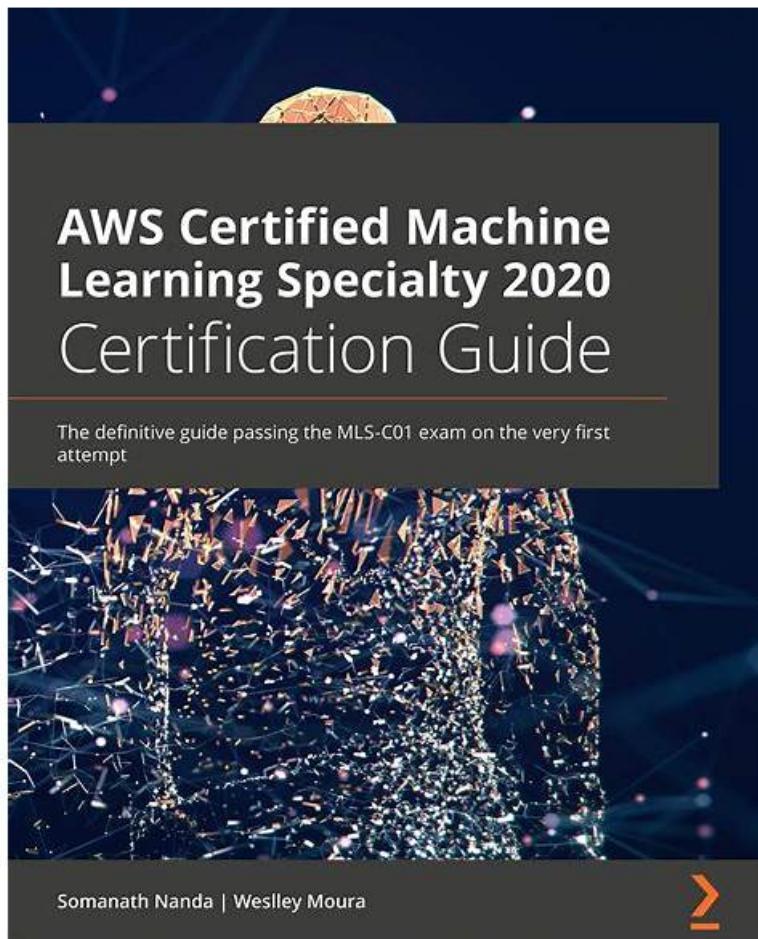


100% Pass 2026 MLS-C01: AWS Certified Machine Learning - Specialty Perfect Valid Exam Papers



BTW, DOWNLOAD part of ActualCollection MLS-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1J-q_dQMhl_-KHtanHHNSPDIjpa0VWtg

Purchasing a product may be a caucious thing for all of us, because we not only need to consider the performance of the product but also need to think about the things after purchasing. Our product will provide free demo for trying, and after you have bought the product of the MLS-C01 exam, we will send you the product by email in ten minutes after we have received the payment. After you bought the practice materials for the MLS-C01 Exam, if you have any question in the process of using, you can ask the service staff for help by email. Free update for having bought product is also available.

Amazon MLS-C01 certification exam is an excellent way for IT professionals to validate their skills in machine learning and differentiate themselves from their peers. By earning this certification, candidates can demonstrate to potential employers or clients that they have the necessary expertise to design, develop and deploy machine learning solutions on AWS. Additionally, the certification can open up new career opportunities and lead to higher salaries.

Amazon MLS-C01 Exam covers a wide range of topics, including data preparation, feature engineering, model training, evaluation and deployment, and machine learning implementation and operations. MLS-C01 exam also covers AWS services such as Amazon SageMaker, AWS Deep Learning AMIs, Amazon S3, Amazon EMR, and Amazon Redshift. MLS-C01 exam is designed to assess the candidate's ability to use these services to build, train, and deploy ML models.

