

AWS-Certified-Machine-Learning-Specialty Exam Current Exam Content- Perfect AWS-Certified-Machine- Learning-Specialty Free Vce Dumps Pass Success



P.S. Free 2025 Amazon AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by Test4Cram
<https://drive.google.com/open?id=1CjgLfLwYmvpwAE2VjZsLYU2wy5UyNGm>

Web-based AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice exam is a convenient format to evaluate and improve preparation for the exam. It is a AWS-Certified-Machine-Learning-Specialty browser-based application, which means you can access it from any operating system with an internet connection and a web browser. Unlike the desktop-based exam simulation software, the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) browser-based practice test requires no plugins and software installation.

Amazon MLS-C01 (AWS Certified Machine Learning - Specialty) Certification Exam is a highly sought-after certification in the field of Machine Learning. It is designed for individuals who have a strong understanding of machine learning concepts and techniques and are interested in showcasing their expertise in this domain. AWS Certified Machine Learning - Specialty certification is specifically designed for professionals who are interested in building, deploying, and maintaining machine learning solutions on Amazon Web Services (AWS) platform.

The AWS-Certified-Machine-Learning-Specialty certification exam is a valuable credential for professionals who want to demonstrate their expertise in machine learning on the AWS platform. AWS Certified Machine Learning - Specialty certification validates an individual's ability to design, implement, and deploy machine learning models on AWS, and passing the exam can lead to higher-paying job opportunities and access to the AWS Certified Community.

The AWS Certified Machine Learning - Specialty Exam covers a wide range of topics related to machine learning, including data preparation, feature engineering, model selection and evaluation, and deployment. It also covers advanced topics such as deep learning, natural language processing, and computer vision. AWS-Certified-Machine-Learning-Specialty Exam consists of multiple-choice questions and is available in English, Simplified Chinese, Korean, and Japanese.

>> **AWS-Certified-Machine-Learning-Specialty Current Exam Content** <<

AWS-Certified-Machine-Learning-Specialty Current Exam Content Authoritative Questions Pool Only at Test4Cram

Test4Cram is fully aware of the fact that preparing successfully for the Amazon AWS-Certified-Machine-Learning-Specialty exam in one go is a necessity because of the expensive registration fee. For applicants like you, success in the AWS Certified Machine Learning - Specialty exam on the first attempt is crucial to saving money and time. Our Free Amazon AWS-Certified-Machine-Learning-Specialty Exam Questions will help you decide fast to buy the premium ones.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q239-Q244):

NEW QUESTION # 239

A Machine Learning Specialist is creating a new natural language processing application that processes a dataset comprised of 1 million sentences. The aim is to then run Word2Vec to generate embeddings of the sentences and enable different types of predictions - Here is an example from the dataset:

"The quck BROWN FOX jumps over the lazy dog "

Which of the following are the operations the Specialist needs to perform to correctly sanitize and prepare the data in a repeatable manner? (Select THREE)

- A. Correct the typography on "quck" to "quick."
- B. One-hot encode all words in the sentence
- C. Remove stop words using an English stopword dictionary.
- D. Normalize all words by making the sentence lowercase
- E. Tokenize the sentence into words.
- F. Perform part-of-speech tagging and keep the action verb and the nouns only

Answer: C,D,E

NEW QUESTION # 240

A manufacturing company needs to identify returned smartphones that have been damaged by moisture. The company has an automated process that produces 2,000 diagnostic values for each phone. The database contains more than five million phone evaluations. The evaluation process is consistent, and there are no missing values in the data. A machine learning (ML) specialist has trained an Amazon SageMaker linear learner ML model to classify phones as moisture damaged or not moisture damaged by using all available features. The model's F1 score is 0.6.

What changes in model training would MOST likely improve the model's F1 score? (Select TWO.)

- A. Continue to use the SageMaker linear learner algorithm. Set the predictor type to regressor.
- B. Continue to use the SageMaker linear learner algorithm. Reduce the number of features with the SageMaker principal component analysis (PCA) algorithm.
- C. Use the SageMaker k-means algorithm with k of less than 1,000 to train the model.
- D. Use the SageMaker k-nearest neighbors (k-NN) algorithm. Set a dimension reduction target of less than 1,000 to train the model.
- E. Continue to use the SageMaker linear learner algorithm. Reduce the number of features with the scikit-learn multi-dimensional scaling (MDS) algorithm.

