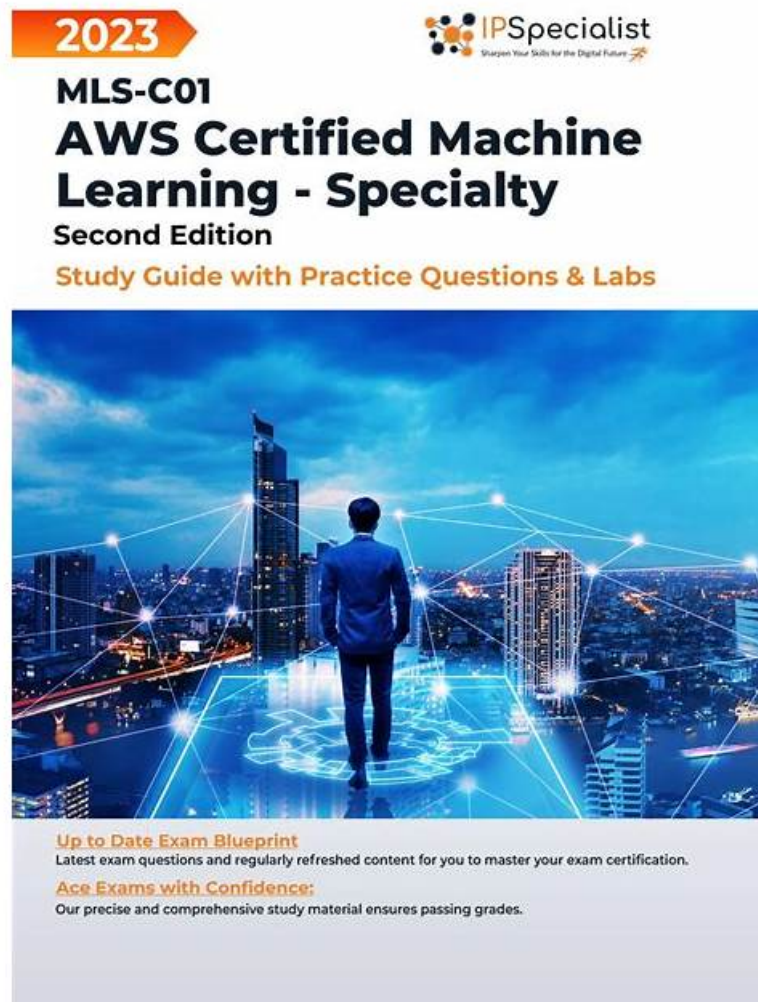


High Hit Rate Latest AWS-Certified-Machine-Learning-Specialty Study Guide, Training AWS-Certified-Machine-Learning-Specialty Materials



BTW, DOWNLOAD part of DumpsActual AWS-Certified-Machine-Learning-Specialty dumps from Cloud Storage:
<https://drive.google.com/open?id=16LcBxyhZQgaiBr2rqtahWytGwDRyqthL>

Regarding the process of globalization, every fighter who seeks a better life needs to keep pace with its tendency to meet challenges. AWS-Certified-Machine-Learning-Specialty certification is a stepping stone for you to stand out from the crowd. Nowadays, having knowledge of the AWS-Certified-Machine-Learning-Specialty study braindumps become widespread, if you grasp solid technological knowledge, you are sure to get a well-paid job and be promoted in a short time. According to our survey, those who have passed the exam with our AWS-Certified-Machine-Learning-Specialty Test Guide convincingly demonstrate their abilities of high quality, raise their professional profile, expand their network and impress prospective employers.

At the fork in the road, we always face many choices. When we choose job, job are also choosing us. Today's era is a time of fierce competition. Our AWS-Certified-Machine-Learning-Specialty exam question can make you stand out in the competition. Why is that? The answer is that you get the certificate. What certificate? Certificates are certifying that you have passed various qualifying examinations. Watch carefully you will find that more and more people are willing to invest time and energy on the AWS-Certified-Machine-Learning-Specialty Exam, because the exam is not achieved overnight, so many people are trying to find a suitable way.

>> **Latest AWS-Certified-Machine-Learning-Specialty Study Guide** <<

Training AWS-Certified-Machine-Learning-Specialty Materials | Latest AWS-Certified-Machine-Learning-Specialty Test Format

There is always a fear of losing AWS-Certified-Machine-Learning-Specialty exam and causes you loss of money and waste time on some unless materials. However, these risks will never exist in our AWS-Certified-Machine-Learning-Specialty exam materials. Your money and exam attempt is bound to award you a sure and definite success with 100% money back guarantee. You can claim for the refund of money if you do not succeed and achieve your target. Our AWS-Certified-Machine-Learning-Specialty exam materials have a most reliable guarantee. We ensure you that you will be paid back in full without any deduction and you can easily pass the AWS-Certified-Machine-Learning-Specialty Exam by using our AWS-Certified-Machine-Learning-Specialty dumps. Moreover, you will get all the updated AWS-Certified-Machine-Learning-Specialty questions with verified answers. If you want to prepare yourself for the real exam, then it is one of the most effect ways to improve your AWS-Certified-Machine-Learning-Specialty exam preparation level.

