How do we compare two ROC curves and what is an example of this comparison?

Authored by

stats writer

June 23, 2024

ROC (Receiver Operating Characteristic) curves are used to evaluate the performance of binary classifiers. They plot the true positive rate against the false positive rate at various threshold settings. These curves can be used to compare the performance of different classifiers or the same classifier with different parameters.

To compare two ROC curves, we first look at their shape. A curve that is closer to the top-left corner of the graph indicates a better performance as it has a higher true positive rate and a lower false positive rate. We can also use the area under the curve (AUC) as a measure of comparison. A higher AUC value signifies a better performance of the classifier.

An example of this comparison could be in the medical field, where we want to compare the performance of two diagnostic tests for a disease. The ROC curves for each test can be plotted and compared to determine which test has a higher accuracy in correctly identifying the presence of the disease. This comparison can help in selecting the more reliable test for diagnosing the disease.

Compare Two ROC Curves (With Example)

One way to visualize the performance of in machine learning is by creating a ROC curve, which stands for "receiver operating characteristic" curve.

This type of curve displays the sensitivity and specificity of a classification model:

Sensitivity: The probability that the model predicts a positive outcome for an observation when the outcome is indeed positive.Specificity: The probability that the model predicts a negative outcome for an observation when the outcome is indeed negative.

The x-axis of a ROC curve represents (1- Specificity) and the y-axis represents the Sensitivity:

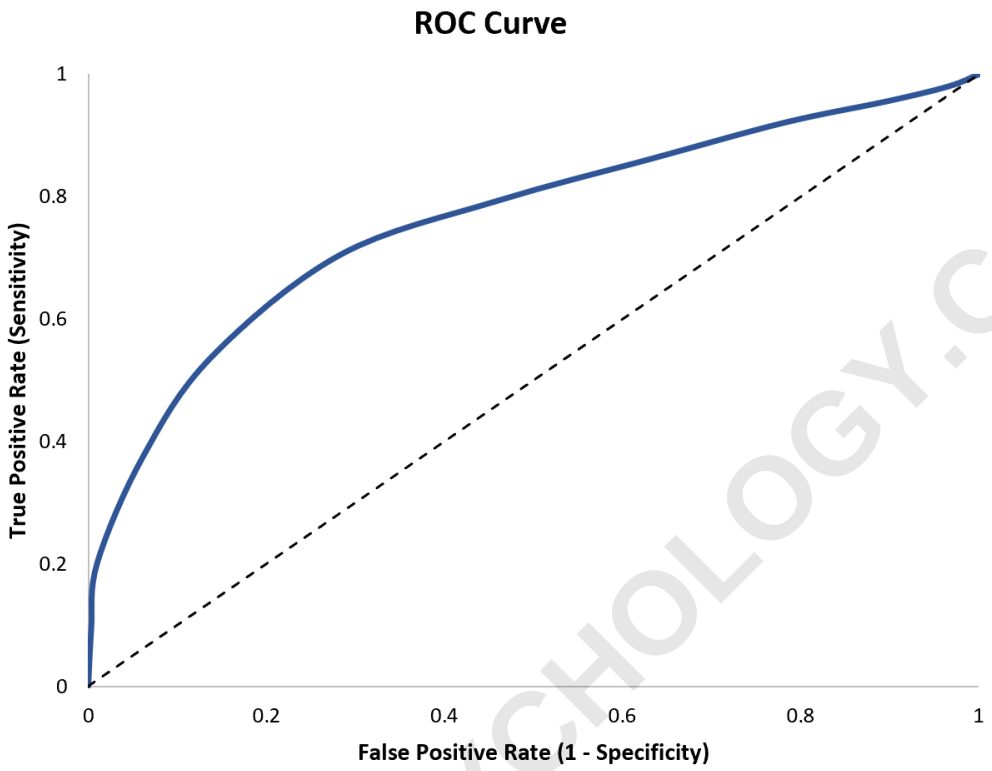

The more that the ROC curve hugs the top left corner of the plot, the better the model does at classifying the data into categories.

To quantify this, we can calculate the AUC (area under the curve) which tells us how much of the plot is located under the curve.

