

[28/July/2018 Updated Valid 80q Professional Data Engineer Practice Test with Free PDF Study Guide Download (Part B)

New Updated Professional Data Engineer Exam Questions from PassLeader Professional Data Engineer PDF dumps! Welcome to download the newest PassLeader Professional Data Engineer VCE dumps: <https://www.passleader.com/professional-data-engineer.html> (80 Q&As) Keywords: Professional Data Engineer exam dumps, Professional Data Engineer exam questions, Professional Data Engineer VCE dumps, Professional Data Engineer PDF dumps, Professional Data Engineer practice tests, Professional Data Engineer study guide, Professional Data Engineer braindumps, **Google Cloud Certifications: Professional Data Engineer Exam** P.S. New Professional Data Engineer dumps PDF: <https://drive.google.com/open?id=1m882ngsiRO1BOHineV4IQUv9jgF5Lpue> P.S. New Professional Cloud Architect dumps PDF: <https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2> NEW QUESTION 16 Why do you need to split a machine learning dataset into training data and test data? A. So you can try two different sets of features. B. To make sure your model is generalized for more than just the training data. C. To allow you to create unit tests in your code. D. So you can use one dataset for a wide model and one for a deep model. Answer: B Explanation: The flaw with evaluating a predictive model on training data is that it does not inform you on how well the model has generalized to new unseen data. A model that is selected for its accuracy on the training dataset rather than its accuracy on an unseen test dataset is very likely to have lower accuracy on an unseen test dataset. The reason is that the model is not as generalized. It has specialized to the structure in the training dataset. This is called overfitting. <https://machinelearningmastery.com/a-simple-intuition-for-overfitting/> NEW QUESTION 17 Which of these numbers are adjusted by a neural network as it learns from a training dataset? (Choose two.) A. Weights B. Biases C. Continuous features D. Input values Answer: AB Explanation: A neural network is a simple mechanism that's implemented with basic math. The only difference between the traditional programming model and a neural network is that you let the computer determine the parameters (weights and bias) by learning from training datasets. <https://cloud.google.com/blog/big-data/2016/07/understanding-neural-networks-with-tensorflow-playground> NEW QUESTION 18 The CUSTOM tier for Cloud Machine Learning Engine allows you to specify the number of which types of cluster nodes? A. Workers B. Masters, workers, and parameter servers C. Workers and parameter servers D. Parameter servers Answer: C Explanation: The CUSTOM tier is not a set tier, but rather enables you to use your own cluster specification. When you use this tier, set values to configure your processing cluster according to these guidelines: - You must set TrainingInput.masterType to specify the type of machine to use for your master node. - You may set TrainingInput.workerCount to specify the number of workers to use. - You may set TrainingInput.parameterServerCount to specify the number of parameter servers to use. - You can specify the type of machine for the master node, but you can't specify more than one master node. https://cloud.google.com/ml-engine/docs/training-overview#job_configuration_parameters NEW QUESTION 19 Which software libraries are supported by Cloud Machine Learning Engine? A. Theano and TensorFlow B. Theano and Torch C. TensorFlow D. TensorFlow and Torch Answer: C Explanation: Cloud ML Engine mainly does two things: - Enables you to train machine learning models at scale by running TensorFlow training applications in the cloud. - Hosts those trained models for you in the cloud so that you can use them to get predictions about new data. https://cloud.google.com/ml-engine/docs/technical-overview#what_it_does NEW QUESTION 20 Which TensorFlow function can you use to configure a categorical column if you don't know all of the possible values for that column? A. categorical_column_with_vocabulary_list B. categorical_column_with_hash_bucket C. categorical_column_with_unknown_values D. sparse_column_with_keys Answer: B Explanation: If you know the set of all possible feature values of a column and there are only a few of them, you can use categorical_column_with_vocabulary_list. Each key in the list will get assigned an auto-incremental ID starting from 0. What if we don't know the set of possible values in advance? Not a problem. We can use categorical_column_with_hash_bucket instead. What will happen is that each possible value in the feature column occupation will be hashed to an integer ID as we encounter them in training. <https://www.tensorflow.org/tutorials/wide> NEW QUESTION 21 Which of the following statements about the Wide & Deep Learning model are true? (Choose two.) A. The wide model is used for memorization, while the deep model is used for generalization. B. A good use for the wide and deep model is a recommender system. C. The wide model is used for generalization, while the deep model is used for memorization. D. A good use for the

wide and deep model is a small-scale linear regression problem. Answer: AB Explanation: Can we teach computers to learn like humans do, by combining the power of memorization and generalization? It's not an easy question to answer, but by jointly training a wide linear model (for memorization) alongside a deep neural network (for generalization), one can combine the strengths of both to bring us one step closer. At Google, we call it Wide & Deep Learning. It's useful for generic large-scale regression and classification problems with sparse inputs (categorical features with a large number of possible feature values), such as recommender systems, search, and ranking problems.

