cloud Infrastructure Documentation, "IAM Policy Verbs").

Question: 9

You are a data scientist using Oracle AutoML to produce a model and you are evaluating the score metric for the model. Which TWO of the following prevailing metrics would you use for evaluating a multiclass classification model?

- A. Mean squared error
- B. Explained variance score
- C. Recall
- D. F1-score
- E. R-squared

Answer: C, D

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand Multiclass Classification: Metrics evaluate how well the model predicts multiple classes.

Evaluate Metrics:

- A . Mean squared error: Used for regression, not classification.
- B . Explained variance score: Regression metric, not suitable.
- C . Recall: Measures true positive rate per class—key for classification.
- D . F1-score: Balances precision and recall—widely used in multiclass.
- E . R-squared: Regression metric, not applicable.

Select Two: Recall (C) and F1-score (D) are standard for multiclass classification.

Oracle AutoML supports metrics like recall and F1-score for multiclass classification, as they assess per-class performance and overall precision-recall balance, respectively. Regression metrics (A, B,E) are irrelevant here. (Oracle Cloud Infrastructure Data Science Documentation, "AutoML Metrics").

Question: 10

You're going to create an Oracle Cloud Infrastructure Anomaly Detection model for multivariate data

- a. Where do you need to store the training data?
 - A. Your local machine
 - B. MySQL database
 - C. Autonomous Data Warehouse
 - D. Object Storage Bucket

Answer: D

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI Anomaly Detection: This service trains models to detect anomalies in multivariate data (e.g., sensor readings), requiring data to be accessible within OCI's ecosystem.

Assess Storage Requirements: The training data must be in a scalable, OCI-compatible location that the Anomaly Detection service can access programmatically.

Evaluate Options:

A . Your local machine: Data on a local machine isn't directly accessible to OCI services without upload, making it impractical for cloud-based training.

B . MySQL database: While OCI supports MySQL, Anomaly Detection doesn't natively integrate with it for training data; it prefers file-based input.

C . Autonomous Data Warehouse: This is a database for analytics, not the default storage for Anomaly Detection training data, which expects CSV/JSON files.

D . Object Storage Bucket: OCI Object Storage is a scalable, durable storage service that Anomaly Detection uses to ingest training data (e.g., CSV files).

Reasoning: Object Storage is the standard for large-scale data in OCI services, offering seamless integration with Anomaly Detection via APIs or SDKs.

Conclusion: D is the correct choice as it aligns with the service's architecture.

The OCI Anomaly Detection service requires training data to be uploaded to an Object Storage Bucket in formats like CSV or JSON. This is explicitly outlined in the official documentation, which states that users must "upload the training dataset to an OCI Object Storage bucket" before creating a data asset for model training. Local storage (A) isn't viable for cloud processing, and databases like MySQL (B) or Autonomous Data Warehouse (C) aren't supported as primary inputs. Object Storage (D) provides the scalability and accessibility needed for multivariate anomaly detection workflows.

: Oracle Cloud Infrastructure Anomaly Detection Documentation, "Preparing Training Data" section.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>