



[Journey to Success] 10,000+ Professionals followed this Roadmap to become a Top Data Scientist 🏆

Download Roadmap

[Home](#)

👤 [Anshul Saini](#) — Published On September 20, 2021 and Last Modified On October 14th, 2021

[Algorithm](#) [Beginner](#) [Guide](#) [Machine Learning](#) [Maths](#)

Find High-Quali

This article was published as a part of the [Data Science Blogathon](#)



Introduction

In this article, I am going to discuss the math intuition behind the Gradient boosting algorithm. It is more popularly known as Gradient boosting Machine or GBM. It is a boosting method and I have talked more about boosting in this [article](#).

Gradient boosting is a method standing out for its prediction speed and accuracy, particularly with large and complex datasets. From Kaggle competitions to machine learning solutions for business, this algorithm has produced the best results. We already know that errors play a major role in any machine learning algorithm. There are mainly two types of error, bias error and variance error. Gradient boost algorithm *helps us minimize bias error* of the model

Before getting into the details of this algorithm we must have some knowledge about AdaBoost Algorithm which is again a boosting method. This algorithm starts by building a decision stump and then assigning equal weights to all the data points. Then it increases the weights for all the points which are misclassified and lowers the weight for those that are easy to classify or are correctly classified. A new decision stump is made for these weighted data points. The idea behind this is to improve the

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

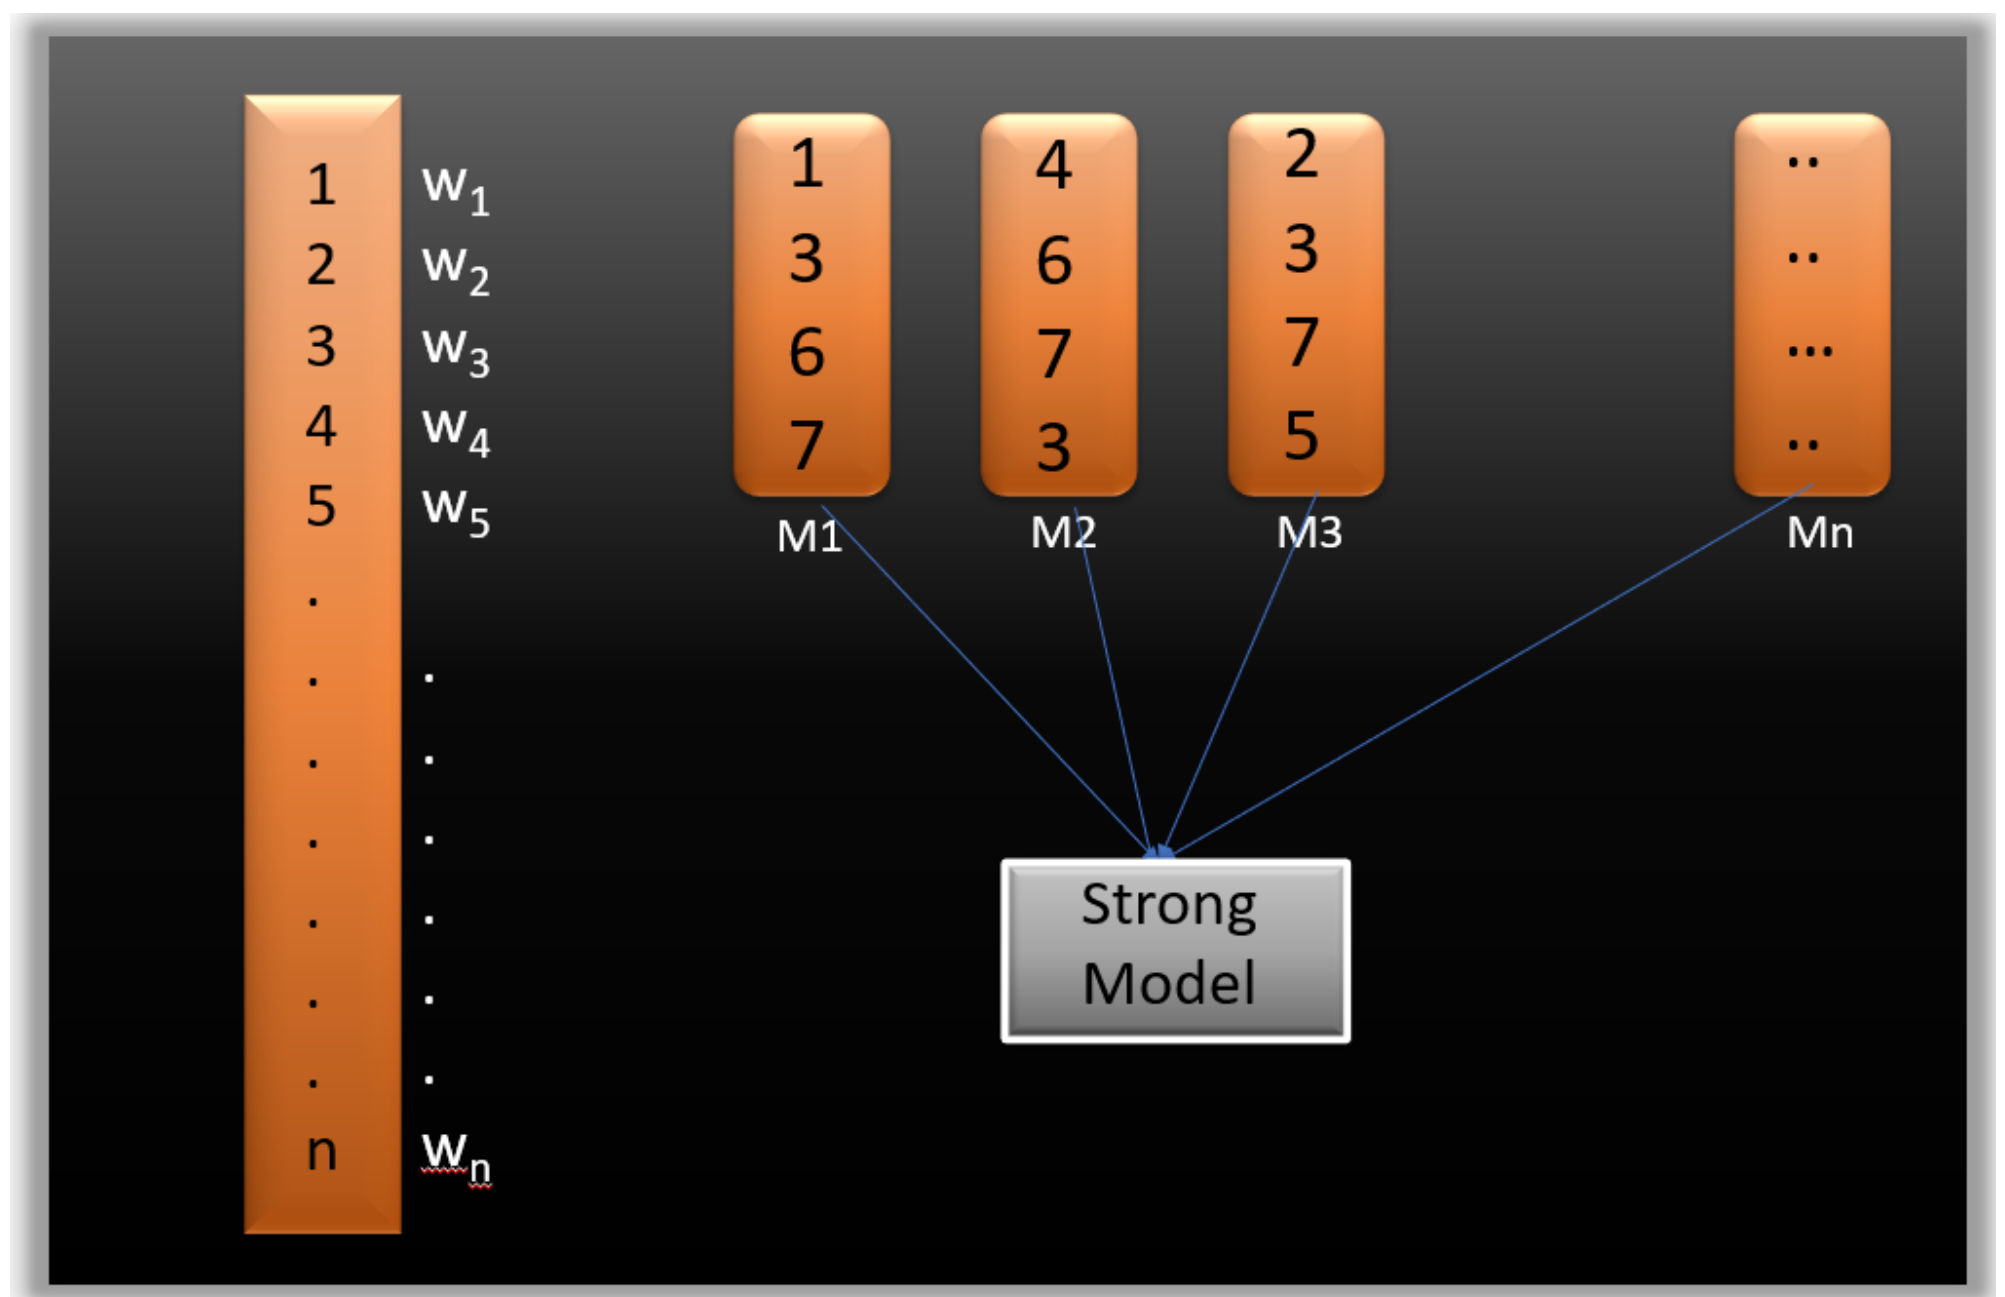
whereas in AdaBoost we can change the base estimator according to our needs.

