Top Exam MLS-C01 Tutorials | Efficient Amazon MLS-C01: AWS Certified Machine Learning - Specialty 100% Pass

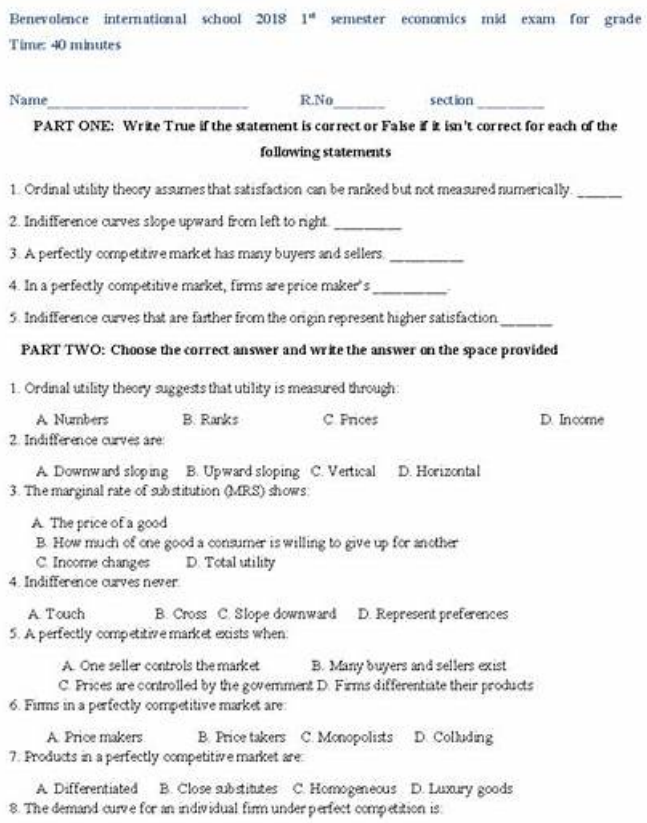

Benevolence international school 2018 1st semester economics mid exam for grade
Time: 40 minutes

Name_____ R.No_____ section _____

PART ONE: Write True if the statement is correct or False if it isn't correct for each of the following statements

1. Ordinal utility theory assumes that satisfaction can be ranked but not measured numerically. _____
2. Indifference curves slope upward from left to right. _____
3. A perfectly competitive market has many buyers and sellers. _____
4. In a perfectly competitive market, firms are price maker's _____
5. Indifference curves that are farther from the origin represent higher satisfaction. _____

PART TWO: Choose the correct answer and write the answer on the space provided

1. Ordinal utility theory suggests that utility is measured through:
   A. Numbers B. Ranks C. Prices D. Income
2. Indifference curves are:
   A. Downward sloping B. Upward sloping C. Vertical D. Horizontal
3. The marginal rate of substitution (MRS) shows:
   A. The price of a good
   B. How much of one good a consumer is willing to give up for another
   C. Income changes D. Total utility
4. Indifference curves never:
   A. Touch B. Cross C. Slope downward D. Represent preferences
5. A perfectly competitive market exists when:
   A. One seller controls the market B. Many buyers and sellers exist
   C. Prices are controlled by the government D. Firms differentiate their products
6. Firms in a perfectly competitive market are:
   A. Price makers B. Price takers C. Monopolists D. Colluding
7. Products in a perfectly competitive market are:
   A. Differentiated B. Close substitutes C. Homogeneous D. Luxury goods
8. The demand curve for an individual firm under perfect competition is:

P.S. Free & New MLS-C01 dumps are available on Google Drive shared by Pass4sureCert: https://drive.google.com/open?id=1J-48OIW4stwVi_hIZAjAi8X0zobQ8Fhc

The Pass4sureCert is one of the top-rated and leading platforms that have been offering a simple, smart, and easiest way to pass the challenging MLS-C01 exam with good scores. The Amazon MLS-C01 Exam Questions are real, valid, and updated. These MLS-C01 exam practice questions are designed and verified by experienced and qualified MLS-C01 exam experts.

Achieving the AWS Certified Machine Learning - Specialty certification can help individuals advance their careers in the field of machine learning and increase their earning potential. AWS Certified Machine Learning - Specialty certification is recognized by industry leaders and can open up new opportunities for professionals in various industries, including healthcare, finance, and retail, among others.

