



DP-100^{Q&As}

Designing and Implementing a Data Science Solution on Azure

Pass Microsoft DP-100 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/dp-100.html>

100% Passing Guarantee
100% Money Back Assurance

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

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



**QUESTION 1**

You deploy a real-time inference service for a trained model.

The deployed model supports a business-critical application, and it is important to be able to monitor the data submitted to the web service and the predictions the data generates.

You need to implement a monitoring solution for the deployed model using minimal administrative effort.

What should you do?

- A. View the explanations for the registered model in Azure ML studio.
- B. Enable Azure Application Insights for the service endpoint and view logged data in the Azure portal.
- C. View the log files generated by the experiment used to train the model.
- D. Create an ML Flow tracking URI that references the endpoint, and view the data logged by ML Flow.

Correct Answer: B

Configure logging with Azure Machine Learning studio

You can also enable Azure Application Insights from Azure Machine Learning studio. When you're ready to deploy your model as a web service, use the following steps to enable Application Insights:

1.

Sign in to the studio at <https://ml.azure.com>.

2.

Go to Models and select the model you want to deploy.

3.

Select +Deploy.

4.

Populate the Deploy model form.

5.

Expand the Advanced menu.

6.

Select Enable Application Insights diagnostics and data collection.



▼ Advanced

Enable Application Insights diagnostics and data collection

☒ Enable Application Insights diagnostics and data collection

Enable SSL

☐ Enable SSL

Max concurrent requests per container

CPU reserve capacity ⓘ

Memory reserve capacity ⓘ

Deploy

Cancel

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-enable-app-insights>

QUESTION 2

HOTSPOT

You are implementing hyperparameter tuning for a model training from a notebook. The notebook is in an Azure Machine Learning workspace. You add code that imports all relevant Python libraries.

You must configure Bayesian sampling over the search space for the `num_hidden_layers` and `batch_size` hyperparameters.

You need to complete the following Python code to configure Bayesian sampling.

Which code segments should you use? To answer, select the appropriate options in the answer area

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
param_sampling = BayesianParameterSampling( {  
    "learning_rate": uniform(0.05, 0.1),  
    "batch_size":  (  (16, 128, 16))  
})
```

A
loguniform
Placeholder

B
range
Placeholder

Correct Answer:

Answer Area