Table of Contents

1. What is Boosting technique?
2. Gradient Boosting Algorithm
3. Gradient Boosting Regressor
4. Example of gradient boosting
5. Gradient Boosting Classifier
6. Implementation using Scikit-learn
7. Parameter Tuning in Gradient Boosting (GBM) in Python
8. End Notes

What is boosting?

While studying machine learning you must have come across this term called Boosting. It is the most misinterpreted term in the field of Data Science. The principle behind boosting algorithms is first we built a model on the training dataset, then a second model is built to rectify the errors present in the first model. Let me try to explain to you what exactly does this means and how does this works.



We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you agree to our [Privacy Policy](#) and [Terms of Use](#).

Gradient Boosting Algorithm: A Complete Guide for Beginners

Now what we do is randomly select observations from the training dataset and feed them to model 1 (M1), we also assume that initially, all the observations have an equal weight that means an equal probability of getting selected.

Remember in ensembling techniques the weak learners combine to make a strong model so here M1, M2, M3....Mn all are weak learners.

Since M1 is a weak learner, it will surely misclassify some of the observations. Now before feeding the observations to M2 what we do is update the weights of the observations which are wrongly classified. You can think of it as a bag that initially contains 10 different color balls but after some time some kid takes out his favorite color ball and put 4 red color balls instead inside the bag. Now off-course the probability of selecting a red ball is higher. This same phenomenon happens in Boosting techniques, when an observation is wrongly classified, its weight get's updated and for those which are correctly classified, their weights get decreased. The probability of selecting a wrongly classified observation gets increased hence in the next model only those observations get selected which were misclassified in model 1.

Similarly, it happens with M2, the wrongly classified weights are again updated and then fed to M3. This procedure is continued until and unless the errors are minimized, and the dataset is predicted correctly. Now when the new datapoint comes in (Test data) it passes through all the models (weak learners) and the class which gets the highest vote is the output for our test data.