Answer: B,D

Explanation:

* Option A is incorrect because reducing the number of features with the SageMaker PCA algorithm can help remove noise and redundancy from the data, and improve the model's performance. PCA is a dimensionality reduction technique that transforms the original features into a smaller set of linearly uncorrelated features called principal components. The SageMaker linear learner algorithm supports PCA as a built-in feature transformation option.

* Option E is incorrect because using the SageMaker k-NN algorithm with a dimension reduction target of less than 1,000 can help the model learn from the similarity of the data points, and improve the model's performance. k-NN is a non-parametric algorithm that classifies an input based on the majority vote of its k nearest neighbors in the feature space. The SageMaker k-NN algorithm supports dimension reduction as a built-in feature transformation option.

* Option B is incorrect because using the scikit-learn MDS algorithm to reduce the number of features is not a feasible option, as MDS is a computationally expensive technique that does not scale well to large datasets. MDS is a dimensionality reduction technique that tries to preserve the pairwise distances between the original data points in a lower-dimensional space.

* Option C is incorrect because setting the predictor type to regressor would change the model's objective from classification to regression, which is not suitable for the given problem. A regressor model would output a continuous value instead of a binary label for each phone.

* Option D is incorrect because using the SageMaker k-means algorithm with k of less than 1,000 would not help the model classify the phones, as k-means is a clustering algorithm that groups the data points into k clusters based on their similarity, without using any labels. A clustering model would not output a binary label for each phone.

Amazon SageMaker Linear Learner Algorithm

Amazon SageMaker K-Nearest Neighbors (k-NN) Algorithm

[Principal Component Analysis - Scikit-learn]

[Multidimensional Scaling - Scikit-learn]

NEW QUESTION # 241

A data science team is working with a tabular dataset that the team stores in Amazon S3. The team wants to experiment with different feature transformations such as categorical feature encoding. Then the team wants to visualize the resulting distribution of the dataset. After the team finds an appropriate set of feature transformations, the team wants to automate the workflow for feature transformations.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Use Amazon SageMaker Data Wrangler preconfigured transformations to explore feature transformations. Use SageMaker Data Wrangler templates for visualization. Export the feature processing workflow to a SageMaker pipeline for automation.
- B. Use AWS Glue Studio with custom code to experiment with different feature transformations. Save the transformations to Amazon S3. Use Amazon QuickSight for visualization. Package the feature processing steps into an AWS Lambda function for automation.
- C. Use an Amazon SageMaker notebook instance to experiment with different feature transformations. Save the transformations to Amazon S3. Use Amazon QuickSight for visualization. Package the feature processing steps into an AWS Lambda function for automation.
- D. Use Amazon SageMaker Data Wrangler preconfigured transformations to experiment with different feature transformations. Save the transformations to Amazon S3. Use Amazon QuickSight for visualization. Package each feature transformation step into a separate AWS Lambda function. Use AWS Step Functions for workflow automation.

Answer: A

Explanation:

The solution A will meet the requirements with the most operational efficiency because it uses Amazon SageMaker Data Wrangler, which is a service that simplifies the process of data preparation and feature engineering for machine learning. The solution A involves the following steps:

Use Amazon SageMaker Data Wrangler preconfigured transformations to explore feature transformations. Amazon SageMaker Data Wrangler provides a visual interface that allows data scientists to apply various transformations to their tabular data, such as encoding categorical features, scaling numerical features, imputing missing values, and more. Amazon SageMaker Data Wrangler also supports custom transformations using Python code or SQL queries¹.

Use SageMaker Data Wrangler templates for visualization. Amazon SageMaker Data Wrangler also provides a set of templates that can generate visualizations of the data, such as histograms, scatter plots, box plots, and more. These visualizations can help data scientists to understand the distribution and characteristics of the data, and to compare the effects of different feature transformations¹.

Export the feature processing workflow to a SageMaker pipeline for automation. Amazon SageMaker Data Wrangler can export the feature processing workflow as a SageMaker pipeline, which is a service that orchestrates and automates machine learning workflows. A SageMaker pipeline can run the feature processing steps as a preprocessing step, and then feed the output to a training step or an inference step. This can reduce the operational overhead of managing the feature processing workflow and ensure its consistency and reproducibility².