The AWS Certified Machine Learning - Specialty certification exam is intended for individuals with a minimum of one year of experience in designing and implementing machine learning solutions on AWS. AWS Certified Machine Learning - Specialty certification is ideal for data scientists, data engineers, software developers, and IT professionals who want to expand their skill set and demonstrate their expertise in the field of machine learning.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q22-Q27):

NEW QUESTION # 22

A retail company collects customer comments about its products from social media, the company website, and customer call logs. A team of data scientists and engineers wants to find common topics and determine which products the customers are referring to in their comments. The team is using natural language processing (NLP) to build a model to help with this classification.

Each product can be classified into multiple categories that the company defines. These categories are related but are not mutually exclusive. For example, if there is mention of "Sample Yogurt" in the document of customer comments, then "Sample Yogurt" should be classified as "yogurt," "snack," and "dairy product." The team is using Amazon Comprehend to train the model and must complete the project as soon as possible.

Which functionality of Amazon Comprehend should the team use to meet these requirements?

- A. Custom classification with multi-class mode
- B. Custom entity recognition
- **C. Custom classification with multi-label mode**
- D. Built-in models

Answer: C

Explanation:

The problem involves assigning multiple related, non-mutually exclusive categories (like "yogurt", "snack", and "dairy product") to customer comments referencing specific products. Amazon Comprehend provides custom classification capabilities, and in this case, the appropriate mode is multi-label, not multi-class.

"Use multi-label mode when each document can belong to more than one class. This is appropriate for situations where the categories are not mutually exclusive." This allows a single document to be tagged with several relevant labels simultaneously, which is precisely the use case described.

NEW QUESTION # 23

A Data Scientist needs to migrate an existing on-premises ETL process to the cloud. The current process runs at regular time intervals and uses PySpark to combine and format multiple large data sources into a single consolidated output for downstream processing. The Data Scientist has been given the following requirements for the cloud solution:

- * Combine multiple data sources
- * Reuse existing PySpark logic
- * Run the solution on the existing schedule
- * Minimize the number of servers that will need to be managed

Which architecture should the Data Scientist use to build this solution?

- A. Use Amazon Kinesis Data Analytics to stream the input data and perform realtime SQL queries against the stream to carry out the required transformations within the stream. Deliver the output results to a "processed" location in Amazon S3 that is accessible for downstream use.

- B. Write the raw data to Amazon S3 Schedule an AWS Lambda function to run on the existing schedule and process the input data from Amazon S3 Write the Lambda logic in Python and implement the existing PySpark logic to perform the ETL process Have the Lambda function output the results to a "processed" location in Amazon S3 that is accessible for downstream use
- C. Write the raw data to Amazon S3 Create an AWS Glue ETL job to perform the ETL processing against the input data Write the ETL job in PySpark to leverage the existing logic Create a new AWS Glue trigger to trigger the ETL job based on the existing schedule Configure the output target of the ETL job to write to a "processed" location in Amazon S3 that is accessible for downstream use.
- D. Write the raw data to Amazon S3 Schedule an AWS Lambda function to submit a Spark step to a persistent Amazon EMR cluster based on the existing schedule Use the existing PySpark logic to run the ETL job on the EMR cluster Output the results to a "processed" location in Amazon S3 that is accessible for downstream use

Answer: D

NEW QUESTION # 24

A company sells thousands of products on a public website and wants to automatically identify products with potential durability problems. The company has 1,000 reviews with date, star rating, review text, review summary, and customer email fields, but many reviews are incomplete and have empty fields. Each review has already been labeled with the correct durability result.

A machine learning specialist must train a model to identify reviews expressing concerns over product durability. The first model needs to be trained and ready to review in 2 days.

What is the MOST direct approach to solve this problem within 2 days?

- A. Train a custom classifier by using Amazon Comprehend.
- B. Build a recurrent neural network (RNN) in Amazon SageMaker by using Gluon and Apache MXNet.
- C. Train a built-in BlazingText model using Word2Vec mode in Amazon SageMaker.
- D. Use a built-in seq2seq model in Amazon SageMaker.