What is a Gradient boosting Algorithm?

The main idea behind this algorithm is to build models sequentially and these subsequent models try to reduce the errors of the previous model. But how do we do that? How do we reduce the error? This is done by building a new model on the errors or residuals of the previous model.

When the target column is continuous, we use **Gradient Boosting Regressor** whereas when it is a classification problem, we use **Gradient Boosting Classifier**. The only difference between the two is the "*Loss function*". The objective here is to minimize this loss function by adding weak learners using gradient descent. Since it is based on loss function hence for regression problems, we'll have different loss functions like Mean squared error (**MSE**) and for classification, we will have different for e.g **log-likelihood**.

Understand Gradient Boosting Algorithm with example

Let's understand the intuition behind Gradient boosting with the help of an example. Here our target column is continuous hence we will use Gradient Boosting Regressor.

Following is a sample from a random dataset where we have to predict the car price based on various features. The target column is price and other features are independent features.

Row No.	Cylinder Number	Car Height	Engine Location	Price
1	Four	48.8	Front	12000
2	Six	48.8	Back	16500
3	Five	52.4	Back	15500

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

Step -1 The first step in gradient boosting is to build a base model to predict the observations in the training dataset. For simplicity we take an average of the target column and assume that to be the predicted value as shown below:

Row No.	Cylinder Number	Car Height	Engine Location	Price	Prediction 1
1	Four	48.8	Front	12000	14500
2	Six	48.8	Back	16500	14500
3	Five	52.4	Back	15500	14500
4	Four	54.3	Front	14000	14500

Image Source: Author

Why did I say we take the average of the target column? Well, there is math involved behind this. Mathematically the first step can be written as:

$$F_0(x) = \arg \min_{\gamma} \sum_{i=1}^n L(y_i, \gamma).$$

Looking at this may give you a headache, but don't worry we will try to understand what is written here.

Here L is our loss function

Gamma is our predicted value

argmin means we have to find a predicted value/gamma for which the loss function is minimum.

Since the target column is continuous our loss function will be:

$$L = \frac{1}{n} \sum_{i=0}^n (y_i - \gamma_i)^2$$

Here y_i is the observed value

And gamma is the predicted value

Now we need to find a minimum value of gamma such that this loss function is minimum. We all have studied how to find minima and maxima in our 12th grade. Did we use to differentiate this loss function and then put it equal to 0 right? Yes, we will do the same here.

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

$$\frac{dL}{d\gamma} = \frac{1}{2} \left(\sum_{i=0} (y_i - \gamma_i) \right) = - \sum_{i=0} (y_i - \gamma_i)$$

Let's see how to do this with the help of our example. Remember that y_i is our observed value and γ_i is our predicted value, by plugging the values in the above formula we get:

$$L = \frac{1}{2}(12000 - \gamma)^2 + \frac{1}{2}(16500 - \gamma)^2 + \frac{1}{2}(15500 - \gamma)^2 + \frac{1}{2}(14000 - \gamma)^2$$

$$\frac{dL}{d\gamma} = \frac{2}{2}(12000 - \gamma)(-1) + \frac{2}{2}(16500 - \gamma)(-1) + \frac{2}{2}(15500 - \gamma)(-1) + \frac{2}{2}(14000 - \gamma)(-1)$$

Now $\frac{dL}{d\gamma} = 0$ and taking $(-)$ common

$$\Rightarrow -[12000 - \gamma + 16500 - \gamma + 15500 - \gamma + 14000 - \gamma] = 0$$

$$\Rightarrow [58000 - 4\gamma] = 0$$

$$\Rightarrow 58000 = 4\gamma$$

$$\Rightarrow \gamma = \frac{58000}{4} = 14500$$

We end up over an average of the observed car price and this is why I asked you to take the average of the target column and assume it to be your first prediction.

Hence for $\gamma=14500$, the loss function will be minimum so this value will become our prediction for the *base model*.

Step-2 The next step is to calculate the pseudo residuals which are (observed value - predicted value)

Row No.	Cylinder Number	Car Height	Engine Location	Price	Prediction 1	Residual 1
1	Four	48.8	Front	12000	14500	-2500
2	Six	48.8	Back	16500	14500	2000
3	Five	52.4	Back	15500	14500	1000
4	Four	54.3	Front	14000	14500	-500

Image Source: Author

Again the question comes why only observed - predicted? Everything is mathematically proved, let's from where did this formula come from. This step can be written as:

$$r_{im} = - \left[\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)} \right]_{F(x)=F_{m-1}(x)} \quad \text{for } i = 1, \dots, n.$$

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#).

Gradient Boosting Algorithm: A Complete Guide for Beginners

$$\frac{dL}{d\gamma} = -(y_i - \gamma_i) = -(Observed - Predicted)$$

If you see the formula of residuals above, we see that the derivative of the loss function is multiplied by a negative sign, so now we get:

$$(Observed - Predicted)$$

The predicted value here is the prediction made by the previous model. In our example the prediction made by the previous model (initial base model prediction) is 14500, to calculate the residuals our formula becomes:

$$(Observed - 14500)$$

In the next step, we will build a model on these pseudo residuals and make predictions. Why do we do this? Because we want to minimize these residuals and minimizing the residuals will eventually improve our model accuracy and prediction power. So, using the Residual as target and the original feature Cylinder number, cylinder height, and Engine location we will generate new predictions. Note that the predictions, in this case, will be the error values, not the predicted car price values since our target column is an error now.

Let's say $h_m(x)$ is our DT made on these residuals.

Step- 4 In this step we find the output values for each leaf of our decision tree. That means there might be a case where 1 leaf gets more than 1 residual, hence we need to find the final output of all the leaves. TO find the output we can simply take the average of all the numbers in a leaf, doesn't matter if there is only 1 number or more than 1.