```
param_sampling = BayesianParameterSampling( {  
    "learning_rate": uniform(0.05, 0.1),  
    "batch_size":  (  (16, 128, 16))  
})
```

A
loguniform
Placeholder

B
range
Placeholder

QUESTION 3**DRAG DROP**

You have been tasked with evaluating the performance of a binary classification model that you created.

You need to choose evaluation metrics to achieve your goal.



Which of the following are the metrics you would choose? Answer by dragging the correct options from the list to the answer area.

Select and Place:

Options

Answer

Precision

Accuracy

Relative Squared
Error

Coefficient of
determination

Relative Absolute
Error

Correct Answer:



Options

Answer

Precision

Accuracy

Relative Squared
Error

Coefficient of
determination

Relative Absolute
Error

The evaluation metrics available for binary classification models are: Accuracy, Precision, Recall, F1 Score, and AUC.

Note: A very natural question is: 'Out of the individuals whom the model, how many were classified correctly (TP)?'

This question can be answered by looking at the Precision of the model, which is the proportion of positives that are classified correctly.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio/evaluate-model-performance>

QUESTION 4

You are a lead data scientist for a project that tracks the health and migration of birds. You create a multi- class image classification deep learning model that uses a set of labeled bird photographs collected by experts. You have 100,000



photographs of birds. All photographs use the JPG format and are stored in an Azure blob container in an Azure subscription.

You need to access the bird photograph files in the Azure blob container from the Azure Machine Learning service workspace that will be used for deep learning model training. You must minimize data movement.

What should you do?

- A. Create an Azure Data Lake store and move the bird photographs to the store.
- B. Create an Azure Cosmos DB database and attach the Azure Blob containing bird photographs storage to the database.
- C. Create and register a dataset by using TabularDataset class that references the Azure blob storage containing bird photographs.
- D. Register the Azure blob storage containing the bird photographs as a datastore in Azure Machine Learning service.
- E. Copy the bird photographs to the blob datastore that was created with your Azure Machine Learning service workspace.

Correct Answer: D

We recommend creating a datastore for an Azure Blob container. When you create a workspace, an Azure blob container and an Azure file share are automatically registered to the workspace.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-access-data>

QUESTION 5

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are creating a model to predict the price of a student's artwork depending on the following variables:

the student's length of education, degree type, and art form.

You start by creating a linear regression model.

You need to evaluate the linear regression model.

Solution: Use the following metrics: Relative Squared Error, Coefficient of Determination, Accuracy, Precision, Recall, F1 score, and AUC.

Does the solution meet the goal?

- A. Yes
- B. No



Correct Answer: B

Relative Squared Error, Coefficient of Determination are good metrics to evaluate the linear regression model, but the others are metrics for classification models.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/evaluate-model>

QUESTION 6

You make use of Azure Machine Learning Studio to develop a linear regression model. You perform an experiment to assess various algorithms. Which of the following is an algorithm that reduces the variances between actual and predicted values?

- A. Fast Forest Quantile Regression
- B. Poisson Regression
- C. Boosted Decision Tree Regression
- D. Linear Regression

Correct Answer: C

Mean absolute error (MAE) measures how close the predictions are to the actual outcomes; thus, a lower score is better.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/algorithm-module-reference/boosted-decision-tree-regression> <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/evaluate-model>
<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/linear-regression>

QUESTION 7

HOTSPOT

You are the owner of an Azure Machine Learning workspace.

You must prevent the creation or deletion of compute resources by using a custom role. You must allow all other operations inside the workspace.

You need to configure the custom role.

How should you complete the configuration? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
{
  "Name": "Data Scientist Custom",
  "IsCustom": true
  "Description": "Description"
  "Actions": [
    {
      "Microsoft.MachineLearningServices/workspaces/*/read"
      "Microsoft.MachineLearningServices/workspaces/computes/*/write"
      "Microsoft.MachineLearningServices/workspaces/delete"
    },
    {
      "Microsoft.MachineLearningServices/workspaces/*/write"
      "Microsoft.MachineLearningServices/workspaces/computes/*/write"
      "Microsoft.MachineLearningServices/workspaces/delete"
    }
  ],
  "NotActions": [
    {
      "Microsoft.MachineLearningServices/workspaces/*/read"
      "Microsoft.MachineLearningServices/workspaces/*/write"
      "Microsoft.MachineLearningServices/workspaces/computes/*/delete"
    },
    {
      "Microsoft.MachineLearningServices/workspaces/*/read"
      "Microsoft.MachineLearningServices/workspaces/*/write"
      "Microsoft.MachineLearningServices/workspaces/computes/*/write"
    }
  ],
  "AssignableScopes": [
    "/subscriptions/<subscription_id>"
  ]
}
```

Correct Answer:

**Answer Area**

```
{  
  "Name": "Data Scientist Custom",  
  "IsCustom": true  
  "Description": "Description"  
  "Actions": [  

```

	▼
Microsoft.MachineLearningServices/workspaces/*/read	
Microsoft.MachineLearningServices/workspaces/computes/*/write	
Microsoft.MachineLearningServices/workspaces/delete	

	▼
Microsoft.MachineLearningServices/workspaces/*/write	
Microsoft.MachineLearningServices/workspaces/computes/*/write	
Microsoft.MachineLearningServices/workspaces/delete	

```
],  
  "NotActions": [  

```

	▼
Microsoft.MachineLearningServices/workspaces/*/read	
Microsoft.MachineLearningServices/workspaces/*/write	
Microsoft.MachineLearningServices/workspaces/computes/*/delete	

	▼
Microsoft.MachineLearningServices/workspaces/*/read	
Microsoft.MachineLearningServices/workspaces/*/write	
Microsoft.MachineLearningServices/workspaces/computes/*/write	

```
],  
  "AssignableScopes": [  
    "/subscriptions/<subscription_id>"  
  ]  
}
```

Box 1: Microsoft.MachineLearningServices/workspaces/*/read

Reader role: Read-only actions in the workspace. Readers can list and view assets, including datastore credentials, in a workspace. Readers can't create or update these assets.

Box 2: Microsoft.MachineLearningServices/workspaces/*/write

If the roles include Actions that have a wildcard (*), the effective permissions are computed by subtracting the NotActions from the allowed Actions.

