Amazon MLS-C01認定試験の出題範囲を絞る参考書

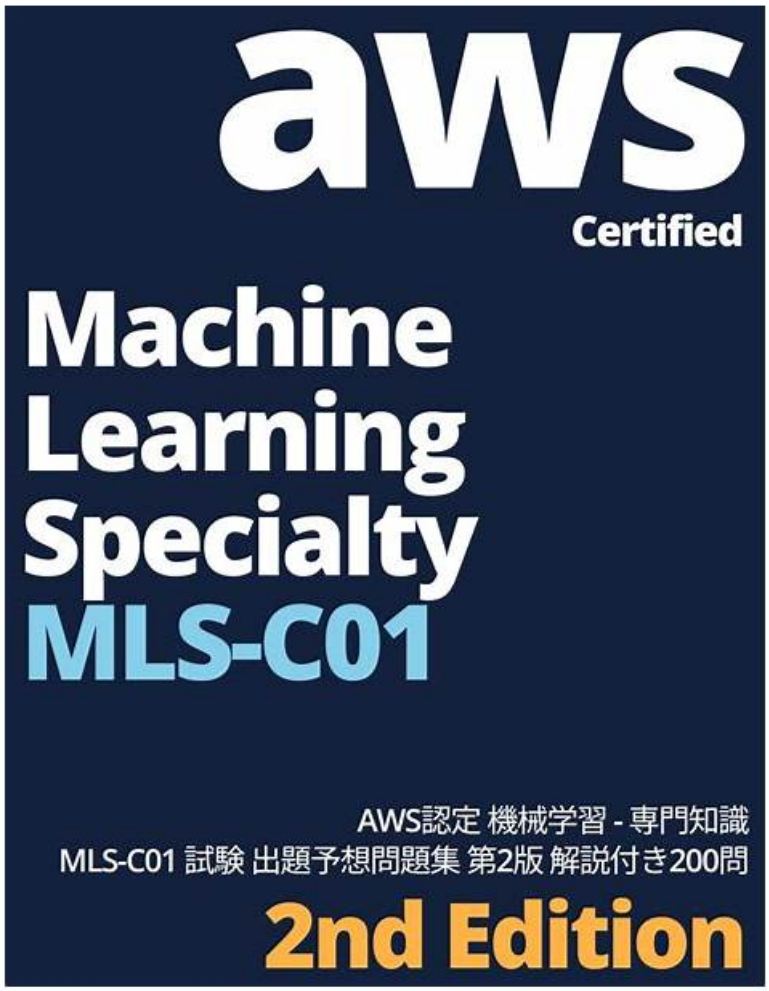

さらに、JPTestKing MLS-C01ダンプの一部が現在無料で提供されています：https://drive.google.com/open?id=1-KtNFaXytkJrvHPpxHgSEq3vMYZvee90

最も短い時間で自分のIT技能を増強したいけれど、質の良い学習教材がないので悩んでいますか。ご心配なく、JPTestKingのAmazonのMLS-C01試験トレーニング資料を手に入れるなら、ITに関する認定試験はなんでも楽に合格できます。JPTestKingの AmazonのMLS-C01試験トレーニング資料は高度に認証されたIT領域の専門家の経験と創造を含めているものです。JPTestKingは君にとって、ベストな選択だといっても良いです。

認定試験は、65の複数選択と多反応の質問で構成されており、候補者は試験を完了するために3時間与えられます。試験はオンラインで入手でき、世界中のどこからでも撮影できます。試験の準備をするために、候補者は、オンライントレーニングコース、練習試験、学習ガイドなど、AWSが提供するさまざまなリソースを利用できます。

Amazon MLS-C01試験の対象となるには、個人が機械学習の経験と、Amazon Sagemaker、Amazon Kinesis、Amazon RedshiftなどのAWSサービスを強く理解している必要があります。この試験は、機械学習でキャリアを進め、AWSサービスを使用して機械学習ソリューションを設計、実装、展開する能力を実証しようとしている専門家向けに設計されています。

>> MLS-C01日本語学習内容 <<

Amazon MLS-C01資格専門知識、MLS-C01日本語版参考書

JPTestKingのAmazonのMLS-C01「AWS Certified Machine Learning - Specialty」トレーニング資料を利用したら、初めて試験を受けるあなたでも一回で試験に合格できることを保証します。JPTestKingのAmazonのMLS-C01トレーニ