>> MLS-C01 Valid Exam Papers <<

Exam MLS-C01 Simulator Fee - MLS-C01 Authentic Exam Questions

We hope this article has given you a good overview of the Amazon MLS-C01 Exam and what you can expect from it. As always,

we recommend you start preparing for your exam as early as possible to give yourself the best chance of success. ActualCollection offers a wide range of study materials and resources to help you prepare, including practice questions, dumps, and a study guide.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q163-Q168):

NEW QUESTION # 163

A Machine Learning Specialist is configuring automatic model tuning in Amazon SageMaker. When using the hyperparameter optimization feature, which of the following guidelines should be followed to improve optimization?

Choose the maximum number of hyperparameters supported by

- A. Amazon SageMaker to search the largest number of combinations possible
- **B. Use log-scaled hyperparameters to allow the hyperparameter space to be searched as quickly as possible**
- C. Specify a very large hyperparameter range to allow Amazon SageMaker to cover every possible value.
- D. Execute only one hyperparameter tuning job at a time and improve tuning through successive rounds of experiments

Answer: B

Explanation:

Explanation

Using log-scaled hyperparameters is a guideline that can improve the automatic model tuning in Amazon SageMaker. Log-scaled hyperparameters are hyperparameters that have values that span several orders of magnitude, such as learning rate, regularization parameter, or number of hidden units. Log-scaled hyperparameters can be specified by using a log-uniform distribution, which assigns equal probability to each order of magnitude within a range. For example, a log-uniform distribution between 0.001 and 1000 can sample values such as 0.001, 0.01, 0.1, 1, 10, 100, or 1000 with equal probability. Using log-scaled hyperparameters can allow the hyperparameter optimization feature to search the hyperparameter space more efficiently and effectively, as it can explore different scales of values and avoid sampling values that are too small or too large. Using log-scaled hyperparameters can also help avoid numerical issues, such as underflow or overflow, that may occur when using linear-scaled hyperparameters. Using log-scaled hyperparameters can be done by setting the ScalingType parameter to Logarithmic when defining the hyperparameter ranges in Amazon SageMaker.¹² The other options are not valid or relevant guidelines for improving the automatic model tuning in Amazon SageMaker. Choosing the maximum number of hyperparameters supported by Amazon SageMaker to search the largest number of combinations possible is not a good practice, as it can increase the time and cost of the tuning job and make it harder to find the optimal values. Amazon SageMaker supports up to 20 hyperparameters for tuning, but it is recommended to choose only the most important and influential hyperparameters for the model and algorithm, and use default or fixed values for the rest.³ Specifying a very large hyperparameter range to allow Amazon SageMaker to cover every possible value is not a good practice, as it can result in sampling values that are irrelevant or impractical for the model and algorithm, and waste the tuning budget. It is recommended to specify a reasonable and realistic hyperparameter range based on the prior knowledge and experience of the model and algorithm, and use the results of the tuning job to refine the range if needed.⁴ Executing only one hyperparameter tuning job at a time and improving tuning through successive rounds of experiments is not a good practice, as it can limit the exploration and exploitation of the hyperparameter space and make the tuning process slower and less efficient. It is recommended to use parallelism and concurrency to run multiple training jobs simultaneously and leverage the Bayesian optimization algorithm that Amazon SageMaker uses to guide the search for the best hyperparameter values.⁵

NEW QUESTION # 164

A bank has collected customer data for 10 years in CSV format. The bank stores the data in an on-premises server. A data science team wants to use Amazon SageMaker to build and train a machine learning (ML) model to predict churn probability. The team will use the historical data. The data scientists want to perform data transformations quickly and to generate data insights before the team builds a model for production.

Which solution will meet these requirements with the LEAST development effort?

- A. Upload the data into the SageMaker Data Wrangler console directly. Allow SageMaker and Amazon QuickSight to access the data that is in an Amazon S3 bucket. Perform data transformations in Data Wrangler and save the transformed data into a second S3 bucket. Use QuickSight to generate data insights.
- B. Upload the data into an Amazon S3 bucket. Allow SageMaker to access the data that is in the bucket. Import the data from the bucket into SageMaker Data Wrangler. Perform data transformations in Data Wrangler. Save the data into a second S3 bucket. Use a SageMaker Studio notebook to generate data insights.
- C. Upload the data into the SageMaker Data Wrangler console directly. Perform data transformations and generate insights within Data Wrangler.
- **D. Upload the data into an Amazon S3 bucket. Allow SageMaker to access the data that is in the bucket.**

Import the data from the S3 bucket into SageMaker Data Wrangler. Perform data transformations and generate insights within Data Wrangler.