The other options are not suitable because:

Option B: Using an Amazon SageMaker notebook instance to experiment with different feature transformations, saving the transformations to Amazon S3, using Amazon QuickSight for visualization, and packaging the feature processing steps into an AWS Lambda function for automation will incur more operational overhead than using Amazon SageMaker Data Wrangler. The data scientist will have to write the code for the feature transformations, the data storage, the data visualization, and the Lambda function. Moreover, AWS Lambda has limitations on the execution time, memory size, and package size, which may not be sufficient for complex feature processing tasks³.

Option C: Using AWS Glue Studio with custom code to experiment with different feature transformations, saving the transformations to Amazon S3, using Amazon QuickSight for visualization, and packaging the feature processing steps into an AWS Lambda function for automation will incur more operational overhead than using Amazon SageMaker Data Wrangler. AWS Glue Studio is a visual interface that allows data engineers to create and run extract, transform, and load (ETL) jobs on AWS Glue. However, AWS Glue Studio does not provide preconfigured transformations or templates for feature engineering or data visualization. The data scientist will have to write custom code for these tasks, as well as for the Lambda function. Moreover, AWS Glue Studio is not integrated with SageMaker pipelines, and it may not be optimized for machine learning workflows⁴.

Option D: Using Amazon SageMaker Data Wrangler preconfigured transformations to experiment with different feature transformations, saving the transformations to Amazon S3, using Amazon QuickSight for visualization, packaging each feature transformation step into a separate AWS Lambda function, and using AWS Step Functions for workflow automation will incur more operational overhead than using Amazon SageMaker Data Wrangler. The data scientist will have to create and manage multiple AWS Lambda functions and AWS Step Functions, which can increase the complexity and cost of the solution. Moreover, AWS Lambda and AWS Step Functions may not be compatible with SageMaker pipelines, and they may not be optimized for machine learning workflows⁵.

References:

- 1: Amazon SageMaker Data Wrangler
- 2: Amazon SageMaker Pipelines
- 3: AWS Lambda
- 4: AWS Glue Studio
- 5: AWS Step Functions

NEW QUESTION # 242

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

Answer: B,C,D

Explanation:

The problem described in the question is a case of overfitting, where the neural network model performs well on the training data but poorly on the test data. This means that the model has learned the noise and specific patterns of the training data, but cannot generalize to new and unseen data. To solve this issue, the Specialist should consider the following changes:

* Choose a lower number of layers: Reducing the number of layers can reduce the complexity and capacity of the neural network model, making it less prone to overfitting. A model with 50 hidden layers is likely too deep for the given data size and task. A simpler model with fewer layers can learn the essential features of the data without memorizing the noise.

* Enable dropout: Dropout is a regularization technique that randomly drops out some units in the neural network during training. This prevents the units from co-adapting too much and forces the model to learn more robust features. Dropout can improve the generalization and test performance of the model by reducing overfitting.

* Enable early stopping: Early stopping is another regularization technique that monitors the validation error during training and stops the training process when the validation error stops decreasing or starts increasing. This prevents the model from overtraining on the training data and reduces overfitting.

References:

- * Deep Learning - Machine Learning Lens
- * How to Avoid Overfitting in Deep Learning Neural Networks
- * How to Identify Overfitting Machine Learning Models in Scikit-Learn

NEW QUESTION # 243

A city wants to monitor its air quality to address the consequences of air pollution. A Machine Learning Specialist needs to forecast the air quality in parts per million of contaminants for the next 2 days in the city. As this is a prototype, only daily data from the last year is available.

Which model is MOST likely to provide the best results in Amazon SageMaker?

- **A. Use the Amazon SageMaker Linear Learner algorithm on the single time series consisting of the full year of data with a predictor_typeof regressor.**
- B. Use the Amazon SageMaker Linear Learner algorithm on the single time series consisting of the full year of data with a predictor_typeof classifier.
- C. Use Amazon SageMaker Random Cut Forest (RCF) on the single time series consisting of the full year of data.
- D. Use the Amazon SageMaker k-Nearest-Neighbors (kNN) algorithm on the single time series consisting of the full year of data with a predictor_typeof regressor.

Answer: A

Explanation:

Explanation/Reference: <https://aws.amazon.com/blogs/machine-learning/build-a-model-to-predict-the-impact-of-weather-on-urban-air-quality-using-amazon-sagemaker/?ref=Welcome.AI>

BONUS!!! Download part of Test4Cram AWS-Certified-Machine-Learning-Specialty dumps for free:
<https://drive.google.com/open?id=1CjgLflLwYmvpwAE2VjZsLYU2wy5UyNGm>