Let's see why do we take the average of all the numbers. Mathematically this step can be represented as:

$$\gamma_m = \arg \min_{\gamma} \sum_{i=1}^n L(y_i, F_{m-1}(x_i) + \gamma h_m(x_i)).$$

Here $h_m(x_i)$ is the DT made on residuals and m is the number of DT. When $m=1$ we are talking about the 1st DT and when it is "M" we are talking about the last DT.

The output value for the leaf is the value of gamma that minimizes the Loss function. The left-hand side "Gamma" is the output value of a particular leaf. On the right-hand side $[F_{m-1}(x_i) + \gamma h_m(x_i)]$ is similar to step 1 but here the difference is that we are taking previous predictions whereas earlier there was no previous prediction.

Let's understand this even better with the help of an example. Suppose this is our regressor tree:

Gradient Boosting Algorithm: A Complete Guide for Beginners

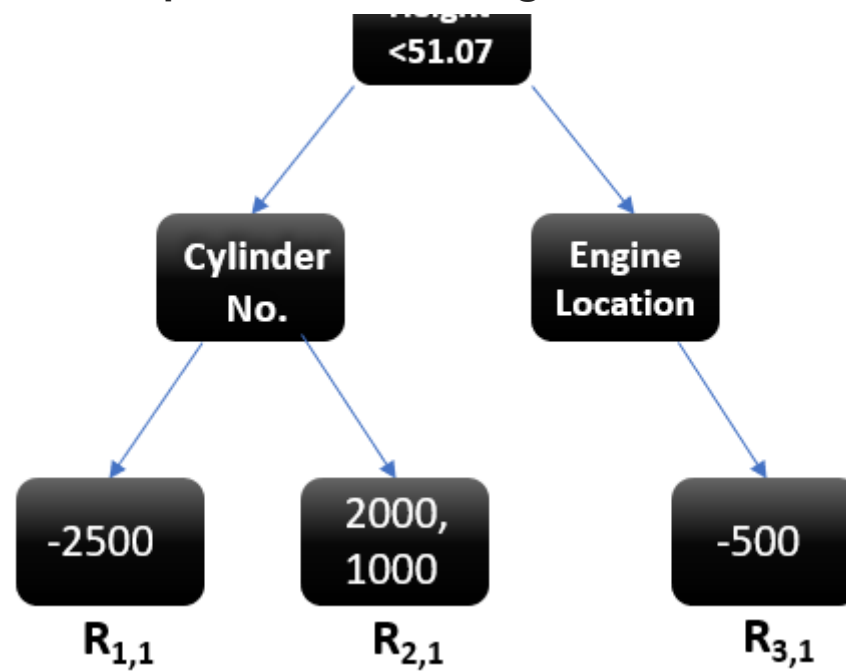


Image Source

We see 1st residual goes in $R_{1,1}$, 2nd and 3rd residuals go in $R_{2,1}$ and 4th residual goes in $R_{3,1}$.

Let's calculate the output for the first leaf that is $R_{1,1}$

$$\gamma_{1,1} = \operatorname{argmin} \frac{1}{2} (12000 - (14500 + \gamma))^2$$

$$\gamma_{1,1} = \operatorname{argmin} \frac{1}{2} (-2500 - \gamma)^2$$

Now we need to find the value for gamma for which this function is minimum. So we find the derivative of this equation w.r.t gamma and put it equal to 0.

$$\frac{d}{d\gamma} \frac{1}{2} (-2500 - \gamma)^2 = 0$$

$$-2500 - \gamma = 0$$

$$\gamma = -2500$$

Hence the leaf $R_{1,1}$ has an output value of -2500. Now let's solve for the $R_{2,1}$

$$\gamma_{2,1} = \operatorname{argmin} \left[\frac{1}{2} (16500 - (14500 + \gamma))^2 + \frac{1}{2} (15500 - (14500 + \gamma))^2 \right]$$

$$\gamma_{2,1} = \operatorname{argmin} \left[\frac{1}{2} (2000 - \gamma)^2 + \frac{1}{2} (1000 - \gamma)^2 \right]$$

Let's take the derivative to get the minimum value of gamma for which this function is minimum:

Gradient Boosting Algorithm: A Complete Guide for Beginners

$$\frac{d}{d\gamma} \left[\frac{1}{2}(2000 - \gamma)^2 + \frac{1}{2}(1000 - \gamma)^2 \right] = 0$$