Answer: D

Explanation:

To prepare and transform historical data efficiently with minimal setup, Amazon SageMaker Data Wrangler is the optimal tool. Data Wrangler simplifies data preprocessing and exploratory data analysis (EDA) by providing a graphical interface for transformations and insights. By first uploading the CSV data to Amazon S3, the data becomes easily accessible to SageMaker and can be imported directly into Data Wrangler.

Once in Data Wrangler, the team can perform required data transformations and generate insights in a single workflow, avoiding the need for additional tools like Amazon QuickSight or further notebook configuration.

This approach provides the simplest and most integrated solution for the data science team.

NEW QUESTION # 165

A company uses a long short-term memory (LSTM) model to evaluate the risk factors of a particular energy sector. The model reviews multi-page text documents to analyze each sentence of the text and categorize it as either a potential risk or no risk. The model is not performing well, even though the Data Scientist has experimented with many different network structures and tuned the corresponding hyperparameters.

Which approach will provide the MAXIMUM performance boost?

- A. Use gated recurrent units (GRUs) instead of LSTM and run the training process until the validation loss stops decreasing.
- B. Reduce the learning rate and run the training process until the training loss stops decreasing.
- C. Initialize the words by word2vec embeddings pretrained on a large collection of news articles related to the energy sector.
- D. Initialize the words by term frequency-inverse document frequency (TF-IDF) vectors pretrained on a large collection of news articles related to the energy sector.

Answer: C

Explanation:

Initializing the words by word2vec embeddings pretrained on a large collection of news articles related to the energy sector will provide the maximum performance boost for the LSTM model. Word2vec is a technique that learns distributed representations of words based on their co-occurrence in a large corpus of text. These representations capture semantic and syntactic similarities between words, which can help the LSTM model better understand the meaning and context of the sentences in the text documents. Using word2vec embeddings that are pretrained on a relevant domain (energy sector) can further improve the performance by reducing the vocabulary mismatch and increasing the coverage of the words in the text documents. References :

* AWS Machine Learning Specialty Exam Guide

* AWS Machine Learning Training - Text Classification with TF-IDF, LSTM, BERT: a comparison of performance

* AWS Machine Learning Training - Machine Learning - Exam Preparation Path

NEW QUESTION # 166

A company is launching a new product and needs to build a mechanism to monitor comments about the company and its new product on social media. The company needs to be able to evaluate the sentiment expressed in social media posts, and visualize trends and configure alarms based on various thresholds.

The company needs to implement this solution quickly, and wants to minimize the infrastructure and data science resources needed to evaluate the messages. The company already has a solution in place to collect posts and store them within an Amazon S3 bucket. What services should the data science team use to deliver this solution?

- A. Train a model in Amazon SageMaker by using the semantic segmentation algorithm to model the semantic content in the corpus of social media posts. Expose an endpoint that can be called by AWS Lambda. Trigger a Lambda function when objects are added to the S3 bucket to invoke the endpoint and record the sentiment in an Amazon DynamoDB table. Schedule a second Lambda function to query recently added records and send an Amazon Simple Notification Service (Amazon SNS) notification to notify analysts of trends.
- B. Trigger an AWS Lambda function when social media posts are added to the S3 bucket. Call Amazon Comprehend for each post to capture the sentiment in the message and record the sentiment in an Amazon DynamoDB table. Schedule a second Lambda function to query recently added records and send an Amazon Simple Notification Service (Amazon SNS) notification to notify analysts of trends.
- C. Train a model in Amazon SageMaker by using the BlazingText algorithm to detect sentiment in the corpus of social media

posts. Expose an endpoint that can be called by AWS Lambda. Trigger a Lambda function when posts are added to the S3 bucket to invoke the endpoint and record the sentiment in an Amazon DynamoDB table and in a custom Amazon CloudWatch metric. Use CloudWatch alarms to notify analysts of trends.

