MLS-C01 Q&As

AWS Certified Machine Learning - Specialty (MLS-C01)

Pass Amazon MLS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/aws-certified-machine-learning-specialty.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A machine learning (ML) specialist is developing a model for a company. The model will classify and predict sequences of objects that are displayed in a video. The ML specialist decides to use a hybrid architecture that consists of a convolutional neural network (CNN) followed by a classifier three-layer recurrent neural network (RNN).

The company developed a similar model previously but trained the model to classify a different set of objects. The ML specialist wants to save time by using the previously trained model and adapting the model for the current use case and set of objects.

Which combination of steps will accomplish this goal with the LEAST amount of effort? (Choose two.)

A. Reinitialize the weights of the entire CNN. Retrain the CNN on the classification task by using the new set of objects.

B. Reinitialize the weights of the entire network. Retrain the entire network on the prediction task by using the new set of objects.

C. Reinitialize the weights of the entire RNN. Retrain the entire model on the prediction task by using the new set of objects.

D. Reinitialize the weights of the last fully connected layer of the CNN. Retrain the CNN on the classification task by using the new set of objects.

E. Reinitialize the weights of the last layer of the RNN. Retrain the entire model on the prediction task by using the new set of objects.

Correct Answer: DE

QUESTION 2

A machine learning engineer is building a bird classification model. The engineer randomly separates a dataset into a training dataset and a validation dataset. During the training phase, the model achieves very high accuracy. However, the model did not generalize well during validation of the validation dataset. The engineer realizes that the original dataset was imbalanced. What should the engineer do to improve the validation accuracy of the model?

A. Perform stratified sampling on the original dataset.

B. Acquire additional data about the majority classes in the original dataset.

C. Use a smaller, randomly sampled version of the training dataset.

D. Perform systematic sampling on the original dataset.

Correct Answer: A

Balanced Class Representation. Stratified sampling divides the original dataset into strata (groups) based on the class labels. It then selects instances from each stratum in a proportional manner, ensuring that the class distribution in the training and validation datasets reflects the original class distribution. Improved Generalization. By having a balanced representation of all classes in the training and validation datasets, the model is exposed to a diverse range of instances during training. This helps the model learn the distinguishing features of each class more effectively, leading to better generalization performance on the validation dataset. Addressing Imbalanced Data. Stratified sampling directly addresses the issue of imbalanced data, which was identified as the root cause of the model\\'s poor generalization performance on the validation dataset.

QUESTION 3

An Machine Learning Specialist discover the following statistics while experimenting on a model.

Experiment 1
Baseline model
Train error = 5%
Test error = 16%

Experiment 2
The Specialist added more layers and neurons to the model and received the following results:
Train error = 5.2%
Test error = 15.7%

Experiment 3
The Specialist reverted back to the original number of neurons form Experiment 1 and implemented regularization in the neural network, which yielded the following results:
Train error = 4.7%
Test error = 9.5%

What can the Specialist from the experiments?

A. The model In Experiment 1 had a high variance error lhat was reduced in Experiment 3 by regularization Experiment 2 shows that there is minimal bias error in Experiment 1

B. The model in Experiment 1 had a high bias error that was reduced in Experiment 3 by regularization Experiment 2 shows that there is minimal variance error in Experiment 1

C. The model in Experiment 1 had a high bias error and a high variance error that were reduced in Experiment 3 by regularization Experiment 2 shows thai high bias cannot be reduced by increasing layers and neurons in the model

D. The model in Experiment 1 had a high random noise error that was reduced in Expenment 3 by regularization Expenment 2 shows that random noise cannot be reduced by increasing layers and neurons in the model

Correct Answer: C

QUESTION 4

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. Modify the bash_profile file in the container and add a bash command to start the training program

B. Use CMD config in the Dockerfile to add the training program as a CMD of the image

C. Configure the training program as an ENTRYPOINT named train

D. Copy the training program to directory /opt/ml/train

Correct Answer: C

"To configure a Docker container to run as an executable, use an ENTRYPOINT instruction in a Dockerfile. SageMaker overrides any default CMD statement in a container by specifying the train argument after the image name"

https://docs.aws.amazon.com/sagemaker/latest/dg/your-algorithms-training-algo-dockerfile.html

QUESTION 5

A trucking company is collecting live image data from its fleet of trucks across the globe. The data is growing rapidly and approximately 100 GB of new data is generated every day. The company wants to explore machine learning uses cases

while ensuring the data is only accessible to specific IAM users.

Which storage option provides the most processing flexibility and will allow access control with IAM?

A. Use a database, such as Amazon DynamoDB, to store the images, and set the IAM policies to restrict access to only the desired IAM users.

B. Use an Amazon S3-backed data lake to store the raw images, and set up the permissions using bucket policies.

C. Setup up Amazon EMR with Hadoop Distributed File System (HDFS) to store the files, and restrict access to the EMR instances using IAM policies.

D. Configure Amazon EFS with IAM policies to make the data available to Amazon EC2 instances owned by the IAM users.

Correct Answer: B

https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-file-systems.html