AWS-Certified-Machine-Learning-Specialty New Dumps Questions & Advanced AWS-Certified-Machine-Learning-Specialty Testing Engine

What's more, part of that BootcampPDF AWS-Certified-Machine-Learning-Specialty dumps now are free: https://drive.google.com/open?id=1eSr2wwOM8Vt5vs4b65TQwD7RDyRBOce8

In order to help you save more time, we will transfer AWS-Certified-Machine-Learning-Specialty test guide to you within 10 minutes online after your payment and guarantee that you can study these AWS-Certified-Machine-Learning-Specialty training materials as soon as possible to avoid time waste. We believe that time is the most valuable things in the world. This is why we are dedicated to improve your study efficiency and production. Here are some advantages of our AWS-Certified-Machine-Learning-Specialty study question and we would appreciate that you can have a look to our AWS-Certified-Machine-Learning-Specialty questions.

The AWS Certified Machine Learning - Specialty certification exam consists of 65 multiple-choice and multiple-answer questions that must be completed within 180 minutes. AWS-Certified-Machine-Learning-Specialty Exam is available in English, Japanese, Korean, and Simplified Chinese. The cost of the exam is $300.

>> AWS-Certified-Machine-Learning-Specialty New Dumps Questions <<

Amazon Offers Amazon AWS-Certified-Machine-Learning-Specialty Dumps with Refund Guaranty

Customers can start using the Amazon AWS-Certified-Machine-Learning-Specialty Exam Questions instantly just after purchasing it from our website for the preparation of the AWS-Certified-Machine-Learning-Specialty certification exam. They can also evaluate the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice test material before buying with a free demo. The users will receive updates 365 days after purchasing. And they will also get a 24/7 support system to help them anytime if they got stuck somewhere or face any issues while preparing for the AWS-Certified-Machine-Learning-Specialty Exam.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q246-Q251):

NEW QUESTION # 246
A medical imaging company wants to train a computer vision model to detect areas of concern on patients' CT scans. The company has a large collection of unlabeled CT scans that are linked to each patient and stored in an Amazon S3 bucket. The scans must be accessible to authorized users only. A machine learning engineer needs to build a labeling pipeline.
Which set of steps should the engineer take to build the labeling pipeline with the LEAST effort?

- A. Create a private workforce and manifest file. Create a labeling job by using the built-in bounding box task type in Amazon SageMaker Ground Truth. Write the labeling instructions.
- B. Create an Amazon Mechanical Turk workforce and manifest file. Create a labeling job by using the built-in image classification task type in Amazon SageMaker Ground Truth. Write the labeling instructions.

- C. Create a workforce with AWS Identity and Access Management (IAM). Build a labeling tool on Amazon EC2 Queue images for labeling by using Amazon Simple Queue Service (Amazon SQS). Write the labeling instructions.
- D. Create a workforce with Amazon Cognito. Build a labeling web application with AWS Amplify. Build a labeling workflow backend using AWS Lambda. Write the labeling instructions.

Answer: A

Explanation:
The engineer should create a private workforce and manifest file, and then create a labeling job by using the built-in bounding box task type in Amazon SageMaker Ground Truth. This will allow the engineer to build the labeling pipeline with the least effort.
A private workforce is a group of workers that you manage and who have access to your labeling tasks. You can use a private workforce to label sensitive data that requires confidentiality, such as medical images. You can create a private workforce by using Amazon Cognito and inviting workers by email. You can also use AWS Single Sign-On or your own authentication system to manage your private workforce.
A manifest file is a JSON file that lists the Amazon S3 locations of your input data. You can use a manifest file to specify the data objects that you want to label in your labeling job. You can create a manifest file by using the AWS CLI, the AWS SDK, or the Amazon SageMaker console.
A labeling job is a process that sends your input data to workers for labeling. You can use the Amazon SageMaker console to create a labeling job and choose from several built-in task types, such as image classification, text classification, semantic segmentation, and bounding box. A bounding box task type allows workers to draw boxes around objects in an image and assign labels to them. This is suitable for object detection tasks, such as identifying areas of concern on CT scans.
References:
Create and Manage Workforces - Amazon SageMaker
Use Input and Output Data - Amazon SageMaker
Create a Labeling Job - Amazon SageMaker
Bounding Box Task Type - Amazon SageMaker

NEW QUESTION # 247
A retail company collects customer comments about its products from social media, the company website, and customer call logs. A team of data scientists and engineers wants to find common topics and determine which products the customers are referring to in their comments. The team is using natural language processing (NLP) to build a model to help with this classification.
Each product can be classified into multiple categories that the company defines. These categories are related but are not mutually exclusive. For example, if there is mention of "Sample Yogurt" in the document of customer comments, then "Sample Yogurt" should be classified as "yogurt," "snack," and "dairy product." The team is using Amazon Comprehend to train the model and must complete the project as soon as possible.
Which functionality of Amazon Comprehend should the team use to meet these requirements?