Box 3: Box 2: Microsoft.MachineLearningServices/workspaces/computes/*/delete Box 4:
Microsoft.MachineLearningServices/workspaces/computes/*/write Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview#how-azure-rbac-determines-if-a-user-has-access-to-a-resource>

QUESTION 8**HOTSPOT**



A biomedical research company plans to enroll people in an experimental medical treatment trial.

You create and train a binary classification model to support selection and admission of patients to the trial. The model includes the following features: Age, Gender, and Ethnicity.

The model returns different performance metrics for people from different ethnic groups.

You need to use Fairlearn to mitigate and minimize disparities for each category in the Ethnicity feature.

Which technique and constraint should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Option	Value
Technique	<div><div></div><div>Grid search</div><div>Outlier detection</div><div>Dimensionality reduction</div></div>
Constraint	<div><div></div><div>Demographic parity</div><div>False-positive rate parity</div></div>

Correct Answer:



Option	Value
Technique	<div><div>Grid search</div><div>Outlier detection</div><div>Dimensionality reduction</div></div>
Constraint	<div><div>Demographic parity</div><div>False-positive rate parity</div></div>

Box 1: Grid Search

Fairlearn open-source package provides postprocessing and reduction unfairness mitigation algorithms: ExponentiatedGradient, GridSearch, and ThresholdOptimizer.

Note: The Fairlearn open-source package provides postprocessing and reduction unfairness mitigation algorithms types:

1.
Reduction: These algorithms take a standard black-box machine learning estimator (e.g., a LightGBM model) and generate a set of retrained models using a sequence of re-weighted training datasets.

2.
Post-processing: These algorithms take an existing classifier and the sensitive feature as input.

Box 2: Demographic parity

The Fairlearn open-source package supports the following types of parity constraints: Demographic parity, Equalized odds, Equal opportunity, and Bounded group loss.

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/concept-fairness-ml>

QUESTION 9**DRAG DROP**

You use Azure Machine Learning to deploy a model as a real-time web service.



You need to create an entry script for the service that ensures that the model is loaded when the service starts and is used to score new data as it is received.

Which functions should you include in the script? To answer, drag the appropriate functions to the correct actions. Each function may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Answer Area

Functions	Action	Function
<code>main()</code>		
<code>score()</code>	Load the model when the service starts.	<input type="text"/>
<code>run()</code>	Use the model to score new data.	<input type="text"/>
<code>init()</code>		
<code>predict()</code>		

Correct Answer:



Answer Area

Functions	Action	Function
<code>main()</code>		
<code>score()</code>	Load the model when the service starts.	<code>init()</code>
	Use the model to score new data.	<code>run()</code>
<code>predict()</code>		

Box 1: `init()`

The entry script has only two required functions, `init()` and `run(data)`. These functions are used to initialize the service at startup and run the model using request data passed in by a client. The rest of the script handles loading and running the

model(s).

Box 2: `run()`

Reference:

<https://docs.microsoft.com/en-us/azure/machine-learning/how-to-deploy-existing-model>

QUESTION 10

HOTSPOT

Complete the sentence by selecting the correct option in the answer area.

Hot Area:



Answer Area

	▼
Probabilistic PCA	
Median	
SMOTE	
Custom substitution value	

is a data cleaning option of the Clean Missing Data module that does not require predictors for each column.

Correct Answer:

Answer Area

	▼
Probabilistic PCA	
Median	
SMOTE	
Custom substitution value	

is a data cleaning option of the Clean Missing Data module that does not require predictors for each column.

Replace using Probabilistic PCA: Compared to other options, such as Multiple Imputation using Chained Equations (MICE), this option has the advantage of not requiring the application of predictors for each column. Instead, it approximates the covariance for the full dataset. Therefore, it might offer better performance for datasets that have missing values in many columns.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/clean-missing-data>

QUESTION 11

You manage an Azure Machine Learning workspace.

You must provide explanations for the behavior of the models with feature importance measures.

You need to configure a Responsible AI dashboard in Azure Machine Learning.

Which dashboard component should you configure?



- A. Counterfactual what-if
- B. Casual inference
- C. Fairness assessment
- D. Interpretability

Correct Answer: D

Use interpretability when you need to:

Determine how trustworthy your AI system's predictions are by understanding what features are most important for the predictions. Approach the debugging of your model by understanding it first and identifying whether the model is using

healthy features or merely false correlations. Uncover potential sources of unfairness by understanding whether the model is basing predictions on sensitive features or on features that are highly correlated with them.