>> Exam MLS-C01 Tutorials <<

Valid Dumps MLS-C01 Ebook & MLS-C01 Valid Test Testking

It is a popular belief that only processional experts can be the leading one to do some adept job. And similarly, only high quality and high accuracy MLS-C01 Exam Questions like ours can give you confidence and reliable backup to get the certificate smoothly because our experts have extracted the most frequent-tested points for your reference. Good practice materials like our AWS

Certified Machine Learning - Specialty study question can educate exam candidates with the most knowledge. Do not make your decisions now will be a pity for good.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q63-Q68):

NEW QUESTION # 63
A credit card company wants to build a credit scoring model to help predict whether a new credit card applicant will default on a credit card payment. The company has collected data from a large number of sources with thousands of raw attributes. Early experiments to train a classification model revealed that many attributes are highly correlated, the large number of features slows down the training speed significantly, and that there are some overfitting issues.
The Data Scientist on this project would like to speed up the model training time without losing a lot of information from the original dataset.
Which feature engineering technique should the Data Scientist use to meet the objectives?

- A. Run self-correlation on all features and remove highly correlated features
- B. Normalize all numerical values to be between 0 and 1
- C. Use an autoencoder or principal component analysis (PCA) to replace original features with new features
- D. Cluster raw data using k-means and use sample data from each cluster to build a new dataset

Answer: C

Explanation:
The best feature engineering technique to speed up the model training time without losing a lot of information from the original dataset is to use an autoencoder or principal component analysis (PCA) to replace original features with new features. An autoencoder is a type of neural network that learns a compressed representation of the input data, called the latent space, by minimizing the reconstruction error between the input and the output. PCA is a statistical technique that reduces the dimensionality of the data by finding a set of orthogonal axes, called the principal components, that capture the maximum variance of the data. Both techniques can help reduce the number of features and remove the noise and redundancy in the data, which can improve the model performance and speed up the training process. References:
* AWS Machine Learning Specialty Exam Guide
* AWS Machine Learning Training - Dimensionality Reduction for Machine Learning
* AWS Machine Learning Training - Deep Learning with Amazon SageMaker

NEW QUESTION # 64
A data engineer is preparing a dataset that a retail company will use to predict the number of visitors to stores.
The data engineer created an Amazon S3 bucket. The engineer subscribed the S3 bucket to an AWS Data Exchange data product for general economic indicators. The data engineer wants to join the economic indicator data to an existing table in Amazon Athena to merge with the business data. All these transformations must finish running in 30-60 minutes.
Which solution will meet these requirements MOST cost-effectively?

- A. Provision an Amazon Redshift cluster. Subscribe to the AWS Data Exchange product and use the product to create an Amazon Redshift Table Merge the data in Amazon Redshift. Write the results back to Amazon S3.
- B. Configure the AWS Data Exchange product as a producer for an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to transfer the data to Amazon S3 Run an AWS Glue job that will merge the existing business data with the Athena table. Write the result set back to Amazon S3.
- C. Use an S3 event on the AWS Data Exchange S3 bucket to invoke an AWS Lambda Function Program the Lambda function to run an AWS Glue job that will merge the existing business data with the Athena table Write the results back to Amazon S3.
- D. Use an S3 event on the AWS Data Exchange S3 bucket to invoke an AWS Lambda function. Program the Lambda function to use Amazon SageMaker Data Wrangler to merge the existing business data with the Athena table. Write the result set back to Amazon S3.

Answer: D

Explanation:
Explanation
The most cost-effective solution is to use an S3 event to trigger a Lambda function that uses SageMaker Data Wrangler to merge the data. This solution avoids the need to provision and manage any additional resources, such as Kinesis streams, Firehose delivery streams, Glue jobs, or Redshift clusters. SageMaker Data Wrangler provides a visual interface to import, prepare, transform, and

analyze data from various sources, including AWS Data Exchange products. It can also export the data preparation workflow to a Python script that can be executed by a Lambda function. This solution can meet the time requirement of 30-60 minutes, depending on the size and complexity of the data.
References:
Using Amazon S3 Event Notifications
Prepare ML Data with Amazon SageMaker Data Wrangler
AWS Lambda Function

NEW QUESTION # 65
A Machine Learning Specialist is packaging a custom ResNet model into a Docker container so the company can leverage Amazon SageMaker for training. The Specialist is using Amazon EC2 P3 instances to train the model and needs to properly configure the Docker container to leverage the NVIDIA GPUs.
What does the Specialist need to do?

- A. Build the Docker container to be NVIDIA-Docker compatible.
- B. Set the GPU flag in the Amazon SageMaker CreateTrainingJob request body
- C. Organize the Docker container's file structure to execute on GPU instances.
- D. Bundle the NVIDIA drivers with the Docker image.

