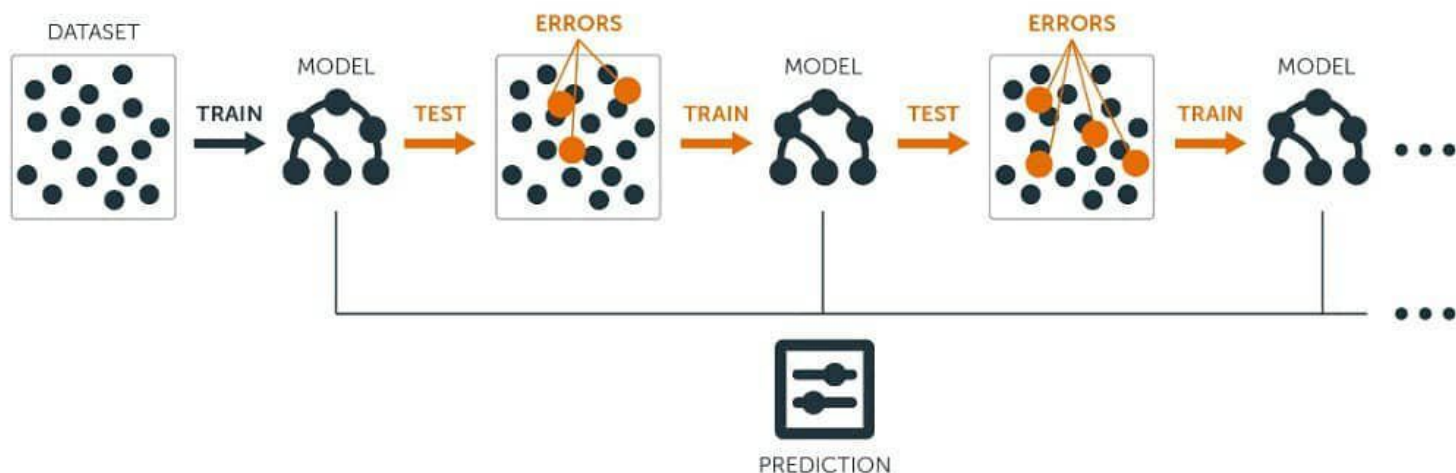


Gradient Boosting

- Boosting is an ensemble technique in which the predictors are not made independently (Bagging), but sequentially.
- The idea of boosting came out of the idea of whether a weak learner can be modified to become better.
- This technique employs the logic in which the subsequent predictors learn from the mistakes of the previous predictors.
- Because new predictors are learning from mistakes committed by previous predictors, it takes less time/iterations to reach close to actual predictions. But we have to choose the stopping criteria carefully or it could lead to overfitting on training data.



Gradient Boosting

- Gradient boosting is a machine learning technique for regression and classification problems, which produces a prediction model in the form of an ensemble of weak prediction models, typically decision trees.
- This method tries to fit the new predictor to the residual errors made by the previous predictor.

Gradient Boosting

Steps of Gradient Boosting.

Fit a simple linear regressor or decision tree on data [call x as input and y as output]

Calculate error residuals. Actual target value, minus predicted target value [$e_1 = y - y_{\text{predicted}1}$]

Fit a new model on error residuals as target variable with same input variables [call it $e_1_{\text{predicted}}$]

Add the predicted residuals to the previous predictions [$y_{\text{predicted}2} = y_{\text{predicted}1} + e_1_{\text{predicted}}$]

Fit another model on residuals that are still left.

i.e. [$e_2 = y - y_{\text{predicted}2}$] and repeat steps 2 to 5 until it starts overfitting or the sum of residuals become constant. Overfitting can be controlled by consistently checking accuracy on validation data.