

[2020-April-NewReal Exam Questions-Braindump2go MLS-C01 VCE and PDF Dumps 81Q Download

2020/April New Braindump2go MLS-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new MLS-C01 Exam Questions!

New QuestionA Machine Learning Specialist built an image classification deep learning model. However, the Specialist ran into an overfitting problem in which the training and testing accuracies were 99% and 75%, respectively. How should the Specialist address this issue and what is the reason behind it?
A. The learning rate should be increased because the optimization process was trapped at a local minimum.
B. The dropout rate at the flatten layer should be increased because the model is not generalized enough.
C. The dimensionality of dense layer next to the flatten layer should be increased because the model is not complex enough.
D. The epoch number should be increased because the optimization process was terminated before it reached the global minimum.
Answer: D
Explanation:
https://www.tensorflow.org/tutorials/keras/overfit_and_underfit

New QuestionA Machine Learning team uses Amazon SageMaker to train an Apache MXNet handwritten digit classifier model using a research dataset. The team wants to receive a notification when the model is overfitting. Auditors want to view the Amazon SageMaker log activity report to ensure there are no unauthorized API calls. What should the Machine Learning team do to address the requirements with the least amount of code and fewest steps?
A. Implement an AWS Lambda function to log Amazon SageMaker API calls to Amazon S3. Add code to push a custom metric to Amazon CloudWatch. Create an alarm in CloudWatch with Amazon SNS to receive a notification when the model is overfitting.
B. Use AWS CloudTrail to log Amazon SageMaker API calls to Amazon S3. Add code to push a custom metric to Amazon CloudWatch. Create an alarm in CloudWatch with Amazon SNS to receive a notification when the model is overfitting.
C. Implement an AWS Lambda function to log Amazon SageMaker API calls to AWS CloudTrail. Add code to push a custom metric to Amazon CloudWatch. Create an alarm in CloudWatch with Amazon SNS to receive a notification when the model is overfitting.
D. Use AWS CloudTrail to log Amazon SageMaker API calls to Amazon S3. Set up Amazon SNS to receive a notification when the model is overfitting.
Answer: C

New QuestionA Machine Learning Specialist is building a prediction model for a large number of features using linear models, such as linear regression and logistic regression. During exploratory data analysis, the Specialist observes that many features are highly correlated with each other. This may make the model unstable. What should be done to reduce the impact of having such a large number of features?
A. Perform one-hot encoding on highly correlated features.
B. Use matrix multiplication on highly correlated features.
C. Create a new feature space using principal component analysis (PCA)
D. Apply the Pearson correlation coefficient.
Answer: C

New QuestionA Machine Learning Specialist is implementing a full Bayesian network on a dataset that describes public transit in New York City. One of the random variables is discrete, and represents the number of minutes New Yorkers wait for a bus given that the buses cycle every 10 minutes, with a mean of 3 minutes. Which prior probability distribution should the ML Specialist use for this variable?
A. Poisson distribution
B. Uniform distribution
C. Normal distribution
D. Binomial distribution
Answer: D

New QuestionA Data Science team within a large company uses Amazon SageMaker notebooks to access data stored in Amazon S3 buckets. The IT Security team is concerned that internet-enabled notebook instances create a security vulnerability where malicious code running on the instances could compromise data privacy. The company mandates that all instances stay within a secured VPC with no internet access, and data communication traffic must stay within the AWS network. How should the Data Science team configure the notebook instance placement to meet these requirements?
A. Associate the Amazon SageMaker notebook with a private subnet in a VPC. Place the Amazon SageMaker endpoint and S3 buckets within the same VPC.
B. Associate the Amazon SageMaker notebook with a private subnet in a VPC. Use IAM policies to grant access to Amazon S3 and Amazon SageMaker.
C. Associate the Amazon SageMaker notebook with a private subnet in a VPC. Ensure the VPC has S3 VPC endpoints and Amazon SageMaker VPC endpoints attached to it.
D. Associate the Amazon SageMaker notebook with a private subnet in a VPC. Ensure the VPC has a NAT gateway and an associated security group allowing only outbound connections to Amazon S3 and Amazon SageMaker.
Answer: D

New QuestionA Machine Learning Specialist has created a deep learning neural network model that performs well on the training data but performs poorly on the test data. Which of the following methods should the Specialist consider using to correct this? (Choose three.)
A. Decrease regularization.
B. Increase regularization.
C. Increase dropout.
D. Decrease dropout.
E. Increase feature combinations.
F. Decrease feature combinations.
Answer: BDE

New QuestionA Data Scientist needs to create a serverless ingestion and analytics solution for high-velocity, real-time streaming data. The ingestion process must buffer and convert incoming records from JSON to a query-optimized, columnar format without data loss. The output datastore must be highly available, and Analysts must be able to run SQL queries against the data and connect to existing business intelligence dashboards. Which solution should the Data Scientist build to satisfy the requirements?
A. Create a schema in the AWS Glue Data Catalog of the incoming data format. Use an Amazon

