AWS-Certified-Machine-Learning-Specialty Latest Cram Materials | AWS-Certified-Machine-Learning-Specialty New Test Bootcamp

BONUS!!! Download part of Easy4Engine AWS-Certified-Machine-Learning-Specialty dumps for free: https://drive.google.com/open?id=15ckdvKl15eJoZcKpBDDXdUFysqz6R7nW

Do you want to obtain your AWS-Certified-Machine-Learning-Specialty study materials as quickly as possible? If you do, then we will be your best choice. You can receive downloading link and password with ten minutes after buying. In addition, AWS-Certified-Machine-Learning-Specialty exam dumps are high quality, because we have experienced experts to edit, and you can pass your exam by using AWS-Certified-Machine-Learning-Specialty Exam Materials of us. In addition, we are pass guarantee and money back guarantee, if you fail to pass the exam by using AWS-Certified-Machine-Learning-Specialty study materials of us, we will give you full refund. And the money will be returned to your payment account.

Earning the AWS Certified Machine Learning - Specialty certification demonstrates to employers and colleagues that you have the skills and knowledge needed to design and deploy machine learning models on the AWS platform. It can help you stand out in a competitive job market and increase your earning potential.

>> AWS-Certified-Machine-Learning-Specialty Latest Cram Materials <<

AWS-Certified-Machine-Learning-Specialty New Test Bootcamp, Reliable AWS-Certified-Machine-Learning-Specialty Exam Guide

The AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty certification offers a great opportunity for beginners and professionals to demonstrate their skills and abilities to perform a certain task. For the complete, comprehensive, for AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty Exam Preparation you can get assistance from AWS Certified Machine Learning - Specialty Exam Questions.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q175-Q180):

NEW QUESTION # 175
A company wants to predict the classification of documents that are created from an application. New documents are saved to an Amazon S3 bucket every 3 seconds. The company has developed three versions of a machine learning (ML) model within Amazon SageMaker to classify document text. The company wants to deploy these three versions to predict the classification of each document.
Which approach will meet these requirements with the LEAST operational overhead?

- A. Configure an S3 event notification that invokes an AWS Lambda function when new documents are created. Configure the Lambda function to create three SageMaker batch transform jobs, one batch transform job for each model for each document.
- B. Deploy all the models to a single SageMaker endpoint. Treat each model as a production variant.
  Configure an S3 event notification that invokes an AWS Lambda function when new documents are created. Configure the Lambda function to call each production variant and return the results of each model.
- C. Deploy each model to its own SageMaker endpoint Configure an S3 event notification that invokes an AWS Lambda function when new documents are created. Configure the Lambda function to call each endpoint and return the results of each model.
- D. Deploy each model to its own SageMaker endpoint. Create three AWS Lambda functions. Configure each Lambda function to call a different endpoint and return the results. Configure three S3 event notifications to invoke the Lambda functions when new documents are created.