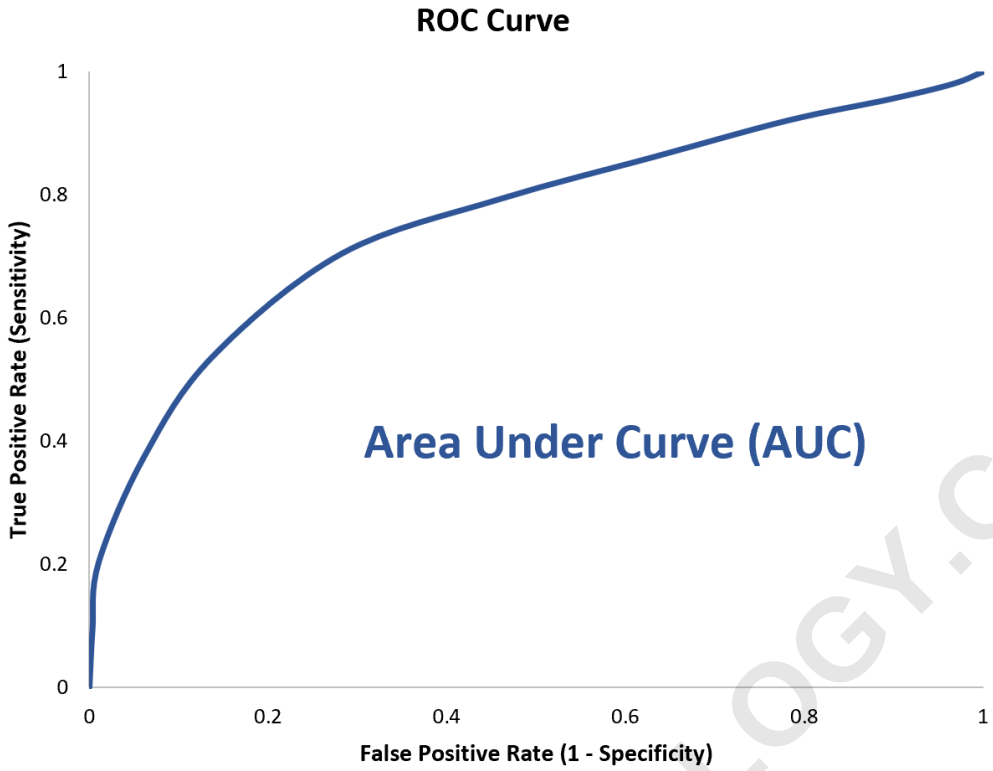

The closer AUC is to 1, the better the model.

When comparing two ROC curves to determine which classification model is best, we often look at which ROC curve "hugs" the top left corner of the plot more and thus has a higher AUC value.

Example: How to Compare Two ROC Curves

Suppose we fit a logistic regression model and a gradient boosted model to a dataset to predict the outcome of some response variable.

Suppose we then create ROC curves to visualize the performance of each model:

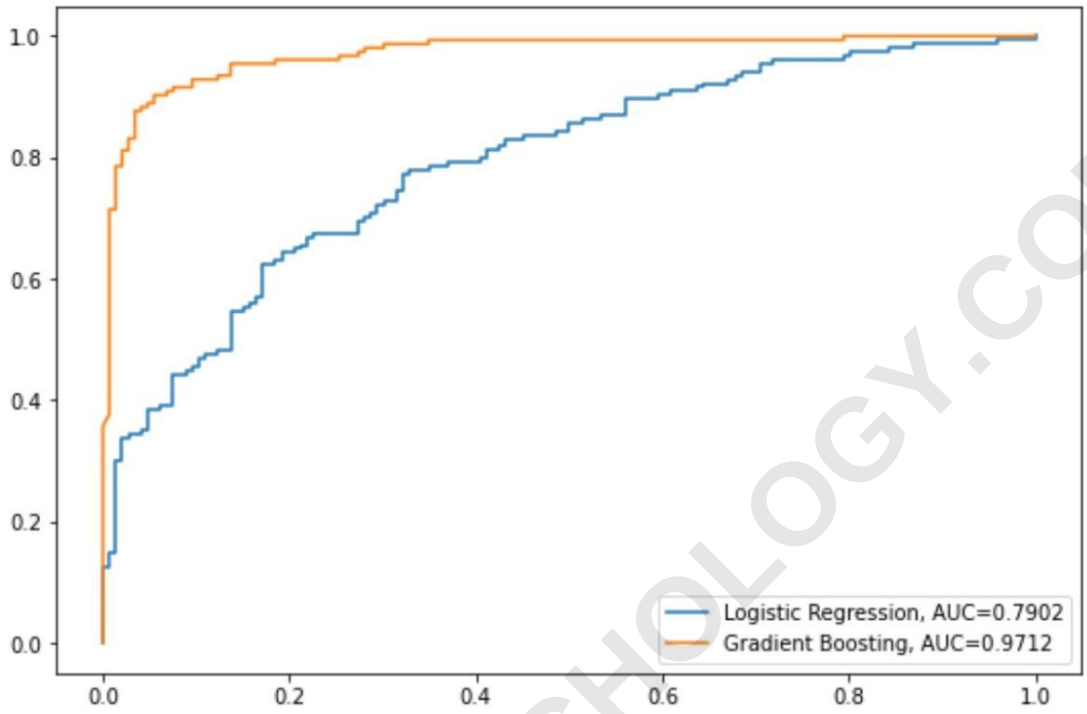

The blue line shows the ROC curve for the logistic regression model and the orange line shows the ROC curve for the gradient boosted model.

From our plot we can see the following AUC values for each model:

AUC of logistic regression model: 0.7902AUC of gradient boosted model: 0.9712

Note: In this example we only compared two ROC

curves, but it's possible to fit several different classification models to a dataset and compare even more ROC curves to determine the best model to use.

The following tutorials provide additional information about classification models and ROC curves: