

Three formats of the Amazon AWS-Certified-Machine-Learning-Specialty Exam Dumps



BTW, DOWNLOAD part of Itcerttest AWS-Certified-Machine-Learning-Specialty dumps from Cloud Storage:
https://drive.google.com/open?id=1kjm_YldrxUhJKb2_kILBfRNWkS69MRQ

The greatest product or service in the world comes from the talents in the organization. Talents have given life to work and have driven companies to move forward. Paying attention to talent development has become the core strategy for today's corporate development. Perhaps you will need our AWS-Certified-Machine-Learning-Specialty Learning Materials. No matter what your ability to improve, our AWS-Certified-Machine-Learning-Specialty practice questions can meet your needs. And with our AWS-Certified-Machine-Learning-Specialty exam questions, you will know you can be better.

To be eligible for the AWS Certified Machine Learning - Specialty exam, candidates should have a minimum of one to two years of experience in machine learning and a solid understanding of AWS services and architecture. Additionally, candidates should be familiar with programming languages such as Python, R, and Java, and have experience working with data processing and analysis tools such as Apache Spark and TensorFlow. Passing AWS-Certified-Machine-Learning-Specialty exam can help professionals showcase their skills and expertise in the field of machine learning and open up new career opportunities.

Amazon MLS-C01 (AWS Certified Machine Learning - Specialty) Exam is a certification exam designed for individuals who are seeking to validate their skills and knowledge in machine learning on the Amazon Web Services (AWS) platform. AWS-Certified-Machine-Learning-Specialty Exam is aimed at professionals who have experience in designing, implementing, deploying, and maintaining machine learning solutions using AWS services.

>> **Examinations AWS-Certified-Machine-Learning-Specialty Actual Questions <<**

Free PDF 2026 Fantastic AWS-Certified-Machine-Learning-Specialty: Examinations AWS Certified Machine Learning - Specialty Actual Questions

You deserve this opportunity to win and try to make some difference in your life if you want to attend the AWS-Certified-Machine-Learning-Specialty exam and get the certification by the help of our AWS-Certified-Machine-Learning-Specialty practice braindumps. As we all know, all companies will pay more attention on the staffs who have more certifications which is a symbol of better understanding and efficiency on the job. Our AWS-Certified-Machine-Learning-Specialty Study Materials have the high pass rate as 98% to 100%, hope you can use it fully and pass the exam smoothly.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q240-Q245):

NEW QUESTION # 240

A company is building a new supervised classification model in an AWS environment. The company's data science team notices that the dataset has a large quantity of variables. All the variables are numeric. The model accuracy for training and validation is low. The model's processing time is affected by high latency. The data science team needs to increase the accuracy of the model and decrease the processing.

How it should the data science team do to meet these requirements?

- A. Use a multiple correspondence analysis (MCA) model
- **B. Use a principal component analysis (PCA) model.**
- C. Create new features and interaction variables.
- D. Apply normalization on the feature set.

Answer: B

Explanation:

The best way to meet the requirements is to use a principal component analysis (PCA) model, which is a technique that reduces the dimensionality of the dataset by transforming the original variables into a smaller set of new variables, called principal components, that capture most of the variance and information in the data. This technique has the following advantages:

- * It can increase the accuracy of the model by removing noise, redundancy, and multicollinearity from the data, and by enhancing the interpretability and generalization of the model.
- * It can decrease the processing time of the model by reducing the number of features and the computational complexity of the model, and by improving the convergence and stability of the model.
- * It is suitable for numeric variables, as it relies on the covariance or correlation matrix of the data, and it can handle a large quantity of variables, as it can extract the most relevant ones.

The other options are not effective or appropriate, because they have the following drawbacks:

- * A: Creating new features and interaction variables can increase the accuracy of the model by capturing more complex and nonlinear relationships in the data, but it can also increase the processing time of the model by adding more features and increasing the computational complexity of the model. Moreover, it can introduce more noise, redundancy, and multicollinearity in the data, which can degrade the performance and interpretability of the model.
- * C: Applying normalization on the feature set can increase the accuracy of the model by scaling the features to a common range and avoiding the dominance of some features over others, but it can also decrease the processing time of the model by reducing the numerical instability and improving the convergence of the model. However, normalization alone is not enough to address the high dimensionality and high latency issues of the dataset, as it does not reduce the number of features or the variance in the data.
- * D: Using a multiple correspondence analysis (MCA) model is not suitable for numeric variables, as it is a technique that reduces the dimensionality of the dataset by transforming the original categorical variables into a smaller set of new variables, called factors, that capture most of the inertia and information in the data. MCA is similar to PCA, but it is designed for nominal or ordinal variables, not for continuous or interval variables.

1: Principal Component Analysis - Amazon SageMaker

2: How to Use PCA for Data Visualization and Improved Performance in Machine Learning | by Pratik Shukla | Towards Data Science

3: Principal Component Analysis (PCA) for Feature Selection and some of its Pitfalls | by Nagesh Singh Chauhan | Towards Data Science

4: How to Reduce Dimensionality with PCA and Train a Support Vector Machine in Python | by James Briggs | Towards Data Science

5: Dimensionality Reduction and Its Applications | by Aniruddha Bhandari | Towards Data Science

6: Principal Component Analysis (PCA) in Python | by Susan Li | Towards Data Science

7: Feature Engineering for Machine Learning | by Dipanjan (DJ) Sarkar | Towards Data Science

8: Feature Engineering - How to Engineer Features and How to Get Good at It | by Parul Pandey | Towards Data Science

[Feature Scaling for Machine Learning: Understanding the Difference Between Normalization vs. Standardization | by Benjamin Obi Tayo Ph.D. | Towards Data Science]
[Why, How and When to Scale your Features | by George Seif | Towards Data Science]
[Normalization vs Dimensionality Reduction | by Saurabh Annadate | Towards Data Science]
[Multiple Correspondence Analysis - Amazon SageMaker]
[Multiple Correspondence Analysis (MCA) | by Raul Eulogio | Towards Data Science]

NEW QUESTION # 241

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 do?

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

Answer: C

NEW QUESTION # 242

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

Answer: A,D,F

Explanation:

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.

References:

* Detect sentiment from customer reviews using Amazon Comprehend | AWS Machine Learning Blog

* How to Apply Machine Learning to Customer Feedback

NEW QUESTION # 243

A Machine Learning Specialist needs to move and transform data in preparation for training. Some of the data needs to be processed in near-real time and other data can be moved hourly. There are existing Amazon EMR MapReduce jobs to clean and feature engineer to perform on the data. Which of the following services can feed data to the MapReduce jobs? (Select TWO)

- A. AWS DMS
- B. Amazon ES
- **C. AWS Data Pipeline**
- **D. Amazon Kinesis**

- E. Amazon Athena

Answer: C,D

Explanation:

Amazon Kinesis and AWS Data Pipeline are two services that can feed data to the Amazon EMR MapReduce jobs. Amazon Kinesis is a service that can ingest, process, and analyze streaming data in real time. Amazon Kinesis can be integrated with Amazon EMR to run MapReduce jobs on streaming data sources, such as web logs, social media, IoT devices, and clickstreams. Amazon Kinesis can handle data that needs to be processed in near-real time, such as for anomaly detection, fraud detection, or dashboarding. AWS Data Pipeline is a service that can orchestrate and automate data movement and transformation across various AWS services and on-premises data sources. AWS Data Pipeline can be integrated with Amazon EMR to run MapReduce jobs on batch data sources, such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon Redshift.

AWS Data Pipeline can handle data that can be moved hourly, such as for data warehousing, reporting, or machine learning.

AWS DMS is not a valid service name. AWS Database Migration Service (AWS DMS) is a service that can migrate data from various sources to various targets, but it does not support streaming data or MapReduce jobs.

Amazon Athena is a service that can query data stored in Amazon S3 using standard SQL, but it does not feed data to Amazon EMR or run MapReduce jobs.

Amazon ES is a service that provides a fully managed Elasticsearch cluster, which can be used for search, analytics, and visualization, but it does not feed data to Amazon EMR or run MapReduce jobs. References:

- * Using Amazon Kinesis with Amazon EMR - [Amazon EMR](#)
- * AWS Data Pipeline - [Amazon Web Services](#)
- * Using AWS Data Pipeline to Run Amazon EMR Jobs - [AWS Data Pipeline](#)

NEW QUESTION # 244

A company's Machine Learning Specialist needs to improve the training speed of a time-series forecasting model using TensorFlow. The training is currently implemented on a single-GPU machine and takes approximately 23 hours to complete. The training needs to be run daily.

The model accuracy is acceptable, but the company anticipates a continuous increase in the size of the training data and a need to update the model on an hourly, rather than a daily, basis. The company also wants to minimize coding effort and infrastructure changes.

What should the Machine Learning Specialist do to the training solution to allow it to scale for future demand?

- A. Do not change the TensorFlow code. Change the machine to one with a more powerful GPU to speed up the training.
- B. **Change the TensorFlow code to implement a Horovod distributed framework supported by Amazon SageMaker. Parallelize the training to as many machines as needed to achieve the business goals.**
- C. Move the training to Amazon EMR and distribute the workload to as many machines as needed to achieve the business goals.
- D. Switch to using a built-in AWS SageMaker DeepAR model. Parallelize the training to as many machines as needed to achieve the business goals.

Answer: B

NEW QUESTION # 245

.....

Among all substantial practice materials with similar themes, our AWS-Certified-Machine-Learning-Specialty practice materials win a majority of credibility for promising customers who are willing to make progress in this line. With excellent quality at attractive price, our AWS-Certified-Machine-Learning-Specialty Exam Questions get high demand of orders in this fierce market. You can just look at the data about the hot hit on the AWS-Certified-Machine-Learning-Specialty study braindumps everyday, and you will know that how popular our AWS-Certified-Machine-Learning-Specialty learning guide is.

New AWS-Certified-Machine-Learning-Specialty Exam Papers: https://www.itcerttest.com/AWS-Certified-Machine-Learning-Specialty_braindumps.html

- Reliable AWS-Certified-Machine-Learning-Specialty Exam Simulator □ AWS-Certified-Machine-Learning-Specialty Valid Exam Vce Free □ Pass4sure AWS-Certified-Machine-Learning-Specialty Dumps Pdf □ Open website ▷ www.pdfdumps.com ▷ and search for ➡ AWS-Certified-Machine-Learning-Specialty □□□ for free download □AWS-Certified-Machine-Learning-Specialty Vce Torrent
- AWS-Certified-Machine-Learning-Specialty Trustworthy Practice □ AWS-Certified-Machine-Learning-Specialty Verified

Answers □ Actual AWS-Certified-Machine-Learning-Specialty Test Pdf □ Open website ⇒ www.pdfvce.com ⇄ and search for ➔ AWS-Certified-Machine-Learning-Specialty □ □ □ for free download □ AWS-Certified-Machine-Learning-Specialty Verified Answers

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by Itcerttest: https://drive.google.com/open?id=1kjm_YldrxUhJKB2_kILBfRNWkS69MRQ