PROFESSIONAL-MACHINE-LEARNING-ENGINEER Q&As

Professional Machine Learning Engineer

Pass Google PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/professional-machine-learning-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

You are developing a classification model to support predictions for your company\\'s various products. The dataset you were given for model development has class imbalance You need to minimize false positives and false negatives What evaluation metric should you use to properly train the model?

A. F1 score

B. Recall

C. Accuracy

D. Precision

Correct Answer: A

QUESTION 2

While conducting an exploratory analysis of a dataset, you discover that categorical feature A has substantial predictive power, but it is sometimes missing. What should you do?

A. Drop feature A if more than 15% of values are missing. Otherwise, use feature A as-is.

B. Compute the mode of feature A and then use it to replace the missing values in feature A.

C. Replace the missing values with the values of the feature with the highest Pearson correlation with feature A.

D. Add an additional class to categorical feature A for missing values. Create a new binary feature that indicates whether feature A is missing.

Correct Answer: D

https://developers.google.com/machine-learning/testing-debugging/common/data-errors

QUESTION 3

You are designing an architecture with a serverless ML system to enrich customer support tickets with informative metadata before they are routed to a support agent. You need a set of models to predict ticket priority, predict ticket resolution time, and perform sentiment analysis to help agents make strategic decisions when they process support requests. Tickets are not expected to have any domain-specific terms or jargon.

The proposed architecture has the following flow:

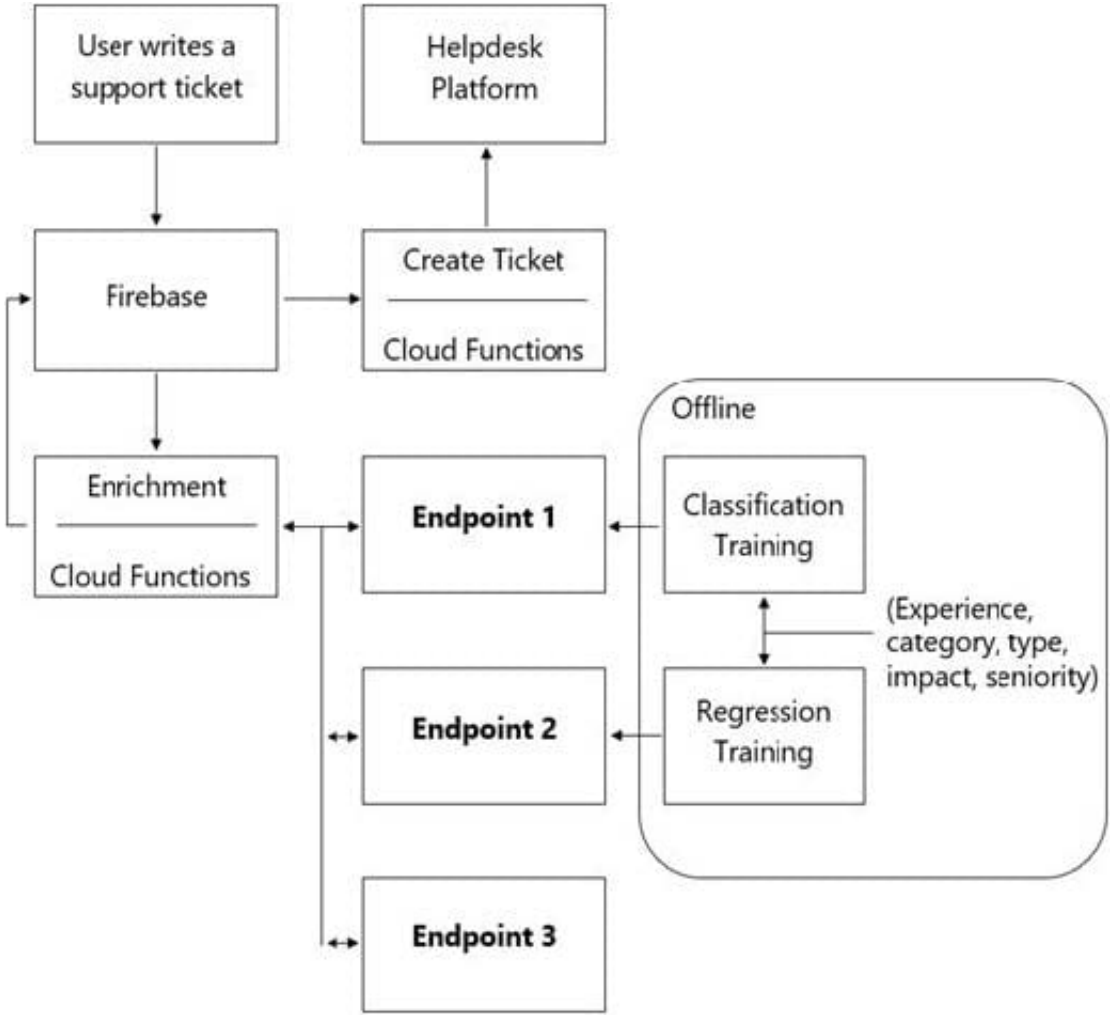

Which endpoints should the Enrichment Cloud Functions call?

A. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Vision

B. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Natural Language

C. 1 = AI Platform, 2 = AI Platform, 3 = Cloud Natural Language API

D. 1 = Cloud Natural Language API, 2 = AI Platform, 3 = Cloud Vision API

Correct Answer: C

https://cloud.google.com/architecture/architecture-of-a-serverless-ml-model#architecture

The architecture has the following flow:

A user writes a ticket to Firebase, which triggers a Cloud Function.

-The Cloud Function calls 3 different endpoints to enrich the ticket:

-An AI Platform endpoint, where the function can predict the priority.

-An AI Platform endpoint, where the function can predict the resolution time.

-The Natural Language API to do sentiment analysis and word salience.

-For each reply, the Cloud Function updates the Firebase real-time database.

-The Cloud Function then creates a ticket into the helpdesk platform using the RESTful API.

QUESTION 4

You work on a team in a data center that is responsible for server maintenance. Your management team wants you to build a predictive maintenance solution that uses monitoring data to detect potential server failures. Incident data has not been labeled yet. What should you do first?

A. Train a time-series model to predict the machines\\' performance values. Configure an alert if a machine\\'s actual performance values significantly differ from the predicted performance values.

B. Develop a simple heuristic (e.g., based on z-score) to label the machines\\' historical performance data. Use this heuristic to monitor server performance in real time.

C. Develop a simple heuristic (e.g., based on z-score) to label the machines\\' historical performance data. Train a model to predict anomalies based on this labeled dataset.

D. Hire a team of qualified analysts to review and label the machines\\' historical performance data. Train a model based on this manually labeled dataset.

Correct Answer: B

https://developers.google.com/machine-learning/guides/rules-of-ml

QUESTION 5

You work at a subscription-based company. You have trained an ensemble of trees and neural networks to predict customer churn, which is the likelihood that customers will not renew their yearly subscription. The average prediction is a 15% churn rate, but for a particular customer the model predicts that they are 70% likely to churn. The customer has a product usage history of 30%, is located in New York City, and became a customer in 1997. You need to explain the difference between the actual prediction, a 70% churn rate, and the average prediction. You want to use Vertex Explainable AI. What should you do?

A. Train local surrogate models to explain individual predictions.

B. Configure sampled Shapley explanations on Vertex Explainable AI.

C. Configure integrated gradients explanations on Vertex Explainable AI.

D. Measure the effect of each feature as the weight of the feature multiplied by the feature value.

Correct Answer: B

https://cloud.google.com/vertex-ai/docs/explainable-ai/overview#compare-methods