Answer: A

Explanation:

Explanation

The most direct approach to solve this problem within 2 days is to train a custom classifier by using Amazon Comprehend. Amazon Comprehend is a natural language processing (NLP) service that can analyze text and extract insights such as sentiment, entities, topics, and syntax. Amazon Comprehend also provides a custom classification feature that allows users to create and train a custom text classifier using their own labeled data.

The custom classifier can then be used to categorize any text document into one or more custom classes. For this use case, the custom classifier can be trained to identify reviews that express concerns over product durability as a class, and use the star rating, review text, and review summary fields as input features. The custom classifier can be created and trained using the Amazon Comprehend console or API, and does not require any coding or machine learning expertise. The training process is fully managed and scalable, and can handle large and complex datasets. The custom classifier can be trained and ready to review in 2 days or less, depending on the size and quality of the dataset.

The other options are not the most direct approaches because:

Option B: Building a recurrent neural network (RNN) in Amazon SageMaker by using Gluon and Apache MXNet is a more complex and time-consuming approach that requires coding and machine learning skills. RNNs are a type of deep learning models that can process sequential data, such as text, and learn long-term dependencies between tokens. Gluon is a high-level API for MXNet that simplifies the development of deep learning models. Amazon SageMaker is a fully managed service that provides tools and frameworks for building, training, and deploying machine learning models. However, to use this approach, the machine learning specialist would have to write custom code to preprocess the data, define the RNN architecture, train the model, and evaluate the results. This would likely take more than 2 days and involve more administrative overhead.

Option C: Training a built-in BlazingText model using Word2Vec mode in Amazon SageMaker is not a suitable approach for text classification. BlazingText is a built-in algorithm in Amazon SageMaker that provides highly optimized implementations of the Word2Vec and text classification algorithms. The Word2Vec algorithm is useful for generating word embeddings, which are dense vector representations of words that capture their semantic and syntactic similarities. However, word embeddings alone are not sufficient for text classification, as they do not account for the context and structure of the text documents. To use this approach, the machine learning specialist would have to combine the word embeddings with another classifier model, such as a logistic regression or a neural network, which would add more complexity and time to the solution.

Option D: Using a built-in seq2seq model in Amazon SageMaker is not a relevant approach for text classification. Seq2seq is a built-in algorithm in Amazon SageMaker that provides a sequence-to-sequence framework for neural machine translation based on MXNet. Seq2seq is a supervised learning algorithm that can generate an output sequence of tokens given an input sequence of

tokens, such as translating a sentence from one language to another. However, seq2seq is not designed for text classification, which requires assigning a label or a category to a text document, not generating another text sequence. To use this approach, the machine learning specialist would have to modify the seq2seq algorithm to fit the text classification task, which would be challenging and inefficient.

References:

Custom Classification - Amazon Comprehend

Build a Text Classification Model with Amazon Comprehend - AWS Machine Learning Blog

Recurrent Neural Networks - Gluon API
BlazingText Algorithm - Amazon SageMaker
Sequence-to-Sequence Algorithm - Amazon SageMaker

NEW QUESTION # 25

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 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.
- **B. Extract text from receipt images by using Amazon Texttract. 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 a deep learning OCR model from the AWS Marketplace. Use the NER deep learning model to extract entities.
- D. Extract text from receipt images by using Amazon Texttract. Use the Amazon SageMaker BlazingText algorithm to train on the text for entities and custom entities.

Answer: B

NEW QUESTION # 26

An interactive online dictionary wants to add a widget that displays words used in similar contexts. A Machine Learning Specialist is asked to provide word features for the downstream nearest neighbor model powering the widget.

What should the Specialist do to meet these requirements?

- A. Create one-hot word encoding vectors.
- B. Create word embedding factors that store edit distance with every other word.
- C. Produce a set of synonyms for every word using Amazon Mechanical Turk.
- **D. Download word embedding's pre-trained on a large corpus.**

Answer: D

Explanation:

Word embeddings are a type of dense representation of words, which encode semantic meaning in a vector form. These embeddings are typically pre-trained on a large corpus of text data, such as a large set of books, news articles, or web pages, and capture the context in which words are used. Word embeddings can be used as features for a nearest neighbor model, which can be used to find words used in similar contexts. Downloading pre-trained word embeddings is a good way to get started quickly and leverage the strengths of these representations, which have been optimized on a large amount of data. This is likely to result in more accurate and reliable features than other options like one-hot encoding, edit distance, or using Amazon Mechanical Turk to produce synonyms.

NEW QUESTION # 27

.....

The evergreen field of Amazon is so attractive that it provides non-stop possibilities for the one who passes the Amazon AWS-Certified-Machine-Learning-Specialty exam. So, to be there on top of the Amazon sector, earning the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) certification is essential. Because of using outdated AWS-