- A. Custom entity recognition
- B. Custom classification with multi-label mode
- C. Built-in models
- D. Custom classification with multi-class mode

Answer: B

Explanation:
The problem involves assigning multiple related, non-mutually exclusive categories (like "yogurt", "snack", and "dairy product") to customer comments referencing specific products. Amazon Comprehend provides custom classification capabilities, and in this case, the appropriate mode is multi-label, not multi-class.
"Use multi-label mode when each document can belong to more than one class. This is appropriate for situations where the categories are not mutually exclusive." This allows a single document to be tagged with several relevant labels simultaneously, which is precisely the use case described.

NEW QUESTION # 248
A car company is developing a machine learning solution to detect whether a car is present in an image. The image dataset consists of one million images. Each image in the dataset is 200 pixels in height by 200 pixels in width. Each image is labeled as either having a car or not having a car.
Which architecture is MOST likely to produce a model that detects whether a car is present in an image with the highest accuracy?

- A. Use a deep convolutional neural network (CNN) classifier with the images as input. Include a linear output layer that

outputs the probability that an image contains a car.

- B. Use a deep multilayer perceptron (MLP) classifier with the images as input. Include a linear output layer that outputs the probability that an image contains a car.
- C. Use a deep multilayer perceptron (MLP) classifier with the images as input. Include a softmax output layer that outputs the probability that an image contains a car.
- D. Use a deep convolutional neural network (CNN) classifier with the images as input. Include a softmax output layer that outputs the probability that an image contains a car.

Answer: A

Explanation:
Explanation
A deep convolutional neural network (CNN) classifier is a suitable architecture for image classification tasks, as it can learn features from the images and reduce the dimensionality of the input. A linear output layer that outputs the probability that an image contains a car is appropriate for a binary classification problem, as it can produce a single scalar value between 0 and 1. A softmax output layer is more suitable for a multi-class classification problem, as it can produce a vector of probabilities that sum up to 1. A deep multilayer perceptron (MLP) classifier is not as effective as a CNN for image classification, as it does not exploit the spatial structure of the images and requires a large number of parameters to process the high-dimensional input. References:
AWS Certified Machine Learning - Specialty Exam Guide
AWS Training - Machine Learning on AWS
AWS Whitepaper - An Overview of Machine Learning on AWS

NEW QUESTION # 249
A Machine Learning Specialist works for a credit card processing company and needs to predict which transactions may be fraudulent in near-real time. Specifically, the Specialist must train a model that returns the probability that a given transaction may be fraudulent How should the Specialist frame this business problem'?

- A. Regression classification
- B. Streaming classification
- C. Binary classification
- D. Multi-category classification

Answer: C

Explanation:
Binary classification is a type of supervised learning problem where the goal is to predict a categorical label that has only two possible values, such as Yes or No, True or False, Positive or Negative. In this case, the label is whether a transaction is fraudulent or not, which is a binary outcome. Binary classification can be used to estimate the probability of an observation belonging to a certain class, such as the probability of a transaction being fraudulent. This can help the business to make decisions based on the risk level of each transaction. References:
Binary Classification - Amazon Machine Learning
AWS Certified Machine Learning - Specialty Sample Questions

NEW QUESTION # 250
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 do1?

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

Answer: B

Explanation:
To leverage the NVIDIA GPUs on Amazon EC2 P3 instances, 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. It automatically configures the container to access the NVIDIA drivers and libraries on the host system. The Specialist does not need

to bundle the NVIDIA drivers with the Docker image, as they are already installed on the EC2 P3 instances. The Specialist does not need to organize the Docker container's file structure to execute on GPU instances, as this is not relevant for GPU compatibility. The Specialist does not need to set the GPU flag in the Amazon SageMaker Create TrainingJob request body, as this is only required for using Elastic Inference accelerators, not EC2 P3 instances.
References: NVIDIA-Docker, Using GPU-Accelerated Containers, Using Elastic Inference in Amazon SageMaker

NEW QUESTION # 251
.....

No matter you are exam candidates of high caliber or newbies, our Amazon AWS-Certified-Machine-Learning-Specialty exam quiz will be your propulsion to gain the best results with least time and reasonable money. Not only because the outstanding content of AWS-Certified-Machine-Learning-Specialty Real Dumps that produced by our professional expert but also for the reason that we have excellent vocational moral to improve our AWS-Certified-Machine-Learning-Specialty learning materials quality.

2025 Latest BootcampPDF AWS-Certified-Machine-Learning-Specialty PDF Dumps and AWS-Certified-Machine-Learning-Specialty Exam Engine Free Share: https://drive.google.com/open?id=1eSr2wwOM8Vt5vs4b65TQwD7RDyRBOce8