$$2000 - \gamma + 1000 - \gamma = 0$$

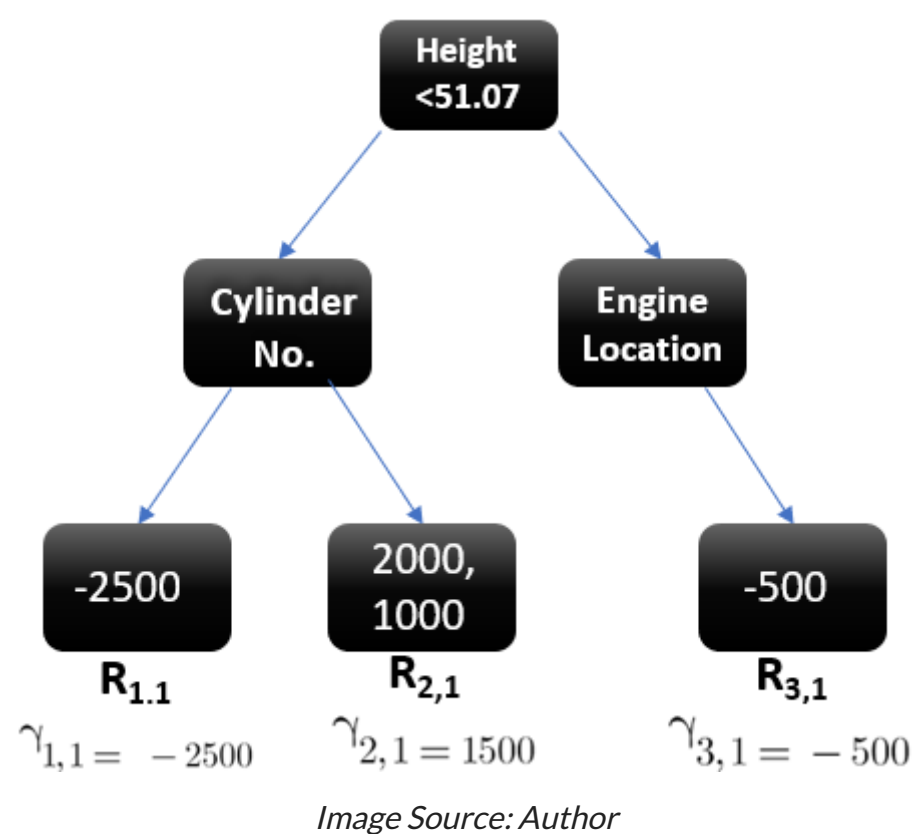
$$3000 - 2\gamma = 0$$

$$\frac{3000}{2} = \gamma$$

$$\gamma = 1500$$

We end up with the **average** of the residuals in the leaf $R_{2,1}$. Hence if we get any leaf with more than 1 residual, we can simply find the average of that leaf and that will be our final output.

Now after calculating the output of all the leaves, we get:



Step-5 This is finally the last step where we have to update the predictions of the previous model. It can be updated as:

Update the model:

$$F_m(x) = F_{m-1}(x) + \nu_m h_m(x)$$

where m is the number of decision trees made.

Since we have just started building our model so our $m=1$. Now to make a new DT our new predictions will be:



Image Source: Author

Here $F_{m-1}(x)$ is the prediction of the base model (previous prediction) since $F_0(x) = F$ is our base model hence the previous

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

improves accuracy in the long run. Let's take $nu=0.1$ in this example.

$H_m(x)$ is the recent DT made on the residuals.

Let's calculate the new prediction now:

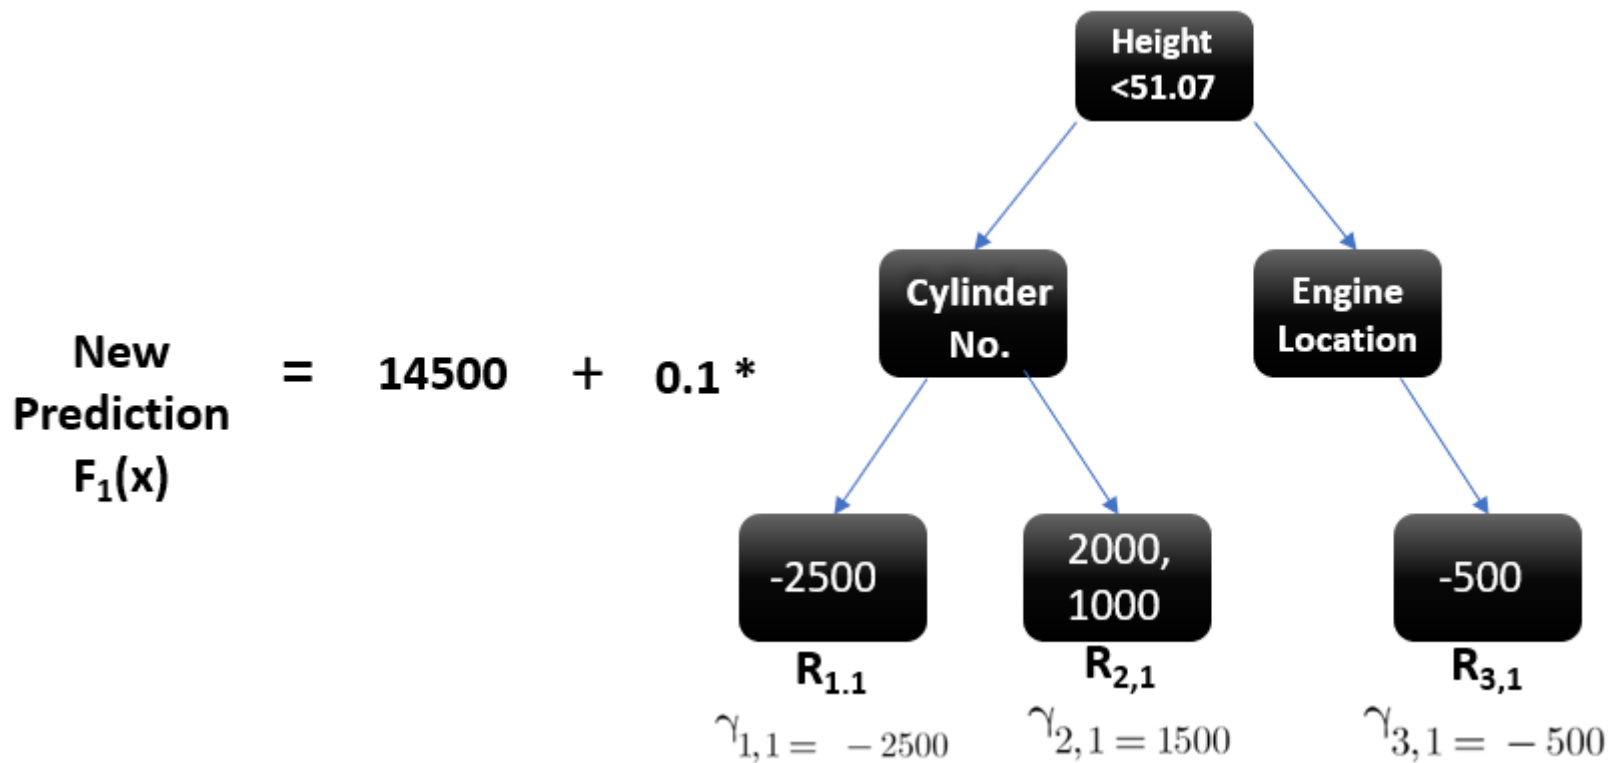


Image Source: Author

Suppose we want to find a prediction of our first data point which has a car height of 48.8. This data point will go through this decision tree and the output it gets will be multiplied with the learning rate and then added to the previous prediction.