Answer: B

Explanation:
The approach that will meet the requirements with the least operational overhead is to deploy all the models to a single SageMaker endpoint, treat each model as a production variant, configure an S3 event notification that invokes an AWS Lambda function when new documents are created, and configure the Lambda function to call each production variant and return the results of each model. This approach involves the following steps:
* Deploy all the models to a single SageMaker endpoint. Amazon SageMaker is a service that can build, train, and deploy machine learning models. Amazon SageMaker can deploy multiple models to a single endpoint, which is a web service that can serve predictions from the models. Each model can be treated as a production variant, which is a version of the model that runs on one or more instances. Amazon SageMaker can distribute the traffic among the production variants according to the specified weights1.
* Treat each model as a production variant. Amazon SageMaker can deploy multiple models to a single endpoint, which is a web service that can serve predictions from the models. Each model can be treated as a production variant, which is a version of the model that runs on one or more instances. Amazon SageMaker can distribute the traffic among the production variants according to the specified weights1.
* Configure an S3 event notification that invokes an AWS Lambda function when new documents are created. Amazon S3 is a service that can store and retrieve any amount of data. Amazon S3 can send event notifications when certain actions occur on the objects in a bucket, such as object creation, deletion, or modification. Amazon S3 can invoke an AWS Lambda function as a destination for the event notifications. AWS Lambda is a service that can run code without provisioning or managing servers2.
* Configure the Lambda function to call each production variant and return the results of each model.
AWS Lambda can execute the code that can call the SageMaker endpoint and specify the production variant to invoke. AWS Lambda can use the AWS SDK or the SageMaker Runtime API to send requests to the endpoint and receive the predictions from the models. AWS Lambda can return the results of each model as a response to the event notification3.
The other options are not suitable because:
* Option A: Configuring an S3 event notification that invokes an AWS Lambda function when new documents are created, configuring the Lambda function to create three SageMaker batch transform jobs, one batch transform job for each model for each document, will incur more operational overhead than using a single SageMaker endpoint. Amazon SageMaker batch transform is a service that can process large datasets in batches and store the predictions in Amazon S3. Amazon SageMaker batch transform is not suitable for real-time inference, as it introduces a delay between the request and the response. Moreover, creating three batch transform jobs for each document will increase the complexity and cost of the solution4.
* Option C: Deploying each model to its own SageMaker endpoint, configuring an S3 event notification that invokes an AWS Lambda function when new documents are created, configuring the Lambda function to call each endpoint and return the results of each model, will incur more operational overhead than using a single SageMaker endpoint. Deploying each model to its own endpoint will increase the number of resources and endpoints to manage and monitor. Moreover, calling each endpoint separately will increase the latency and network traffic of the solution5.
* Option D: Deploying each model to its own SageMaker endpoint, creating three AWS Lambda functions, configuring each Lambda function to call a different endpoint and return the results, configuring three S3 event notifications to invoke the Lambda functions when new documents are created, will incur more operational overhead than using a single SageMaker endpoint and a single Lambda function. Deploying each model to its own endpoint will increase the number of resources and endpoints to manage and monitor. Creating three Lambda functions will increase the complexity and cost of the solution. Configuring three S3 event notifications will increase the number of triggers and destinations to manage and monitor6.
1: Deploying Multiple Models to a Single Endpoint - Amazon SageMaker
2: Configuring Amazon S3 Event Notifications - Amazon Simple Storage Service
3: Invoke an Endpoint - Amazon SageMaker
4: Get Inferences for an Entire Dataset with Batch Transform - Amazon SageMaker
5: Deploy a Model - Amazon SageMaker
6: AWS Lambda

NEW QUESTION # 176
A company is converting a large number of unstructured paper receipts into images. The company wants to create a model based on natural language processing (NLP) to find relevant entities such as date, location, and notes, as well as some custom entities such as receipt numbers.
The company is using optical character recognition (OCR) to extract text for data labeling. However, documents are in different structures and formats, and the company is facing challenges with setting up the manual workflows for each document type. Additionally, the company trained a named entity recognition (NER) model for custom entity detection using a small sample size. This model has a very low confidence score and will require retraining with a large dataset.
Which solution for text extraction and entity detection will require the LEAST amount of effort?

- A. Extract text from receipt images by using Amazon Textract. Use Amazon Comprehend for entity detection, and use Amazon Comprehend custom entity recognition for custom entity detection.
- B. Extract text from receipt images by using a deep learning OCR model from the AWS Marketplace. Use Amazon Comprehend for entity detection, and use Amazon Comprehend custom entity recognition for custom entity detection.
- C. Extract text from receipt images by using Amazon Textract. Use the Amazon SageMaker BlazingText algorithm to train on the text for entities and custom entities.
- D. Extract text from receipt images by using a deep learning OCR model from the AWS Marketplace. Use the NER deep learning model to extract entities.