<https://research.googleblog.com/2016/06/wide-deep-learning-better-together-with.html> NEW QUESTION 22 To run a TensorFlow training job on your own computer using Cloud Machine Learning Engine, what would your command start with? A. gcloud ml-engine local train B. gcloud ml-engine jobs submit training C. gcloud ml-engine jobs submit training local D. you can't run a TensorFlow program on your own computer using Cloud ML Engine Answer: A Explanation: gcloud ml-engine local train - run a Cloud ML Engine training job locally: This command runs the specified module in an environment similar to that of a live Cloud ML Engine Training Job. This is especially useful in the case of testing distributed models, as it allows you to validate that you are properly interacting with the Cloud ML Engine cluster configuration. <https://cloud.google.com/sdk/gcloud/reference/ml-engine/local/train>

NEW QUESTION 23 If you want to create a machine learning model that predicts the price of a particular stock based on its recent price history, what type of estimator should you use? A. Unsupervised learning B. Regressor C. Classifier D. Clustering estimator Answer: B Explanation: B: Regression is the supervised learning task for modeling and predicting continuous, numeric variables. Examples include predicting real-estate prices, stock price movements, or student test scores. C: Classification is the supervised learning task for modeling and predicting categorical variables. Examples include predicting employee churn, email spam, financial fraud, or student letter grades. D: Clustering is an unsupervised learning task for finding natural groupings of observations (i.e. clusters) based on the inherent structure within your dataset. Examples include customer segmentation, grouping similar items in e-commerce, and social network analysis. <https://elitedatascience.com/machine-learning-algorithms>

NEW QUESTION 24 Suppose you have a dataset of images that are each labeled as to whether or not they contain a human face. To create a neural network that recognizes human faces in images using this labeled dataset, what approach would likely be the most effective? A. Use K-means Clustering to detect faces in the pixels. B. Use feature engineering to add features for eyes, noses, and mouths to the input data. C. Use deep learning by creating a neural network with multiple hidden layers to automatically detect features of faces. D. Build a neural network with an input layer of pixels, a hidden layer, and an output layer with two categories. Answer: C Explanation: Traditional machine learning relies on shallow nets, composed of one input and one output layer, and at most one hidden layer in between. More than three layers (including input and output) qualifies as 'deep' learning. So deep is a strictly defined, technical term that means more than one hidden layer. In deep-learning networks, each layer of nodes trains on a distinct set of features based on the previous layer's output. The further you advance into the neural net, the more complex the features your nodes can recognize, since they aggregate and recombine features from the previous layer. A neural network with only one hidden layer would be unable to automatically recognize high-level features of faces, such as eyes, because it wouldn't be able to "build" these features using previous hidden layers that detect low-level features, such as lines. Feature engineering is difficult to perform on raw image data. K-means Clustering is an unsupervised learning method used to categorize unlabeled data. <https://deeplearning4j.org/neuralnet-overview>

NEW QUESTION 25 What are two of the characteristics of using online prediction rather than batch prediction? (Choose two.) A. It is optimized to handle a high volume of data instances in a job and to run more complex models. B. Predictions are returned in the response message. C. Predictions are written to output files in a Cloud Storage location that you specify. D. It is optimized to minimize the latency of serving predictions. Answer: BD Explanation: Online prediction: - Optimized to minimize the latency of serving predictions. - Predictions returned in the response message. - Optimized to handle a high volume of instances in a job and to run more complex models. - Predictions written to output files in a Cloud Storage location that you specify.

https://cloud.google.com/ml-engine/docs/tensorflow/prediction-overview#online_prediction_versus_batch_prediction NEW QUESTION 26 Which of these are examples of a value in a sparse vector? (Choose two.) A. [0, 5, 0, 0, 0] B. [0, 0, 0, 1, 0, 0, 1] C. [0, 1] D. [1, 0, 0, 0, 0, 0, 0] Answer: CD Explanation: Categorical features in linear models are typically translated into a sparse vector in which each possible value has a corresponding index or id. For example, if there are only three possible eye colors you can represent 'eye_color' as a length 3 vector: 'brown' would become [1, 0, 0], 'blue' would become [0, 1, 0] and 'green' would become [0, 0, 1]. These vectors

are called "sparse" because they may be very long, with many zeros, when the set of possible values is very large (such as all English words). [0, 0, 0, 1, 0, 0, 1] is not a sparse vector because it has two 1s in it. A sparse vector contains only a single 1. [0, 5, 0, 0, 0, 0] is not a sparse vector because it has a 5 in it. Sparse vectors only contain 0s and 1s.