Etc.

Note: Responsible AI dashboard components

The Responsible AI dashboard brings together, in a comprehensive view, various new and pre-existing tools. The dashboard integrates these tools with Azure Machine Learning CLI v2, Azure Machine Learning Python SDK v2, and Azure

Machine Learning studio. The tools include:

*

Model interpretability (importance values for aggregate and individual features), to understand your model's predictions and how those overall and individual predictions are made.

*

Counterfactual what-if, to observe how feature perturbations would affect your model predictions while providing the closest data points with opposing or different model predictions.

*

Causal analysis, to use historical data to view the causal effects of treatment features on real-world outcomes.

*

Etc.

Reference: <https://learn.microsoft.com/en-us/azure/machine-learning/how-to-machine-learning-interpretability>
<https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai-dashboard>

QUESTION 12

You create an Azure Machine Learning workspace. You are implementing hyperparameter tuning for a model training from a notebook.

You must configure a Bandit termination policy that provides the following outcome:



If the value of the primary metric of AUC is 0.8 at the point of evaluation intervals, any run with the primary metric value below 0.66 will be terminated.

You need to identify which Bandit termination policy configuration to use.

What should you identify?

- A. Set slack_amount to 0.2.
- B. Set slack_factor to 0.1.
- C. Set slack_factor to 0.2.
- D. Set slack_amount to 0.1.

Correct Answer: C

Bandit policy

Bandit policy is based on slack factor/slack amount and evaluation interval. Bandit policy ends a job when the primary metric isn't within the specified slack factor/slack amount of the most successful job.

Specify the following configuration parameters:

slack_factor or slack_amount: the slack allowed with respect to the best performing training job. slack_factor specifies the allowable slack as a ratio. slack_amount specifies the allowable slack as an absolute amount, instead of a ratio.

For example, consider a Bandit policy applied at interval 10. Assume that the best performing job at interval 10 reported a primary metric is 0.8 with a goal to maximize the primary metric. If the policy specifies a slack_factor of 0.2, any training

jobs whose best metric at interval 10 is less than 0.66 ($0.8/(1+\text{slack_factor})$) will be terminated.

Reference: <https://learn.microsoft.com/en-us/azure/machine-learning/how-to-tune-hyperparameters>

QUESTION 13

You are solving a classification task.

The dataset is imbalanced.

You need to select an Azure Machine Learning Studio module to improve the classification accuracy.

Which module should you use?

- A. Permutation Feature Importance
- B. Filter Based Feature Selection
- C. Fisher Linear Discriminant Analysis
- D. Synthetic Minority Oversampling Technique (SMOTE)

Correct Answer: D



Use the SMOTE module in Azure Machine Learning Studio (classic) to increase the number of underrepresented cases in a dataset used for machine learning. SMOTE is a better way of increasing the number of rare cases than simply duplicating existing cases.

You connect the SMOTE module to a dataset that is imbalanced. There are many reasons why a dataset might be imbalanced: the category you are targeting might be very rare in the population, or the data might simply be difficult to collect. Typically, you use SMOTE when the class you want to analyze is under- represented.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/smote>

QUESTION 14

You create a multi-class image classification deep learning model.

You train the model by using PyTorch version 1.2.

You need to ensure that the correct version of PyTorch can be identified for the inferencing environment when the model is deployed.

What should you do?

- A. Save the model locally as a .pt file, and deploy the model as a local web service.
- B. Deploy the model on computer that is configured to use the default Azure Machine Learning conda environment.
- C. Register the model with a .pt file extension and the default version property.
- D. Register the model, specifying the model_framework and model_framework_version properties.

Correct Answer: D

framework_version: The PyTorch version to be used for executing training code.

Reference: <https://docs.microsoft.com/en-us/python/api/azureml-train-core/azureml.train.dnn.pytorch?view=azure-ml-py>

QUESTION 15

You train and register an Azure Machine Learning model

You plan to deploy the model to an online endpoint

You need to ensure that applications will be able to use the authentication method with a non-expiring artifact to access the model.

Solution: Create a managed online endpoint with the default authentication settings. Deploy the model to the online endpoint.

Does the solution meet the goal?

- A. Yes
- B. No



Correct Answer: B

[Latest DP-100 Dumps](#)

[DP-100 Exam Questions](#)

[DP-100 Braindumps](#)