Now let's say $m=2$ which means we have built 2 decision trees and now we want to have new predictions.

This time we will add the previous prediction that is $F_1(x)$ to the new DT made on residuals. We will iterate through these steps again and again till the **loss is negligible**.

I am taking a hypothetical example here just to make you understand how this predicts for a new dataset:

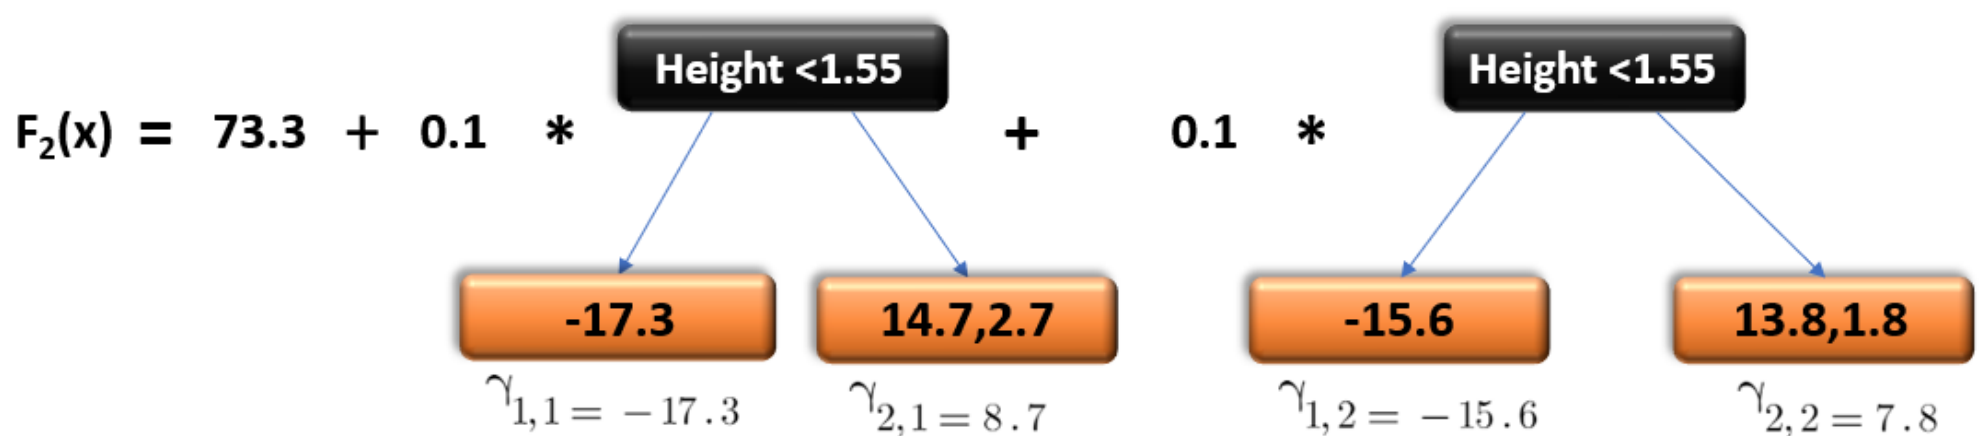


Image Source: Author

If a new data point says height = 1.40 comes, it'll go through all the trees and then will give the prediction. Here we have only 2 trees hence the datapoint will go through these 2 trees and the final output will be $F_2(x)$.

Gradient Boosting Algorithm: A Complete Guide for Beginners

A gradient boosting classifier is used when the target column is binary. All the steps explained in the Gradient boosting regressor are used here, the only difference is we change the loss function. Earlier we used Mean squared error when the target column was continuous but this time, we will use log-likelihood as our loss function.

Let's see how this loss function works, to read more about log-likelihood I recommend you to go through this [article](#) where I have given each detail you need to understand this.

The loss function for the classification problem is given below:

$$L = - \sum_{i=1}^n y_i \log(p) + (1 - p) \log(1 - p)$$

Our first step in the gradient boosting algorithm was to initialize the model with some constant value, there we used the average of the target column but here we'll use log(odds) to get that constant value. The question comes why log(odds)?

When we differentiate this loss function, we will get a function of log(odds) and then we need to find a value of log(odds) for which the loss function is minimum.

Confused right? Okay let's see how it works:

Let's first transform this loss function so that it is a function of log(odds), I'll tell you later why we did this transformation.

$$\begin{aligned} L &= - \left[\sum_{i=1}^n y_i \log(p) + (1 - y_i) \log(1 - p) \right] \\ &= -y * \log(p) - (1 - y) * \log(1 - p) \\ &= -y * \log(p) - \log(1 - p) + y * \log(1 - p) \\ &= y * \left[\log(p) - \log(1 - p) \right] - \log(1 - p) \\ &= -y * \left[\frac{\log(p)}{\log(1 - p)} \right] - \log(1 - p) \\ &= -y * \log\left(\frac{p}{1 - p}\right) - \log(1 - p) \end{aligned}$$