https://www.tensorflow.org/tutorials/linear#feature_columns_and_transformations NEW QUESTION 27 How can you get a neural network to learn about relationships between categories in a categorical feature? A. Create a multi-hot column B. Create a one-hot column C. Create a hash bucket

D. Create an embedding column Answer: D Explanation: There are two problems with one-hot encoding. First, it has high dimensionality, meaning that instead of having just one value, like a continuous feature, it has many values, or dimensions. This makes computation more time-consuming, especially if a feature has a very large number of categories. The second problem is that it doesn't encode any relationships between the categories. They are completely independent from each other, so the network has no way of knowing which ones are similar to each other. Both of these problems can be solved by representing a categorical feature with an embedding column. The idea is that each category has a smaller vector with, let's say, 5 values in it. But unlike a one-hot vector, the values are not usually 0. The values are weights, similar to the weights that are used for basic features in a neural network. The difference is that each category has a set of weights (5 of them in this case). You can think of each value in the embedding vector as a feature of the category. So, if two categories are very similar to each other, then their embedding vectors should be very similar too.

<https://cloudacademy.com/google/introduction-to-google-cloud-machine-learning-engine-course/a-wide-and-deep-model.html>

NEW QUESTION 28 If a dataset contains rows with individual people and columns for year of birth, country, and income, how many of the columns are continuous and how many are categorical? A. 1 continuous and 2 categorical B. 3 categorical C. 3 continuous D. 2 continuous and 1 categorical Answer: D Explanation: The columns can be grouped into two types -- categorical and continuous columns: 1. A column is called categorical if its value can only be one of the categories in a finite set. For example, the native country of a person (U.S., India, Japan, etc.) or the education level (high school, college, etc.) are categorical columns. 2. A column is called continuous if its value can be any numerical value in a continuous range. For example, the capital gain of a person (e.g. \$14,084) is a continuous column. Year of birth and income are continuous columns. Country is a categorical column. You could use bucketization to turn year of birth and/or income into categorical features, but the raw columns are continuous.

https://www.tensorflow.org/tutorials/wide#reading_the_census_data NEW QUESTION 29 Which of the following are examples of hyperparameters? (Choose two.) A. Number of hidden layers B. Number of nodes in each hidden layer C. Biases D. Weights Answer: AB Explanation: If model parameters are variables that get adjusted by training with existing data, your hyperparameters are the variables about the training process itself. For example, part of setting up a deep neural network is deciding how many "hidden" layers of nodes to use between the input layer and the output layer, as well as how many nodes each layer should use. These variables are not directly related to the training data at all. They are configuration variables. Another difference is that parameters change during a training job, while the hyperparameters are usually constant during a job. Weights and biases are variables that get adjusted during the training process, so they are not hyperparameters. <https://cloud.google.com/ml-engine/docs/hyperparameter-tuning-overview>

NEW QUESTION 30 Which of the following are feature engineering techniques? (Choose two.) A. Hidden feature layers B. Feature prioritization C. Crossed feature columns D. Bucketization of a continuous feature Answer: CD Explanation: Selecting and crafting the right set of feature columns is key to learning an effective model. Bucketization is a process of dividing the entire range of a continuous feature into a set of consecutive bins/buckets, and then converting the original numerical feature into a bucket ID (as a categorical feature) depending on which bucket that value falls into. Using each base feature column separately may not be enough to explain the data. To learn the differences between different feature combinations, we can add crossed feature columns to the model.

https://www.tensorflow.org/tutorials/wide#selecting_and_engineering_features_for_the_model Download the newest PassLeader Professional Data Engineer dumps from passleader.com now! 100% Pass Guarantee! Professional Data Engineer PDF dumps & Professional Data Engineer VCE dumps: <https://www.passleader.com/professional-data-engineer.html> (80 Q&As) (New Questions Are 100% Available and Wrong Answers Have Been Corrected! Free VCE simulator!) P.S. New Professional Data Engineer dumps PDF: <https://drive.google.com/open?id=1m882ngsiRO1BOHineV4IQUv9jgF5Lpue> P.S. New Professional Cloud Architect dumps PDF: <https://drive.google.com/open?id=19jt3GbCmVz-pmGbZv8zjAu0NH7423IQ2>