ング資料を利用しても合格しないのなら、我々は全額で返金することができます。あなたに他の同じ値段の製品を無料に送って差し上げます。

AWS Certified Machine Learning - Specialty試験は、65問の多肢選択と多肢選択式の質問から構成され、受験者は3時間で試験を完了する必要があります。試験は、AWSプラットフォーム上での機械学習ソリューションの設計、実装、展開、およびメンテナンス能力を評価します。試験に合格するためには、候補者はAWSプラットフォームでの機械学習ソリューションの開発とメンテナンスに1年以上の経験が必要です。試験に合格すると、AWS Certified Machine Learning - Specialty認定を取得し、グローバルに認められ、AWSプラットフォーム上の機械学習分野での専門知識を証明することができます。

Amazon AWS Certified Machine Learning - Specialty 認定 MLS-C01 試験問題 (Q261-Q266):

質問 # 261
A retail company is selling products through a global online marketplace. The company wants to use machine learning (ML) to analyze customer feedback and identify specific areas for improvement. A developer has built a tool that collects customer reviews from the online marketplace and stores them in an Amazon S3 bucket. This process yields a dataset of 40 reviews. A data scientist building the ML models must identify additional sources of data to increase the size of the dataset.
Which data sources should the data scientist use to augment the dataset of reviews? (Choose three.)

- A. Product sales revenue figures for the company
- B. Instruction manuals for the company's products
- C. A publicly available collection of customer reviews
- D. A publicly available collection of news articles
- E. Emails exchanged by customers and the company's customer service agents
- F. Social media posts containing the name of the company or its products

正解：C、E、F

解説：
The data sources that the data scientist should use to augment the dataset of reviews are those that contain relevant and diverse customer feedback about the company or its products. Emails exchanged by customers and the company's customer service agents can provide valuable insights into the issues and complaints that customers have, as well as the solutions and responses that the company offers. Social media posts containing the name of the company or its products can capture the opinions and sentiments of customers and potential customers, as well as their reactions to marketing campaigns and product launches. A publicly available collection of customer reviews can provide a large and varied sample of feedback from different online platforms and marketplaces, which can help to generalize the ML models and avoid bias.
Detect sentiment from customer reviews using Amazon Comprehend | AWS Machine Learning Blog How to Apply Machine Learning to Customer Feedback

質問 # 262
A health care company is planning to use neural networks to classify their X-ray images into normal and abnormal classes. The labeled data is divided into a training set of 1,000 images and a test set of 200 images.
The initial training of a neural network model with 50 hidden layers yielded 99% accuracy on the training set, but only 55% accuracy on the test set.
What changes should the Specialist consider to solve this issue? (Choose three.)

- A. Choose a higher number of layers
- B. Choose a smaller learning rate
- C. Enable early stopping
- D. Enable dropout
- E. Choose a lower number of layers
- F. Include all the images from the test set in the training set

正解：A、D、F

質問 # 263
A Machine Learning Specialist needs to create a data repository to hold a large amount of time-based training data for a new model.

In the source system, new files are added every hour Throughout a single 24-hour period, the volume of hourly updates will change significantly. The Specialist always wants to train on the last 24 hours of the data Which type of data repository is the MOST cost-effective solution?

- A. An Amazon RDS database with hourly table partitions
- B. An Amazon S3 data lake with hourly object prefixes
- C. An Amazon EBS-backed Amazon EC2 instance with hourly directories
- D. An Amazon EMR cluster with hourly hive partitions on Amazon EBS volumes

正解：B

質問 # 264
An automotive company uses computer vision in its autonomous cars. The company trained its object detection models successfully by using transfer learning from a convolutional neural network (CNN). The company trained the models by using PyTorch through the Amazon SageMaker SDK.
The vehicles have limited hardware and compute power. The company wants to optimize the model to reduce memory, battery, and hardware consumption without a significant sacrifice in accuracy.
Which solution will improve the computational efficiency of the models?