Gradient Boosting Algorithm: A Complete Guide for Beginners

$$\begin{aligned}
 \text{Now, } \log(1-p) &= \log\left(1 - \frac{e^{\log(\text{odds})}}{1 + e^{\log(\text{odds})}}\right) \\
 &= \log\left(\frac{1 + e^{\log(\text{odds})}}{1 + e^{\log(\text{odds})}} - \frac{e^{\log(\text{odds})}}{1 + e^{\log(\text{odds})}}\right) = \log\left(\frac{1}{1 + e^{\log(\text{odds})}}\right) \\
 &= \log(1) - \log\left(1 + e^{\log(\text{odds})}\right) \\
 &= 0 - \log\left(1 + e^{\log(\text{odds})}\right)
 \end{aligned}$$

$$\text{Hence, } -y * \log(\text{odds}) - \left(-\log\left(1 + e^{\log(\text{odds})}\right)\right)$$

$$L = -y * \log(\text{odds}) + \log\left(1 + e^{\log(\text{odds})}\right)$$

Now this is our loss function, and we need to minimize it, for this, we take the derivative of this w.r.t to $\log(\text{odds})$ and then put it equal to 0,

$$\frac{dL}{d[\log(\text{odds})]} = -y + \frac{e^{\log(\text{odds})}}{1 + e^{\log(\text{odds})}}$$

We know $\frac{e^{\log(\text{odds})}}{1 + e^{\log(\text{odds})}} = p$, hence we can substitute p

$$\frac{dL}{d[\log(\text{odds})]} = -y + p$$

Here y are the observed values

You must be wondering that why did we transform the loss function into the function of $\log(\text{odds})$. Actually, sometimes it is easy to use the function of $\log(\text{odds})$, and sometimes it's easy to use the function of predicted probability " p ".

It is not compulsory to transform the loss function, we did this just to have easy calculations.

Hence the minimum value of this loss function will be our first prediction (base model prediction)

Now in the Gradient boosting regressor our next step was to calculate the pseudo residuals where we multiplied the derivative

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you agree to our [Privacy Policy](#) and [Terms of Use](#).

Gradient Boosting Algorithm: A Complete Guide for Beginners

$$\frac{dL}{d[\log(\text{odds})]} = -y + p$$

$$\frac{dL}{d[\log(\text{odds})]} = -(-y + p) = (y - p) = (\text{observed} - \text{predicted})$$

After finding the residuals we can build a decision tree with all independent variables and target variables as "Residuals".

Now when we have our first decision tree, we find the final output of the leaves because there might be a case where a leaf gets more than 1 residuals, so we need to calculate the final output value. The math behind this step is out of the scope of this article so I will mention the direct formula to calculate the output of a leaf:

$$\gamma = \frac{\sum_{i=1}^n \text{Residual}_i}{\sum_{i=1}^n [\text{Previous probability}_i \times (1 - \text{Previous probability}_i)]}$$

Finally, we are ready to get new predictions by adding our base model with the new tree we made on residuals.

There are a few variations of gradient boosting and a couple of them are momentarily clarified in the coming article.

Implementation Using scikit-learn

For implementation on a dataset, we will be using the Income Evaluation dataset, which has information about an individual's personal life and an output of 50K or <=50. The dataset can be found here (<https://www.kaggle.com/lodetomasi1995/income-classification>)

The task here is to classify the income of an individual, when given the required inputs about his personal life.

First, let's import all required libraries.

Gradient Boosting Algorithm: A Complete Guide for Beginners

```

from sklearn.ensemble import GradientBoostingClassifier

import numpy as np

import pandas as pd

from sklearn.preprocessing import StandardScaler

from sklearn.model_selection import train_test_split

from sklearn.metrics import accuracy_score, confusion_matrix

from sklearn import preprocessing

import warnings

warnings.filterwarnings("ignore")

```

Now let's read the dataset and look at the columns to understand the information better.

```

df = pd.read_csv('income_evaluation.csv')
df.head()

```

	age	workclass	fnlwgt	education	education-num	marital-status	occupation	relationship	race	sex	capital-gain	capital-loss	hours-per-week	native-country	income
0	39	State-gov	77516	Bachelors	13	Never-married	Adm-clerical	Not-in-family	White	Male	2174	0	40	United-States	<=50K
1	50	Self-emp-not-inc	83311	Bachelors	13	Married-civ-spouse	Exec-managerial	Husband	White	Male	0	0	13	United-States	<=50K
2	38	Private	215646	HS-grad	9	Divorced	Handlers-cleaners	Not-in-family	White	Male	0	0	40	United-States	<=50K
3	53	Private	234721	11th	7	Married-civ-spouse	Handlers-cleaners	Husband	Black	Male	0	0	40	United-States	<=50K
4	28	Private	338409	Bachelors	13	Married-civ-spouse	Prof-specialty	Wife	Black	Female	0	0	40	Cuba	<=50K

I have already done the data preprocessing part and you can look whole code [here](#). Here my main aim is to tell you how to implement this on python. Now for training and testing our model, the data has to be divided into train and test data.

We will also scale the data to lie between 0 and 1.

```
# Split dataset into test and train data
```

```
X_train, X_test, y_train, y_test = train_test_split(df.drop('income', axis=1), df['income'],
test_size=0.2)
```

Now let's go ahead with defining the Gradient Boosting Classifier along with its hyperparameters. Next, we will fit this model on the training data.

```
# Define Gradient Boosting Classifier with hyperparameters
```

```
gbc=GradientBoostingClassifier(n_estimators=500, learning_rate=0.05, random_state=100, max_features=5 )
```

```
# Fit train data to GBC
```

```
gbc.fit(X_train, y_train)
```

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

Below, you can see the confusion matrix of the model, which gives a report of the number of classifications and misclassifications.

```
# Confusion matrix will give number of correct and incorrect classifications
```

```
print(confusion_matrix(y_test, gbc.predict(X_test)))
          [[4107  801]
           [ 533 4195]]
```

The number of misclassifications by the Gradient Boosting Classifier are 1334, compared to 8302 correct classifications. The model has performed decently.

