



Across

3. In machine learning, a mechanism for bucketing categorical data
6. The primary algorithm for performing gradient descent on neural networks.
9. Abbreviation for independently and identically distributed
12. The more common label in a class-imbalanced dataset.
13. Applying a constraint to an algorithm to ensure one or more definitions of fairness are satisfied.
18. A process used, as part of training, to

Down

1. A post-prediction adjustment, typically to account for prediction bias.
2. A TensorFlow programming environment in which operations run immediately.
4. Obtaining an understanding of data by considering samples, measurement, and visualization.
5. An ensemble approach to finding the decision tree that best fits the training data
7. state-action value function
8. Loss function based on the absolute value of the difference between the values that a model is predicting and the actual values of the labels

evaluate the quality of a machine learning model using the validation set.

19. A coefficient for a feature in a linear model, or an edge in a deep network.
20. A column-oriented data analysis API.
21. Abbreviation for generative adversarial network

10. A metric that your algorithm is trying to optimize.
11. The recommended format for saving and recovering TensorFlow models.
14. A statistical way of comparing two (or more) techniques, typically an incumbent against a new rival.
15. When one number in your model becomes a NaN during training, which causes many or all other numbers in your model to eventually become a NaN.
16. Q-learning In reinforcement learning, implementing Q-learning by using a table to store the Q-functions
17. A popular Python machine learning API