

[New Exams! Valid DP-100 Exam VCE and PDF Free Download in Braindump2go[Q1-Q11]

July/2019 Braindump2go DP-100 Dumps with PDF and VCE New Updated Today! Following are some new DP-100 Exam Questions:
1. |2019 Latest DP-100 Exam Dumps (PDF & VCE) Instant Download: <https://www.braindump2go.com/dp-100.html>
2. |2019 Latest DP-100 Exam Questions & Answers Instant

Download: https://drive.google.com/drive/folders/1GRXSnO2A4MYVb3Cfs4F_07I919k9_LAD?usp=sharing
Case Study: 1
Overview You are a data scientist in a company that provides data science for professional sporting events. Models will be global and local market data to meet the following business goals:
? Understand sentiment of mobile device users at sporting events based on audio from crowd reactions.
? Access a user's tendency to respond to an advertisement.
? Customize styles of ads served on mobile devices.
? Use video to detect penalty events.
Current environment Requirements?
Media used for penalty event detection will be provided by consumer devices. Media may include images and videos captured during the sporting event and shared using social media. The images and videos will have varying sizes and formats.
The data available for model building comprises of seven years of sporting event media. The sporting event media includes: recorded videos, transcripts of radio commentary, and logs from related social media feeds captured during the sporting events.
Crowd sentiment will include audio recordings submitted by event attendees in both mono and stereo.
Formats.
Advertisements? Ad response models must be trained at the beginning of each event and applied during the sporting event.
Market segmentation models must optimize for similar ad response history.
Sampling must guarantee mutual and collective exclusivity local and global segmentation models that share the same features.
Local market segmentation models will be applied before determining a user's propensity to respond to an advertisement.
Data scientists must be able to detect model degradation and decay.
Ad response models must support non linear boundaries features.

The ad propensity model uses a cut threshold is 0.45 and retrains occur if weighted Kappa deviates from 0.1 +/-5%.
The ad propensity model uses cost factors shown in the following diagram:
The ad propensity model uses proposed cost factors shown in the following diagram:
Performance curves of current and proposed cost factor scenarios are shown in the following diagram:

Penalty detection and sentiment Findings?
Data scientists must build an intelligent solution by using multiple machine learning models for penalty event detection.
Data scientists must build notebooks in a local environment using automatic feature engineering and model building in machine learning pipelines.
Notebooks must be deployed to retrain by using Spark instances with dynamic worker allocation.
Notebooks must execute with the same code on new Spark instances to recode only the source of the data.
Global penalty detection models must be trained by using dynamic runtime graph computation during training.
Local penalty detection models must be written by using BrainScript.
Experiments for local crowd sentiment models must combine local penalty detection data.
Crowd sentiment models must identify known sounds such as cheers and known catch phrases. Individual crowd sentiment models will detect similar sounds.
All shared features for local models are continuous variables.
Shared features must use double precision. Subsequent layers must have aggregate running mean and standard deviation metrics Available.segments

During the initial weeks in production, the following was observed: Ad response rates declined. Drops were not consistent across ad styles. The distribution of features across training and production data are not consistent.
Analysis shows that of the 100 numeric features on user location and behavior, the 47 features that come from location sources are being used as raw features. A suggested experiment to remedy the bias and variance issue is to engineer 10 linearly uncorrected features.
Penalty detection and sentiment?
Initial data discovery shows a wide range of densities of target states in training data used for crowd sentiment models.

All penalty detection models show inference phases using a Stochastic Gradient Descent (SGD) are running too slow.
Audio samples show that the length of a catch phrase varies between 25%-47%, depending on region.
The performance of the global penalty detection models show lower variance but higher bias when comparing training and validation sets. Before implementing any feature changes, you must confirm the bias and variance using all training and validation cases.
Question: 1 You need to resolve the local machine learning pipeline performance issue. What should you do?
A. Increase Graphic Processing Units (GPUs).
B. Increase the learning rate.
C. Increase the training iterations.
D. Increase Central Processing Units (CPUs).
Answer: A

Question: 2 DRAG DROP You need to modify the inputs for the global penalty event model to address the bias and variance issue. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
Answer: Select the Location data Select the Behavior data Perform a primary component Analysis (PCA) Add a K-Means clustering module with 10 clusters. Bin the New data Build ratios.
Question: 3 You need to select an environment that will meet the business and data requirements. Which environment should you use?
A. Azure HDInsight with Spark MLlib
B. Azure Cognitive Services
C. Azure Machine Learning Studio
D. Microsoft Machine Learning Server
Answer: D

Question: 4 HOTSPOT You are using C-Support Vector classification to do a multi-class classification with

an unbalanced training dataset. The C-Support Vector classification using Python code shown below: You need to evaluate the C-Support Vector classification code. Which evaluation statement should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point. Answer: Explanation: Box 1: Automatically adjust weights inversely proportional to class frequencies in the input data. The 'balanced' mode uses the values of y to automatically adjust weights inversely proportional to class frequencies in the input data as $n_samples / (n_classes * np.bincount(y))$. Box 2: Penalty parameter C of the error term. Parameter: C : float, optional (default=1.0) Penalty parameter C of the error term. References:

