



1Z0-1110-22^{Q&As}

Oracle Cloud Infrastructure Data Science 2022 Professional

Pass Oracle 1Z0-1110-22 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-1110-22.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Which of the following TWO non-open source JupyterLab extensions has Oracle Cloud Infrastructure (OCI) Data Science developed and added to the notebook session experience?

- A. Environment Explorer
- B. Table of Contents
- C. Command Palette
- D. Notebook Examples
- E. Terminal

Correct Answer: AD

QUESTION 2

You are a data scientist with a set of text and image files that need annotation, and you want to use Oracle Cloud Infrastructure (OCI) Data Labeling. Which of the following THREE annotation classes are supported by the tool.?

- A. Object Detection
- B. Named Entity Extraction
- C. Classification (single/multi label)
- D. Key-Point and Landmark
- E. Polygonal Segmentation
- F. Semantic Segmentation

Correct Answer: ABC

QUESTION 3

You have built a machine model to predict whether a bank customer is going to default on a loan. You want to use Local Interpretable Model-Agnostic Explanations (LIME) to understand a specific prediction. What is the key idea behind LIME?

- A. Model-agnostic techniques are more interpretable than techniques that are dependent on the types of models.
- B. Local explanation techniques are model agnostic, while global explanation techniques are not.
- C. Global behavior of a machine learning model may be complex, while the local behavior may be approximated with a simpler surrogate model.
- D. Global and local behaviors of machine learning models are similar.



Correct Answer: C

QUESTION 4

You have an embarrassingly parallel or distributed batch job on a large amount of data running using Data Science Jobs

What would be the best approach to run the workload?

- A. Create the job in Data Science Jobs and then start the number of simultaneous job runs required for your workload.
- B. Create the job in Data Science Jobs and start a job run. When it is done, start a new job run until you achieve the number of runs required.
- C. Reconfigure the job run because Data science jobs does not support embarrassingly parallel.
- D. Create a new job for every job run that you have to run in parallel, because the Date Science Jobs service can have only one job run per job.

Correct Answer: A

QUESTION 5

Using Oracle AutoML, you are tuning hyperparameters on a supported model class and have specified a time budget. AutoML terminates computation once the time budget is exhausted. What would you expect AutoML to return in case the time budget is exhausted before hyperparameter tuning is completed?

- A. A hyperparameter configuration with a minimum learning rate is returned.
- B. The current best-known hyperparameter configuration is returned.
- C. A random hyperparameter configuration is returned.
- D. The last generated hyperparameter configuration is returned.

Correct Answer: B

[1Z0-1110-22 VCE Dumps](#)

[1Z0-1110-22 Study Guide](#)

[1Z0-1110-22 Braindumps](#)