Answer: A

NEW QUESTION # 177
A company is building a demand forecasting model based on machine learning (ML). In the development stage, an ML specialist uses an Amazon SageMaker notebook to perform feature engineering during work hours that consumes low amounts of CPU and memory resources. A data engineer uses the same notebook to perform data preprocessing once a day on average that requires very high memory and completes in only 2 hours. The data preprocessing is not configured to use GPU. All the processes are running well on an ml.m5.
4xlarge notebook instance.
The company receives an AWS Budgets alert that the billing for this month exceeds the allocated budget.
Which solution will result in the MOST cost savings?

- A. Keep the notebook instance type and size the same. Stop the notebook when it is not in use. Run data preprocessing on a P3 instance type with the same memory as the ml.m5.4xlarge instance by using Amazon SageMaker Processing.
- B. Change the notebook instance type to a smaller general-purpose instance. Stop the notebook when it is not in use. Run data preprocessing on an ml. r5 instance with the same memory size as the ml.m5.
4xlarge instance by using Amazon SageMaker Processing.
- C. Change the notebook instance type to a memory optimized instance with the same vCPU number as the ml.m5.4xlarge instance has. Stop the notebook when it is not in use. Run both data preprocessing and feature engineering development on that instance.
- D. Change the notebook instance type to a smaller general-purpose instance. Stop the notebook when it is not in use. Run data preprocessing on an R5 instance with the same memory size as the ml.m5.4xlarge instance by using the Reserved Instance option.

Answer: B

Explanation:
The best solution to reduce the cost of the notebook instance and the data preprocessing job is to change the notebook instance type to a smaller general-purpose instance, stop the notebook when it is not in use, and run data preprocessing on an ml.r5 instance with the same memory size as the ml.m5.4xlarge instance by using Amazon SageMaker Processing. This solution will result in the most cost savings because:
* Changing the notebook instance type to a smaller general-purpose instance will reduce the hourly cost of running the notebook, since the feature engineering development does not require high CPU and memory resources. For example, an ml.t3.medium instance costs $0.0464 per hour, while an ml.m5.
4xlarge instance costs $0.888 per hour1.
* Stopping the notebook when it is not in use will also reduce the cost, since the notebook will only incur charges when it is running. For example, if the notebook is used for 8 hours per day, 5 days per week, then stopping it when it is not in use will save about 76% of the monthly cost compared to leaving it running all the time2.
* Running data preprocessing on an ml.r5 instance with the same memory size as the ml.m5.4xlarge instance by using Amazon

SageMaker Processing will reduce the cost of the data preprocessing job, since the ml.r5 instance is optimized for memory-intensive workloads and has a lower cost per GB of memory than the ml.m5 instance. For example, an ml.r5.4xlarge instance has 128 GB of memory and costs $1.008 per hour, while an ml.m5.4xlarge instance has 64 GB of memory and costs $0.888 per hour1. Therefore, the ml.r5.4xlarge instance can process the same amount of data in half the time and at a lower cost than the ml.m5.4xlarge instance. Moreover, using Amazon SageMaker Processing will allow the data preprocessing job to run on a separate, fully managed infrastructure that can be scaled up or down as needed, without affecting the notebook instance.
The other options are not as effective as option C for the following reasons:
* Option A is not optimal because changing the notebook instance type to a memory optimized instance with the same vCPU number as the ml.m5.4xlarge instance has will not reduce the cost of the notebook, since the memory optimized instances have a higher cost per vCPU than the general-purpose instances. For example, an ml.r5.4xlarge instance has 16 vCPUs and costs $1.008 per hour, while an ml.
m5.4xlarge instance has 16 vCPUs and costs $0.888 per hour1. Moreover, running both data preprocessing and feature engineering development on the same instance will not take advantage of the scalability and flexibility of Amazon SageMaker Processing.
* Option B is not suitable because running data preprocessing on a P3 instance type with the same memory as the ml.m5.4xlarge instance by using Amazon SageMaker Processing will not reduce the cost of the data preprocessing job, since the P3 instance type is optimized for GPU-based workloads and has a higher cost per GB of memory than the ml.m5 or ml.r5 instance types. For example, an ml.p3.
2xlarge instance has 61 GB of memory and costs $3.06 per hour, while an ml.m5.4xlarge instance has
64 GB of memory and costs $0.888 per hour1. Moreover, the data preprocessing job does not require GPU, so using a P3 instance type will be wasteful and inefficient.
* Option D is not feasible because running data preprocessing on an R5 instance with the same memory size as the ml.m5.4xlarge instance by using the Reserved Instance option will not reduce the cost of the data preprocessing job, since the Reserved Instance option requires a commitment to a consistent amount of usage for a period of 1 or 3 years3. However, the data preprocessing job only runs once a day on average and completes in only 2 hours, so it does not have a consistent or predictable usage pattern. Therefore, using the Reserved Instance option will not provide any cost savings and may incur additional charges for unused capacity.
Amazon SageMaker Pricing
Manage Notebook Instances - Amazon SageMaker
Amazon EC2 Pricing - Reserved Instances