Kinesis Data Firehose delivery stream to stream the data and transform the data to Apache Parquet or ORC format using the AWS Glue Data Catalog before delivering to Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena, and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.B. Write each JSON record to a staging location in Amazon S3. Use the S3 Put event to trigger an AWS Lambda function that transforms the data into Apache Parquet or ORC format and writes the data to a processed data location in Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena, and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.C. Write each JSON record to a staging location in Amazon S3. Use the S3 Put event to trigger an AWS Lambda function that transforms the data into Apache Parquet or ORC format and inserts it into an Amazon RDS PostgreSQL database. Have the Analysts query and run dashboards from the RDS database.D. Use Amazon Kinesis Data Analytics to ingest the streaming data and perform real-time SQL queries to convert the records to Apache Parquet before delivering to Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.**Answer: A**New QuestionAn online reseller has a large, multi-column dataset with one column missing 30% of its data. A Machine Learning Specialist believes that certain columns in the dataset could be used to reconstruct the missing data.Which reconstruction approach should the Specialist use to preserve the integrity of the dataset?A. Listwise deletionB. Last observation carried forwardC. Multiple imputationD. Mean substitution**Answer: C****Explanation:**

<https://worldwidescience.org/topicpages/i/imputing+missing+values.html>New QuestionA company is setting up an Amazon SageMaker environment. The corporate data security policy does not allow communication over the internet.How can the company enable the Amazon SageMaker service without enabling direct internet access to Amazon SageMaker notebook instances?A. Create a NAT gateway within the corporate VPC.B. Route Amazon SageMaker traffic through an on-premises network.C. Create Amazon SageMaker VPC interface endpoints within the corporate VPC.D. Create VPC peering with Amazon VPC hosting Amazon SageMaker.**Answer: A****Explanation:**<https://docs.aws.amazon.com/sagemaker/latest/dg/sagemaker-dg.pdf> (46)New QuestionMachine Learning Specialist is training a model to identify the make and model of vehicles in images. The Specialist wants to use transfer learning and an existing model trained on images of general objects. The Specialist collated a large custom dataset of pictures containing different vehicle makes and models.What should the Specialist do to initialize the model to re-train it with the custom data?A. Initialize the model with random weights in all layers including the last fully connected layer.B. Initialize the model with pre-trained weights in all layers and replace the last fully connected layer.C. Initialize the model with random weights in all layers and replace the last fully connected layer.D. Initialize the model with pre-trained weights in all layers including the last fully connected layer.**Answer: B**New QuestionAn office security agency conducted a successful pilot using 100 cameras installed at key locations within the main office. Images from the cameras were uploaded to Amazon S3 and tagged using Amazon Rekognition, and the results were stored in Amazon ES. The agency is now looking to expand the pilot into a full production system using thousands of video cameras in its office locations globally. The goal is to identify activities performed by non-employees in real timeWhich solution should the agency consider?A. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection of known employees, and alert when non-employees are detected.B. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Image to detect faces from a collection of known employees and alert when non-employees are detected.C. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection on each stream, and alert when non-employees are detected.D. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, run an AWS Lambda function to capture image fragments and then call Amazon Rekognition Image to detect faces from a collection of known employees, and alert when non-employees are detected.**Answer: D****Explanation:**

<https://aws.amazon.com/blogs/machine-learning/video-analytics-in-the-cloud-and-at-the-edge-with-aws-deeplens-and-kinesis-video-streams/>New QuestionA Marketing Manager at a pet insurance company plans to launch a targeted marketing campaign on social media to acquire new customers. Currently, the company has the following data in Amazon Aurora:- Profiles for all past and existing customers- Profiles for all past and existing insured pets- Policy-level information- Premiums received- Claims paidWhat steps should be taken to implement a machine learning model to identify potential new customers on social media?A. Use regression on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social mediaB. Use clustering on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social mediaC. Use a recommendation engine on customer profile data to understand key characteristics of consumer segments. Find

similar profiles on social media.D. Use a decision tree classifier engine on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media.**Answer: C**New QuestionA manufacturing company has a large set of labeled historical sales data. The manufacturer would like to predict how many units of a particular part should be produced each quarter.Which machine learning approach should be used to solve this problem?A. Logistic regressionB. Random Cut Forest (RCF)C. Principal component analysis (PCA)D. Linear regression**Answer: B**Resources From:1.2020 Latest Braindump2go MLS-C01 Exam Dumps (PDF & VCE) Free Share:<https://www.braindump2go.com/mls-c01.html>2.2020 Latest Braindump2go MLS-C01 PDF and MLS-C01 VCE Dumps Free Share:
<https://drive.google.com/drive/folders/1eX--L9LzE21hzqPikigeo1QoAGNWL4vd?usp=sharing>3.2020 Latest MLS-C01 Exam Questions from:<https://od.lk/fl/NDZfMTI1MDEyN18>Free Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!