Let's check the accuracy

```
# Accuracy of model
```

```
print("GBC accuracy is %2.2f" % accuracy_score(
    y_test, gbc.predict(X_test)))
```

```
GBC accuracy is 0.86
```

Let's check the classification report also:

```
from sklearn.metrics import classification_report
```

```
pred=gbc.predict(X_test)
```

```
print(classification_report(y_test, pred))
```

	precision	recall	f1-score	support
0	0.89	0.84	0.86	4908
1	0.84	0.89	0.86	4728
accuracy			0.86	9636
macro avg	0.86	0.86	0.86	9636
weighted avg	0.86	0.86	0.86	9636

The accuracy is 86%, which is pretty good but this can be improved by tuning the hyperparameters or processing the data to remove outliers.

This however gives us the basic idea behind gradient boosting and its underlying working principles.

Parameter Tuning in Gradient Boosting (GBM) in Python

Tuning n_estimators and Learning rate

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

rate as low values always work better, given that we train on sufficient number of trees. A high number of trees can be computationally expensive that's why I have taken few number of trees here.

```
from sklearn.model_selection import GridSearchCV

grid = {

    'learning_rate':[0.01,0.05,0.1],

    'n_estimators':np.arange(100,500,100),

}

gb = GradientBoostingClassifier()

gb_cv = GridSearchCV(gb, grid, cv = 4)

gb_cv.fit(X_train,y_train)

print("Best Parameters:",gb_cv.best_params_)

print("Train Score:",gb_cv.best_score_)

print("Test Score:",gb_cv.score(X_test,y_test))

Best Parameters: {'learning_rate': 0.1, 'n_estimators': 400}
Train Score: 0.8973320876154883
Test Score: 0.8974465434917999
```

We see the accuracy increased from 86 to 89 after tuning n_estimators and learning rate. Also the “true positive” and the “true negative” rate improved.

	precision	recall	f1-score	support
0	0.91	0.88	0.89	4908
1	0.88	0.91	0.89	4728
accuracy			0.89	9636
macro avg	0.89	0.89	0.89	9636
weighted avg	0.89	0.89	0.89	9636

We can also tune max_depth parameter which you must have heard in decision trees and random forests.

Gradient Boosting Algorithm: A Complete Guide for Beginners

```
gb = GradientBoostingClassifier(learning_rate=0.1,n_estimators=400)

gb_cv = GridSearchCV(gb, grid, cv = 4)

gb_cv.fit(X_train,y_train)

print("Best Parameters:",gb_cv.best_params_)

print("Train Score:",gb_cv.best_score_)

print("Test Score:",gb_cv.score(X_test,y_test))
```

```
Best Parameters: {'max_depth': 7}
Train Score: 0.9138861970645884
Test Score: 0.9124117891241179
```

The accuracy has increased even more when we tuned the parameter “max_depth”.

End Notes

I hope you got an understanding of how the Gradient Boosting algorithm works under the hood. I have tried to show you the math behind this is the easiest way possible.

In the next article, I will explain Xtreme Gradient Boosting (XGB), which is again a new technique to combine various models and to improve our accuracy score. It is just an extension of the gradient boost algorithm.

Let me know if you have any queries in the comments below.

About the Author

I am an undergraduate student currently in my last year majoring in Statistics (Bachelors of Statistics) and have a strong interest in the field of data science, machine learning, and artificial intelligence. I enjoy diving into data to discover trends and other valuable insights about the data. I am constantly learning and motivated to try new things.

I am open to collaboration and work.

For any doubt and queries, feel free to contact me on [Email](#)

Connect with me on [LinkedIn](#) and [Twitter](#)

The media shown in this article are not owned by Analytics Vidhya and are used at the Author’s discretion.

[blogathon](#) [Gradient Boosting](#) [machine learning](#) [Supervised Learning](#)

About the Author



[Anshul Saini](#)


I am an undergraduate student currently in my last year majoring in Statistics (Bachelors of Statistics) and have a strong interest in the field of data science, machine learning, and artificial intelligence. I enjoy diving into data to discover trends and other valuable insights about the data. I am constantly learning and motivated to try

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you

agree to our [Privacy Policy](#) and [Terms of Use](#). Accept

Gradient Boosting Algorithm: A Complete Guide for Beginners

Our Top Authors



Download

Analytics Vidhya App for the Latest blog/Article



Previous Post

[How to Develop a Virtual Keyboard Using OpenCV](#)

Next Post

[Beginner's Guide to Recursion in Python](#)

One thought on "Gradient Boosting Algorithm: A Complete Guide for Beginners"



YesNoSpin says:
January 13, 2023 at 2:55 pm

This is a great guide for beginners!

[Reply](#)

Leave a Reply

Your email address will not be published. Required fields are marked *

Notify me of follow-up comments by email.


Notify me of new posts by email.

Submit

We use cookies on Analytics Vidhya websites to deliver our services, analyze web traffic, and improve your experience on the site. By using Analytics Vidhya, you agree to our [Privacy Policy](#) and [Terms of Use](#). [Accept](#)

Gradient Boosting Algorithm: A Complete Guide for Beginners

Top Resources




[How to Read and Write With CSV Files in Python...](#)

[Harika Bonthu](#) - AUG 21, 2021



[Understand Random Forest Algorithms With Examples \(Updated 2023\)](#)

[Sruthi E R](#) - JUN 17, 2021



[Understanding & Interpreting Confusion Matrices for Machine Learning \(Updated 2023\)](#)

[Aniruddha Bhandari](#) - APR 17, 2020



[Feature Selection Techniques in Machine Learning \(Updated 2023\)](#)

[Aman Gupta](#) - OCT 10, 2020



Analytics Vidhya

[About Us](#)

[Our Team](#)

[Careers](#)

[Contact us](#)

Companies

[Post Jobs](#)

[Trainings](#)

[Hiring Hackathons](#)

[Advertising](#)

Data Scientists

[Blog](#)

[Hackathon](#)

[Discussions](#)

[Apply Jobs](#)

[Visit us](#)