NEW QUESTION # 178
A company wants to conduct targeted marketing to sell solar panels to homeowners. The company wants to use machine learning (ML) technologies to identify which houses already have solar panels. The company has collected 8,000 satellite images as training data and will use Amazon SageMaker Ground Truth to label the data.
The company has a small internal team that is working on the project. The internal team has no ML expertise and no ML experience. Which solution will meet these requirements with the LEAST amount of effort from the internal team?

- A. Set up a public workforce. Use the public workforce to label the data. Use the SageMaker Object Detection algorithm to train a model. Use SageMaker batch transform for inference.
- B. Set up a private workforce that consists of the internal team. Use the private workforce and the SageMaker Ground Truth active learning feature to label the data. Use Amazon Rekognition Custom Labels for model training and hosting.
- C. Set up a private workforce that consists of the internal team. Use the private workforce and the SageMaker Ground Truth active learning feature to label the data. Use the SageMaker Object Detection algorithm to train a model. Use SageMaker batch transform for inference.
- D. Set up a private workforce that consists of the internal team. Use the private workforce to label the data. Use Amazon Rekognition Custom Labels for model training and hosting.

Answer: B

Explanation:
The solution A will meet the requirements with the least amount of effort from the internal team because it uses Amazon SageMaker Ground Truth and Amazon Rekognition Custom Labels, which are fully managed services that can provide the desired functionality.
The solution A involves the following steps:
Set up a private workforce that consists of the internal team. Use the private workforce and the SageMaker Ground Truth active learning feature to label the data. Amazon SageMaker Ground Truth is a service that can create high-quality training datasets for machine learning by using human labelers. A private workforce is a group of labelers that the company can manage and control. The internal team can use the private workforce to label the satellite images as having solar panels or not. The SageMaker Ground Truth active learning feature can reduce the labeling effort by using a machine learning model to automatically label the easy examples and only send the difficult ones to the human labelers1.
Use Amazon Rekognition Custom Labels for model training and hosting. Amazon Rekognition Custom Labels is a service that can

