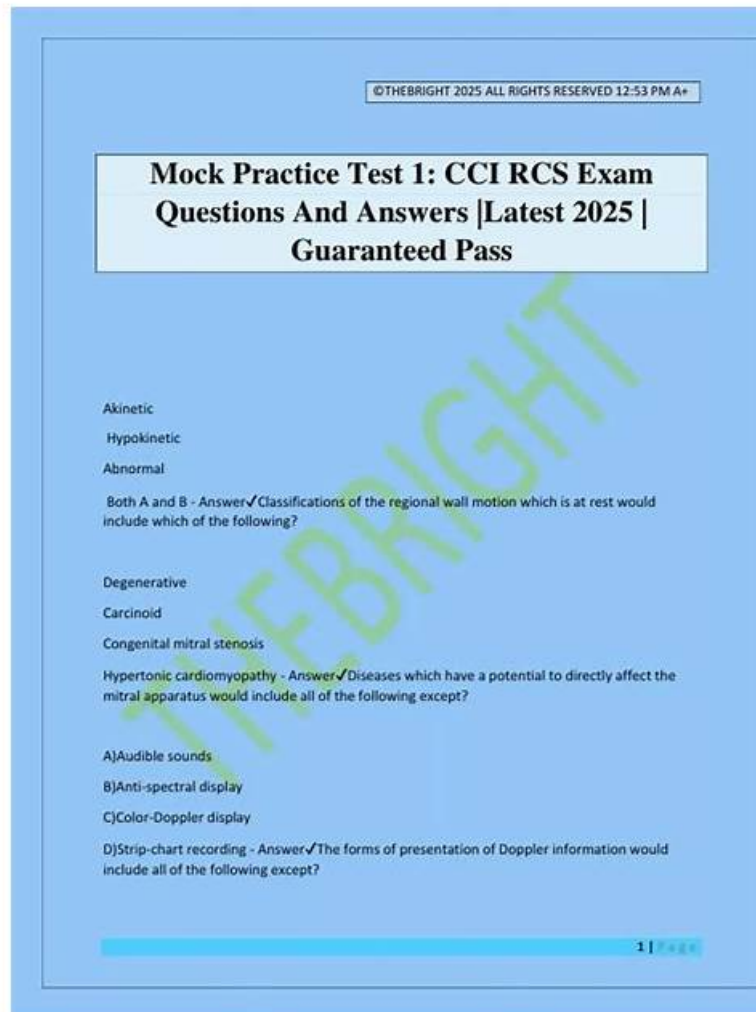


# AI-300 Practice Mock | Pass Guaranteed | Refund Guaranteed



If you would like to use all kinds of electronic devices to prepare for the AI-300 exam, with the online app version of our AI-300 study materials, you can just feel free to practice the questions in our AI-300 training materials no matter you are using your mobile phone, personal computer, or tablet PC. In addition, another strong point of the online app version is that it is convenient for you to use even though you are in offline environment. In other words, you can prepare for your AI-300 Exam with under the guidance of our AI-300 training materials anywhere at any time.

In order to meet different needs for AI-300 exam bootcamp, three versions are available. You can choose the most suitable one according to your own exam needs. All three have free demo for you to have a try before buying. AI-300 PDF version is printable, you can study them anytime. AI-300 Soft test engine supports MS operating system, and have two modes for practice, and it can also stimulate the real exam environment, therefore, this version can build you exam confidence. AI-300 Online test engine is convenient to learn, and it also supports offline practice.

>> AI-300 Practice Mock <<

## AI-300 Test Sample Online & AI-300 Valid Test Sample

The Operationalizing Machine Learning and Generative AI Solutions (AI-300) certification examination is an essential component of professional development, and passing this Microsoft AI-300 test can increase career options and a rise in salary. Nonetheless, getting ready for the Prepare for your AI-300 Exam may be difficult, and many working professionals have trouble locating the AI-

300 practice questions they need to succeed in this endeavor.

## Microsoft Operationalizing Machine Learning and Generative AI Solutions Sample Questions (Q10-Q15):

### NEW QUESTION # 10

You create a binary classification model. You use the Fairlearn package to assess model fairness.

You must eliminate the need to retrain the model.

You need to implement the Fairlearn package.

Which algorithm should you use?

- A. `fairlearn.preprocessing.CorrelationRemover`
- B. `fairlearn.reductions.ExponentiatedGradient`
- C. `fairlearn.reductions.GridSearch`
- D. `fairlearn.postprocessing.ThresholdOptimizer`

**Answer: D**

Explanation:

The best algorithm to implement within the Fairlearn package for a binary classification model without retraining is the `ThresholdOptimizer`.