<https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html> Question: 5 DRAG DROP You need to define a process for penalty event detection. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. Answer: Question: 6 DRAG DROP You need to define an evaluation strategy for the crowd sentiment models. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. Answer: Explanation: Scenario: Experiments for local crowd sentiment models must combine local penalty detection data. Crowd sentiment models must identify known sounds such as cheers and known catch phrases. Individual crowd sentiment models will detect similar sounds. Note: Evaluate the change in correlation between model error rate and centroid distance. In machine learning, a nearest centroid classifier or nearest prototype classifier is a classification model that assigns to observations the label of the class of training samples whose mean (centroid) is closest to the observation. References:

https://en.wikipedia.org/wiki/Nearest_centroid_classifier

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/sweep-clustering> Question: 7 HOTSPOT You need to build a feature extraction strategy for the local models. How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point. Answer: Question: 8 You need to implement a scaling strategy for the local penalty detection data. Which normalization type should you use? A. Streaming B. Weight C. Batch D. Cosine Answer: C Explanation: Post batch normalization statistics (PBN) is the Microsoft Cognitive Toolkit (CNTK) version of how to evaluate the population mean and variance of Batch Normalization which could be used in inference. Original Paper. In CNTK, custom networks are defined using the BrainScriptNetworkBuilder and described in the CNTK network description language "BrainScript." Scenario: Local penalty detection models must be written by using BrainScript. References:

<https://docs.microsoft.com/en-us/cognitive-toolkit/post-batch-normalization-statistics> Question: 9 HOTSPOT You need to use the Python language to build a sampling strategy for the global penalty detection models. How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point. Answer: Explanation: Box 1: import torch as deeplearninglib Box 2: ..DistributedSampler(Sampler).. DistributedSampler(Sampler): Sampler that restricts data loading to a subset of the dataset. It is especially useful in conjunction with `class:`torch.nn.parallel.DistributedDataParallel``. In such case, each process can pass a DistributedSampler instance as a DataLoader sampler, and load a subset of the original dataset that is exclusive to it. Scenario: Sampling must guarantee mutual and collective exclusivity between local and global segmentation models that share the same features. Box 3: optimizer = deeplearninglib.train.GradientDescentOptimizer(learning_rate=0.10) Incorrect Answers: ..SGD.. Scenario: All penalty detection models show inference phases using a Stochastic Gradient Descent (SGD) are running too slow. Box 4: .. nn.parallel.DistributedDataParallel.. DistributedSampler(Sampler): The sampler that restricts data loading to a subset of the dataset. It is especially useful in conjunction with `class:`torch.nn.parallel.DistributedDataParallel``. References:

<https://github.com/pytorch/pytorch/blob/master/torch/utils/data/distributed.py> Question: 10 You need to implement a feature engineering strategy for the crowd sentiment local models. What should you do? A. Apply an analysis of variance (ANOVA). B. Apply a Pearson correlation coefficient. C. Apply a Spearman correlation coefficient. D. Apply a linear discriminant analysis. Answer: D Explanation: The linear discriminant analysis method works only on continuous variables, not categorical or ordinal variables. Linear discriminant analysis is similar to analysis of variance (ANOVA) in that it works by comparing the means of the variables. Scenario: Data scientists must build notebooks in a local environment using automatic feature engineering and model building in machine learning pipelines. Experiments for local crowd sentiment models must combine local penalty detection data. All shared features for local models are continuous variables. Incorrect Answers: B: The Pearson correlation coefficient, sometimes called Pearson's R test, is a statistical value that measures the linear relationship between two variables. By examining the coefficient values, you can infer something about the strength of the relationship between the two variables, and whether they are positively correlated or negatively correlated. C: Spearman's correlation coefficient is designed for use with non-parametric and non-normally distributed data. Spearman's coefficient is a nonparametric measure of statistical dependence between two variables, and is sometimes denoted by the Greek letter rho. The Spearman's coefficient expresses the degree to which

two variables are monotonically related. It is also called Spearman rank correlation, because it can be used with ordinal variables.

References:<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/fisher-linear-discriminant-analysis><https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/compute-linear-correlation>

Question: 11 DRAG DROP You need to define a modeling strategy for ad response. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. Answer: Explanation: Step 1: Implement a K-Means Clustering model Step 2: Use the cluster as a feature in a Decision jungle model. Decision jungles are non-parametric models, which can represent non-linear decision boundaries. Step 3: Use the raw score as a feature in a Score Matchbox Recommender model The goal of creating a recommendation system is to recommend one or more "items" to "users" of the system. Examples of an item could be a movie, restaurant, book, or song. A user could be a person, group of persons, or other entity with item preferences. Scenario: Ad response rated declined. Ad response models must be trained at the beginning of each event and applied during the sporting event. Market segmentation models must optimize for similar ad response history. Ad response models must support non-linear boundaries of features. References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/multiclass-decision-jungle>

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/score-matchbox-recommender>

!!!RECOMMEND!!! 1. |2019 Latest DP-100 Exam Dumps (PDF & VCE) Instant

Download: <https://www.braindump2go.com/dp-100.html> 2. |2019 Latest DP-100 Study Guide Video Instant Download: YouTube

Video: [YouTube.com/watch?v=1VR23KIX-IE](https://www.youtube.com/watch?v=1VR23KIX-IE)