Answer: A

Explanation:
To leverage the NVIDIA GPUs on Amazon EC2 P3 instances for training a custom ResNet model using Amazon SageMaker, the Machine Learning Specialist needs to build the Docker container to be NVIDIA- Docker compatible. NVIDIA-Docker is a tool that enables GPU-accelerated containers to run on Docker.
NVIDIA-Docker can automatically configure the Docker container with the necessary drivers, libraries, and environment variables to access the NVIDIA GPUs. NVIDIA-Docker can also isolate the GPU resources and ensure that each container has exclusive access to a GPU.
To build a Docker container that is NVIDIA-Docker compatible, the Machine Learning Specialist needs to follow these steps:
* Install the NVIDIA Container Toolkit on the host machine that runs Docker. This toolkit includes the NVIDIA Container Runtime, which is a modified version of the Docker runtime that supports GPU hardware.
* Use the base image provided by NVIDIA as the first line of the Dockerfile. The base image contains the NVIDIA drivers and CUDA toolkit that are required for GPU-accelerated applications. The base image can be specified as FROM nvcr.io/nvidia/cuda:tag, where tag is the version of CUDA and the operating system.
* Install the required dependencies and frameworks for the ResNet model, such as PyTorch, torchvision, etc., in the Dockerfile.
* Copy the ResNet model code and any other necessary files to the Docker container in the Dockerfile.
* Build the Docker image using the docker build command.
* Push the Docker image to a repository, such as Amazon Elastic Container Registry (Amazon ECR), using the docker push command.
* Specify the Docker image URI and the instance type (ml.p3.xlarge) in the Amazon SageMaker CreateTrainingJob request body.
The other options are not valid or sufficient for building a Docker container that can leverage the NVIDIA GPUs on Amazon EC2 P3 instances. Bundling the NVIDIA drivers with the Docker image is not a good option, as it can cause driver conflicts and compatibility issues with the host machine and the NVIDIA GPUs.
Organizing the Docker container's file structure to execute on GPU instances is not a good option, as it does not ensure that the Docker container can access the NVIDIA GPUs and the CUDA toolkit. Setting the GPU flag in the Amazon SageMaker CreateTrainingJob request body is not a good option, as it does not apply to custom Docker containers, but only to built-in algorithms and frameworks that support GPU instances.

NEW QUESTION # 66
A retail chain has been ingesting purchasing records from its network of 20,000 stores to Amazon S3 using Amazon Kinesis Data Firehose. To support training an improved machine learning model, training records will require new but simple transformations, and some attributes will be combined. The model needs to be retrained daily.
Given the large number of stores and the legacy data ingestion, which change will require the LEAST amount of development effort?

- A. Insert an Amazon Kinesis Data Analytics stream downstream of the Kinesis Data Firehose stream that transforms raw record attributes into simple transformed values using SQL.
- B. Spin up a fleet of Amazon EC2 instances with the transformation logic, have them transform the data records accumulating on Amazon S3, and output the transformed records to Amazon S3.

- C. Deploy an Amazon EMR cluster running Apache Spark with the transformation logic, and have the cluster run each day on the accumulating records in Amazon S3, outputting new/transformed records to Amazon S3.
- D. Require that the stores to switch to capturing their data locally on AWS Storage Gateway for loading into Amazon S3, then use AWS Glue to do the transformation.

Answer: A

Explanation:
Explanation

NEW QUESTION # 67
A Machine Learning Specialist wants to bring a custom algorithm to Amazon SageMaker. The Specialist implements the algorithm in a Docker container supported by Amazon SageMaker.
How should the Specialist package the Docker container so that Amazon SageMaker can launch the training correctly?

- A. Use CMD configin the Dockerfile to add the training program as a CMD of the image
- B. Configure the training program as an ENTRYPOINTnamed train
- C. Modify the bash_profile file in the container and add a bashcommand to start the training program
- D. Copy the training program to directory /opt/ml/train

Answer: A

NEW QUESTION # 68
.....

In order to cater to different kinds of needs of customers, three versions for MLS-C01 learning materials are available. You can choose one you prefer according to your own needs. MLS-C01 PDF version is printable and you can study anywhere and anyplace. MLS-C01 Soft test engine supports MS operating system and have two modes for practice. In addition, MLS-C01 Soft test engine can simulate the real exam environment, and your confidence for the exam can be strengthened through this version. MLS-C01 Online test engine is convenient and easy to study, it supports all web browsers, and it has testing history and performance review, so that you can have a general review before next training.

- Latest MLS-C01 Study Notes □ MLS-C01 Valid Exam Pass4sure □ MLS-C01 Valid Test Cost □ Search for ➡ MLS-C01 □□□ on ▹ www.easy4engine.com ◃ immediately to obtain a free download ֍MLS-C01 Test Preparation
- lailatuanday.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, drnesmaelsersawy.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, smartmaths.com.ng, Disposable vapes

P.S. Free 2026 Amazon MLS-C01 dumps are available on Google Drive shared by Pass4sureCert: https://drive.google.com/open?id=1J-48OIW4stwVi_hIZAjAi8X0zobQ8Fhc