Why `ThresholdOptimizer`?

This algorithm is a post-processing technique specifically designed to mitigate unfairness after a model has already been trained. By setting the `prefit` parameter to `True`, it treats your existing model as a "black box," adjusting its decision thresholds for different groups to satisfy fairness constraints (like demographic parity or equalized odds) without changing the underlying model weights.

Reference:

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

### NEW QUESTION # 11

Hotspot Question

You are monitoring a fine-tuned large language model deployed in Microsoft Foundry.

You evaluate the model before and after fine-tuning by using the same evaluation dataset.

You review the following evaluation results:

You need to determine whether the fine-tuned model shows improved performance without introducing regression. For each of the following statements, select Yes if the statement is true.

Otherwise, select No.

NOTE: Each correct selection is worth one point.

**Answer:**

Explanation:

### NEW QUESTION # 12

Hotspot Question

You manage an Azure Machine Learning workspace named `workspace1` by using the Python SDK v2.

The default datastore of `workspace1` contains a folder named `sample_data`. The folder structure contains the following content:

You write Python SDK v2 code to materialize the data from the files in the `sample_data` folder into a Pandas data frame.

You need to complete the Python SDK v2 code to use the `MLTable` folder as the materialization blueprint.

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

NOTE: Each correct selection is worth one point

**Answer:**

Explanation:

### NEW QUESTION # 13

An organization validates generative AI applications during CI/CD Microsoft Foundry. Evaluation must run automatically and block releases when quality thresholds are NOT met. Manual evaluation is no longer acceptable. Evaluation must use both predefined quality metrics and custom safety checks. You need to implement an automated evaluation workflow that supports both built-in and custom metrics. What should you do?

- A. Enable application tracing to collect runtime telemetry.
- B. Review evaluation results manually after deployment.
- C. Monitor latency metrics during model inference.
- **D. Implement an evaluation step by using GitHub Actions.**

**Answer: D**

Explanation:

To implement an automated evaluation step in GitHub Actions for Microsoft Foundry AI, you can use the Microsoft Foundry Evaluation GitHub Action (or the Azure AI Evaluation SDK).

This setup allows you to run both built-in metrics (like groundedness or coherence) and custom safety checks, then fail the build if scores fall below your defined thresholds.

Implementation Steps

#### 1. Define Your Evaluators

You need to configure which metrics to use. Microsoft Foundry supports two main types:

Built-in Metrics: Pre-trained models that score quality (coherence, fluency) and safety (hate, violence, self-harm).

Custom Metrics: Python-based evaluators you define to check domain-specific requirements.

#### 2. Configure the GitHub Actions Workflow

Create a .yml file in your .github/workflows directory. This workflow will:

Trigger on a pull request or commit.

Authenticate with Azure/Foundry.

Run Evaluation using the microsoft/ai-agent-evals action.

Enforce Thresholds to block the release if quality is insufficient.

Key Components for "Block Release" Logic

To ensure the release is blocked, your workflow must include a gating step that interprets the evaluation results:

Reference:

<https://learn.microsoft.com/en-us/training/modules/automated-evaluation-genaiops>

### NEW QUESTION # 14

A team schedules weekly retraining of a model using Azure ML pipelines. They also want retraining triggered automatically when production data significantly deviates from training data distribution, without duplicating pipeline logic. What should they implement?

- A. Azure Function to retrain model manually
- B. Two independent pipelines with shared scripts
- C. Notebook-based retraining process
- **D. One pipeline triggered by schedule and data drift alerts**

**Answer: D**

Explanation:

Using a single pipeline triggered by both a schedule and data drift alerts ensures consistent retraining logic and avoids duplication.

This approach minimizes operational overhead and maintenance complexity. Creating multiple pipelines can lead to inconsistencies, duplicated code, and increased effort when updating retraining logic or dependencies.

### NEW QUESTION # 15

.....

I am proud to tell you that our company is definitely one of the most authoritative companies in the international market for AI-300 exam. What's more, we will provide the most considerate after sale service for our customers in twenty four hours a day seven days a week, therefore, our company is really the best choice for you to buy the AI-300 Training Materials. You can just feel rest assured that our after sale service staffs are always here waiting for offering you our services on our AI-300 exam questions. Please feel free to contact us. You will be surprised by our good AI-300 study guide.