- D. Trigger an AWS Lambda function when social media posts are added to the S3 bucket. Call Amazon Comprehend for each post to capture the sentiment in the message and record the sentiment in a custom Amazon CloudWatch metric and in S3. Use CloudWatch alarms to notify analysts of trends.

Answer: D

Explanation:

The solution that uses Amazon Comprehend and Amazon CloudWatch is the most suitable for the given scenario. Amazon Comprehend is a natural language processing (NLP) service that can analyze text and extract insights such as sentiment, entities, topics, and syntax. Amazon CloudWatch is a monitoring and observability service that can collect and track metrics, create dashboards, and set alarms based on various thresholds. By using these services, the data science team can quickly and easily implement a solution to monitor the sentiment of social media posts without requiring much infrastructure or data science resources. The solution also meets the requirements of storing the sentiment in both S3 and CloudWatch, and using CloudWatch alarms to notify analysts of trends.

References:

Amazon Comprehend
Amazon CloudWatch

NEW QUESTION # 167

A company that manufactures mobile devices wants to determine and calibrate the appropriate sales price for its devices. The company is collecting the relevant data and is determining data features that it can use to train machine learning (ML) models. There are more than 1,000 features, and the company wants to determine the primary features that contribute to the sales price. Which techniques should the company use for feature selection? (Choose three.)

- A. Data augmentation
- B. Univariate selection
- C. Feature importance with a tree-based classifier
- D. Correlation plot with heat maps
- E. Data binning
- F. Data scaling with standardization and normalization

Answer: B,C,D

Explanation:

Feature selection is the process of selecting a subset of extracted features that are relevant and contribute to minimizing the error rate of a trained model. Some techniques for feature selection are:

Correlation plot with heat maps: This technique visualizes the correlation between features using a color-coded matrix. Features that are highly correlated with each other or with the target variable can be identified and removed to reduce redundancy and noise.

Univariate selection: This technique evaluates each feature individually based on a statistical test, such as chi-square, ANOVA, or mutual information, and selects the features that have the highest scores or p-values. This technique is simple and fast, but it does not consider the interactions between features.

Feature importance with a tree-based classifier: This technique uses a tree-based classifier, such as random forest or gradient boosting, to rank the features based on their importance in splitting the nodes. Features that have low importance scores can be dropped from the model. This technique can capture the non-linear relationships and interactions between features.

The other options are not techniques for feature selection, but rather for feature engineering, which is the process of creating, transforming, or extracting features from the original data. Feature engineering can improve the performance and interpretability of the model, but it does not reduce the number of features.

Data scaling with standardization and normalization: This technique transforms the features to have a common scale, such as zero mean and unit variance, or a range between 0 and 1. This technique can help some algorithms, such as k-means or logistic regression, to converge faster and avoid numerical instability, but it does not change the number of features.

Data binning: This technique groups the continuous features into discrete bins or categories based on some criteria, such as equal width, equal frequency, or clustering. This technique can reduce the noise and outliers in the data, and also create ordinal or nominal features that can be used for some algorithms, such as decision trees or naive Bayes, but it does not reduce the number of features.

Data augmentation: This technique generates new data from the existing data by applying some transformations, such as rotation, flipping, cropping, or noise addition. This technique can increase the size and diversity of the data, and help prevent overfitting, but it does not reduce the number of features.

References:

Feature engineering - Machine Learning Lens

Amazon SageMaker Autopilot now provides feature selection and the ability to change data types while creating an AutoML experiment Feature Selection in Machine Learning | Baeldung on Computer Science Feature Selection in Machine Learning: An easy Introduction

NEW QUESTION # 168

Our MLS-C01 exam guide are not only rich and varied in test questions, but also of high quality. A very high hit rate gives you a good chance of passing the final MLS-C01 exam. According to past statistics, 98 % - 99 % of the users who have used our MLS-C01 Study Materials can pass the exam successfully. So without doubt, you will be our nest passer as well as long as you buy our MLS-C01 practice braindumps.

Exam MLS-C01 Simulator Fee: <https://www.actualcollection.com/MLS-C01-exam-questions.html>

BONUS!!! Download part of ActualCollection MLS-C01 dumps for free: https://drive.google.com/open?id=1J-q_dQMlhI-KHtanHHNSPDIjpaOVWtg