

# ADAPTIVE BOOSTING

AdaBoost, short for Adaptive Boosting, stands as a cornerstone in the realm of machine learning, offering a potent ensemble learning approach for classification or regression tasks.

## **Weak Learners:**

- Simple decision trees with a single split.
- These weak learners, while individually less capable, are sequentially combined to form a robust classifier.
- Focuses on refining its model iteratively, with each new weak learner addressing the mistakes of its predecessors.

## **Weights and Error:**

- Training by assigning weights to instances in the training set.
- Weights evolve with each iteration, emphasizing the instances that were previously misclassified.
- Allows the algorithm to learn and adapt, gradually improving its performance.

## **Combining Classifiers:**

- Combines the predictions of weak learners judiciously.
- Weighted sum of these learners' predictions, where the weight assigned to each learner reflects its accuracy.
- Results in a powerful, accurate model that outperforms individual weak learners.

## **Boosting:**

- Adheres to the boosting strategy, a technique that accentuates misclassified instances.

- Ensures a targeted and adaptive learning process, enhancing the overall performance of the ensemble.

### **Training Process and Iterative Refinement:**

- Each weak learner refines the model's understanding of the data.
- Adjustments to the weights of misclassified instances
- This helps algorithm toward better generalization and robustness.

### **Predictions and Model Robustness:**

- Aggregating the individual predictions of weak learners.
- Less prone to overfitting.
- Boasts robustness and adaptability, making it a go-to choice for diverse machine learning applications.

### **Limitations:**

- Sensitivity to noisy data and outliers.
- Careful evaluation of data quality is crucial for optimal performance.
- Strength lies in its adaptability, but this very feature can become a limitation if not handled judiciously.

## **Conclusion:**

In conclusion, AdaBoost stands tall as a versatile and powerful algorithm, leveraging the synergy of weak learners to create robust classification models. Its adaptive nature and remarkable performance make it an asset in the machine learning arsenal, providing a blueprint for tackling complex real-world challenges with finesse.

## **Example:**

### **Initialization:**

We start with equal weights for all instances:  $w_1 = w_2 = w_3 = w_4 = 0.25$ .

### **Weak Learner (Decision Stump):**

We use a decision stump as a weak learner.

Let's say the first decision stump (weak learner) focuses on the  $X_1$  feature and splits the data based on  $X_1 > 2.5$ .

The errors are calculated: Instance 1 is misclassified, so its weight increases ( $w_1 = 0.5$ ), while the weights of correctly classified instances decrease ( $w_2 = w_3 = w_4 = 0.167$ ).

## Calculate the Weak Learner's Weight:

The weight of the weak learner ( $\alpha$ ) is calculated based on its error rate. In this case, let's assume the error rate is 0.25.

$$\text{Alpha } (\alpha) = 0.5 * \log\left(\frac{1 - 0.25}{0.25}\right) \approx 0.6496.$$

## Update Instance Weights:

We update the weights of instances based on whether they were correctly or incorrectly classified by the weak learner.

$$w1 = w1 * \exp(-\alpha) \approx 0.5 * \exp(-0.6496) \approx 0.2416$$

$$w2 = w2 * \exp(\alpha) \approx 0.167 * \exp(0.6496) \approx 0.4025$$

$$w3 = w3 * \exp(\alpha) \approx 0.167 * \exp(0.6496) \approx 0.4025$$

$$w4 = w4 * \exp(\alpha) \approx 0.167 * \exp(0.6496) \approx 0.4025$$

## Iterative Process:

Repeat steps 2-4 with a new weak learner. The algorithm focuses on instances that were previously misclassified, adjusting weights and building a strong ensemble.

## Final Prediction:

The final prediction is made by combining the weak learners' predictions, each weighted by its  $\alpha$  value.

This iterative process continues until a predetermined number of weak learners are trained. The final ensemble combines their outputs to make robust predictions.

## Adaptive boosting from Scikit Learn:

### Parameters:

- **base\_estimator (default=DecisionTreeClassifier):**
  - The base estimator, or weak learner, used for the ensemble.
  - It should be a classifier, and its default is a decision tree.
- **n\_estimators:**
  - The maximum number of weak learners (estimators) to train.
  - This controls the number of iterations or rounds in the boosting process.
- **learning\_rate (default=1.0):**
  - The contribution of each weak learner to the final combination.
  - A smaller value requires more weak learners but can improve generalization.
- **algorithm (default='SAMME.R'):**
  - The boosting algorithm to use. 'SAMME.R' is for real boosting, while 'SAMME' is for discrete boosting.
- **random\_state:**
  - Controls the random seed for reproducibility.

```
1 from sklearn.model_selection import train_test_split, cross_val_score
2 from sklearn.datasets import load_iris
3 from sklearn.ensemble import AdaBoostClassifier
4 from sklearn.tree import DecisionTreeClassifier
5 import warnings
6
7
8 warnings.filterwarnings("ignore", category=FutureWarning)
9
10 # Load the Iris dataset
11 iris = load_iris()
12 X, y = iris.data, iris.target
13
14 # Split the data into training and testing sets
15 X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)
16
17 # Create a base estimator (Decision Tree Classifier)
18 base_estimator = DecisionTreeClassifier(max_depth=10) # Use a more complex base estimator
19
20 # Create AdaBoostClassifier
21 adaboost_clf = AdaBoostClassifier(base_estimator=base_estimator, n_estimators=50, learning_rate=1.0, random_state=42)
22
23 # Cross-validation
24 cv_scores = cross_val_score(adaboost_clf, X, y, cv=5)
25 print(f"Cross-Validation Scores: {cv_scores}")
26 print(f"Mean Accuracy: {cv_scores.mean()}")
27
28 # # Train the AdaBoostClassifier on the full training set
29 # adaboost_clf.fit(X_train, y_train)
30
31 # # Evaluate on the test set
32 # test_accuracy = adaboost_clf.score(X_test, y_test)
33 # print(f"Test Accuracy: {test_accuracy}")
34
```

```
📄 Cross-Validation Scores: [0.96666667 0.96666667 0.9        0.93333333 1.        ]
Mean Accuracy: 0.9533333333333334
```