- A. Use Amazon SageMaker Debugger to gain visibility into the training weights, gradients, biases, and activation outputs. Compute the filter ranks based on the training information. Apply pruning to remove the low-ranking filters. Set the new weights based on the pruned set of filters. Run a new training job with the pruned model.
- B. Use Amazon CloudWatch metrics to gain visibility into the SageMaker training weights, gradients, biases, and activation outputs. Compute the filter ranks based on the training information. Apply pruning to remove the low-ranking filters. Set new weights based on the pruned set of filters. Run a new training job with the pruned model.
- C. Use Amazon SageMaker Ground Truth to build and run data labeling workflows. Collect a larger labeled dataset with the labelling workflows. Run a new training job that uses the new labeled data with previous training data.
- D. Use Amazon SageMaker Model Monitor to gain visibility into the ModelLatency metric and OverheadLatency metric of the model after the company deploys the model. Increase the model learning rate. Run a new training job.

正解：A

解説：
The solution C will improve the computational efficiency of the models because it uses Amazon SageMaker Debugger and pruning, which are techniques that can reduce the size and complexity of the convolutional neural network (CNN) models. The solution C involves the following steps:
Use Amazon SageMaker Debugger to gain visibility into the training weights, gradients, biases, and activation outputs. Amazon SageMaker Debugger is a service that can capture and analyze the tensors that are emitted during the training process of machine learning models. Amazon SageMaker Debugger can provide insights into the model performance, quality, and convergence. Amazon SageMaker Debugger can also help to identify and diagnose issues such as overfitting, underfitting, vanishing gradients, and exploding gradients1.
Compute the filter ranks based on the training information. Filter ranking is a technique that can measure the importance of each filter in a convolutional layer based on some criterion, such as the average percentage of zero activations or the L1-norm of the filter weights. Filter ranking can help to identify the filters that have little or no contribution to the model output, and thus can be removed without affecting the model accuracy2.
Apply pruning to remove the low-ranking filters. Pruning is a technique that can reduce the size and complexity of a neural network by removing the redundant or irrelevant parts of the network, such as neurons, connections, or filters. Pruning can help to improve the computational efficiency, memory usage, and inference speed of the model, as well as to prevent overfitting and improve generalization3.
Set the new weights based on the pruned set of filters. After pruning, the model will have a smaller and simpler architecture, with fewer filters in each convolutional layer. The new weights of the model can be set based on the pruned set of filters, either by initializing them randomly or by fine-tuning them from the original weights4.
Run a new training job with the pruned model. The pruned model can be trained again with the same or a different dataset, using the same or a different framework or algorithm. The new training job can use the same or a different configuration of Amazon SageMaker, such as the instance type, the hyperparameters, or the data ingestion mode. The new training job can also use Amazon SageMaker Debugger to monitor and analyze the training process and the model quality5.
The other options are not suitable because:
Option A: Using Amazon CloudWatch metrics to gain visibility into the SageMaker training weights, gradients, biases, and activation outputs will not be as effective as using Amazon SageMaker Debugger. Amazon CloudWatch is a service that can monitor and observe the operational health and performance of AWS resources and applications. Amazon CloudWatch can provide metrics,

alarms, dashboards, and logs for various AWS services, including Amazon SageMaker. However, Amazon CloudWatch does not provide the same level of granularity and detail as Amazon SageMaker Debugger for the tensors that are emitted during the training process of machine learning models. Amazon CloudWatch metrics are mainly focused on the resource utilization and the training progress, not on the model performance, quality, and convergence6.
Option B: Using Amazon SageMaker Ground Truth to build and run data labeling workflows and collecting a larger labeled dataset with the labeling workflows will not improve the computational efficiency of the models. Amazon SageMaker Ground Truth is a service that can create high-quality training datasets for machine learning by using human labelers. A larger labeled dataset can help to improve the model accuracy and generalization, but it will not reduce the memory, battery, and hardware consumption of the model. Moreover, a larger labeled dataset may increase the training time and cost of the model7.
Option D: Using Amazon SageMaker Model Monitor to gain visibility into the ModelLatency metric and OverheadLatency metric of the model after the company deploys the model and increasing the model learning rate will not improve the computational efficiency of the models. Amazon SageMaker Model Monitor is a service that can monitor and analyze the quality and performance of machine learning models that are deployed on Amazon SageMaker endpoints. The ModelLatency metric and the OverheadLatency metric can measure the inference latency of the model and the endpoint, respectively. However, these metrics do not provide any information about the training weights, gradients, biases, and activation outputs of the model, which are needed for pruning. Moreover, increasing the model learning rate will not reduce the size and complexity of the model, but it may affect the model convergence and accuracy.
References:
1: Amazon SageMaker Debugger
2: Pruning Convolutional Neural Networks for Resource Efficient Inference
3: Pruning Neural Networks: A Survey
4: Learning both Weights and Connections for Efficient Neural Networks
5: Amazon SageMaker Training Jobs
6: Amazon CloudWatch Metrics for Amazon SageMaker
7: Amazon SageMaker Ground Truth
8: Amazon SageMaker Model Monitor

質問 # 265
A retail company is ingesting purchasing records from its network of 20,000 stores to Amazon S3 by using Amazon Kinesis Data Firehose. The company uses a small, server-based application in each store to send the data to AWS over the internet. The company uses this data to train a machine learning model that is retrained each day. The company's data science team has identified existing attributes on these records that could be combined to create an improved model.
Which change will create the required transformed records with the LEAST operational overhead?

- A. Deploy an Amazon S3 File Gateway in the stores. Update the in-store software to deliver data to the S3 File Gateway. Use a scheduled daily AWS Glue job to transform the data that the S3 File Gateway delivers to Amazon S3.
- B. Deploy an Amazon EMR cluster that runs Apache Spark and includes the transformation logic. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to launch the cluster each day and transform the records that accumulate in Amazon S3. Deliver the transformed records to Amazon S3.
- C. Launch a fleet of Amazon EC2 instances that include the transformation logic. Configure the EC2 instances with a daily cron job to transform the records that accumulate in Amazon S3. Deliver the transformed records to Amazon S3.
- D. Create an AWS Lambda function that can transform the incoming records. Enable data transformation on the ingestion Kinesis Data Firehose delivery stream. Use the Lambda function as the invocation target.

正解：D

解説：
The solution A will create the required transformed records with the least operational overhead because it uses AWS Lambda and Amazon Kinesis Data Firehose, which are fully managed services that can provide the desired functionality. The solution A involves the following steps:
Create an AWS Lambda function that can transform the incoming records. AWS Lambda is a service that can run code without provisioning or managing servers. AWS Lambda can execute the transformation logic on the purchasing records and add the new attributes to the records1.
Enable data transformation on the ingestion Kinesis Data Firehose delivery stream. Use the Lambda function as the invocation target. Amazon Kinesis Data Firehose is a service that can capture, transform, and load streaming data into AWS data stores. Amazon Kinesis Data Firehose can enable data transformation and invoke the Lambda function to process the incoming records before delivering them to Amazon S3. This can reduce the operational overhead of managing the transformation process and the data storage2.
The other options are not suitable because:
Option B: Deploying an Amazon EMR cluster that runs Apache Spark and includes the transformation logic, using Amazon

EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to launch the cluster each day and transform the records that accumulate in Amazon S3, and delivering the transformed records to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the Amazon EMR cluster, the Apache Spark application, the AWS Lambda function, and the Amazon EventBridge rule. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day3.
Option C: Deploying an Amazon S3 File Gateway in the stores, updating the in-store software to deliver data to the S3 File Gateway, and using a scheduled daily AWS Glue job to transform the data that the S3 File Gateway delivers to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the S3 File Gateway, the in-store software, and the AWS Glue job. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day4.
Option D: Launching a fleet of Amazon EC2 instances that include the transformation logic, configuring the EC2 instances with a daily cron job to transform the records that accumulate in Amazon S3, and delivering the transformed records to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the EC2 instances, the transformation code, and the cron job. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day5.
References:
1: AWS Lambda
2: Amazon Kinesis Data Firehose
3: Amazon EMR
4: Amazon S3 File Gateway
5: Amazon EC2

質問 # 266
.....

2026年JPTestKingの最新MLS-C01 PDFダンプおよびMLS-C01試験エンジンの無料共有：https://drive.google.com/open?id=1-KtNFaXytkJrvHPpxHgSEq3vMYZvee90