train and deploy custom machine learning models for image analysis. Amazon Rekognition Custom Labels can use the labeled data from SageMaker Ground Truth to train a model that can detect solar panels in satellite images. Amazon Rekognition Custom Labels can also host the model and provide an API endpoint for inference2.
The other options are not suitable because:
Option B: Setting up a private workforce that consists of the internal team, using the private workforce to label the data, and using Amazon Rekognition Custom Labels for model training and hosting will incur more effort from the internal team than using SageMaker Ground Truth active learning feature. The internal team will have to label all the images manually, without the assistance of the machine learning model that can automate some of the labeling tasks1.
Option C: Setting up a private workforce that consists of the internal team, using the private workforce and the SageMaker Ground Truth active learning feature to label the data, using the SageMaker Object Detection algorithm to train a model, and using SageMaker batch transform for inference will incur more operational overhead than using Amazon Rekognition Custom Labels. The company will have to manage the SageMaker training job, the model artifact, and the batch transform job. Moreover, SageMaker batch transform is not suitable for real-time inference, as it processes the data in batches and stores the results in Amazon S33.
Option D: Setting up a public workforce, using the public workforce to label the data, using the SageMaker Object Detection algorithm to train a model, and using SageMaker batch transform for inference will incur more operational overhead and cost than using a private workforce and Amazon Rekognition Custom Labels. A public workforce is a group of labelers from Amazon Mechanical Turk, a crowdsourcing marketplace. The company will have to pay the public workforce for each labeling task, and it may not have full control over the quality and security of the labeled data. The company will also have to manage the SageMaker training job, the model artifact, and the batch transform job, as explained in option C4.
References:
1: Amazon SageMaker Ground Truth
2: Amazon Rekognition Custom Labels
3: Amazon SageMaker Object Detection
4: Amazon Mechanical Turk

NEW QUESTION # 179
A large mobile network operating company is building a machine learning model to predict customers who are likely to unsubscribe from the service. The company plans to offer an incentive for these customers as the cost of churn is far greater than the cost of the incentive.
The model produces the following confusion matrix after evaluating on a test dataset of 100 customers:
Based on the model evaluation results, why is this a viable model for production?

n = 100	PREDICTED CHURN Yes	PREDICTED CHURN No
ACTUAL Churn Yes	10	4
Actual No	10	76

- A. The model is 86% accurate and the cost incurred by the company as a result of false positives is less than the false negatives.
- B. The precision of the model is 86%, which is less than the accuracy of the model.
- C. The model is 86% accurate and the cost incurred by the company as a result of false negatives is less than the false positives.
- D. The precision of the model is 86%, which is greater than the accuracy of the model.

Answer: A

Explanation:
Explanation
Based on the model evaluation results, this is a viable model for production because the model is 86% accurate and the cost incurred by the company as a result of false positives is less than the false negatives. The accuracy of the model is the proportion of correct predictions out of the total predictions, which can be calculated by adding the true positives and true negatives and dividing by the total number of observations. In this case, the accuracy of the model is (10 + 76) / 100 = 0.86, which means that the model correctly predicted
86% of the customers' churn status. The cost incurred by the company as a result of false positives and false negatives is the loss or damage that the company suffers when the model makes incorrect predictions. A false positive is when the model predicts that a customer will churn, but the customer actually does not churn. A false negative is when the model predicts that a customer will not churn, but the customer actually churns. In this case, the cost of a false positive is the incentive that the company offers to the customer who is predicted to churn, which is a relatively low cost. The cost of a false negative is the revenue that the company loses

when the customer churns, which is a relatively high cost. Therefore, the cost of a false positive is less than the cost of a false negative, and the company would prefer to have more false positives than false negatives. The model has 10 false positives and 4 false negatives, which means that the company's cost is lower than if the model had more false negatives and fewer false positives.

NEW QUESTION # 180

......

As we discussed above that the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam preparation material is available in three different formats. One of them is Amazon AWS-Certified-Machine-Learning-Specialty PDF questions format which is portable. Users of this format can print AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) real exam questions in this file to study without accessing any device. Furthermore, smart devices like laptops, smartphones, and tablets support the AWS-Certified-Machine-Learning-Specialty PDF Questions. Hence, you can carry this material to any place and revise AWS-Certified-Machine-Learning-Specialty exam questions conveniently without time restrictions.

AWS-Certified-Machine-Learning-Specialty New Test Bootcamp: https://www.easy4engine.com/AWS-Certified-Machine-Learning-Specialty-test-engine.html

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, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, ncon.edu.sa, 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, myportal.utt.edu.tt, 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, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by Easy4Engine: https://drive.google.com/open?id=15ckdvKl15eJoZcKpBDDXdUFysqz6R7nW