

What is the calculation for Mean Absolute Error (MAE) in R?

Authored by
stats writer

April 28, 2024

RECOMMENDED CITATION

stats writer (2024). *What is the calculation for Mean Absolute Error (MAE) in R?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=140540>

The Mean Absolute Error (MAE) is a commonly used metric for evaluating the accuracy of a regression model in R. It measures the average absolute difference between the actual and predicted values of a variable. The calculation for MAE involves taking the average of the absolute differences between the predicted and actual values. This metric is useful in assessing the overall performance of a model, as it gives equal weight to all errors, regardless of their direction. In R, the MAE is calculated by taking the sum of the absolute differences and dividing it by the total number of observations. A lower MAE value indicates a better fit between the model and the data.

Calculate Mean Absolute Error in R

In statistics, the mean absolute error (MAE) is a way to measure the accuracy of a given model. It is calculated as:

$$\text{MAE} = (1/n) * \sum |y_i - x_i|$$

where:

Σ : A Greek symbol that means "sum"
 y_i : The observed value for the i th observation
 x_i : The predicted value for the i th observation
 n : The total number of observations

We can calculate the mean absolute error in R by using the `mae(actual, predicted)` function from the Metrics package.

This tutorial provides two examples of how to use this function in practice.

Example 1: Calculate Mean Absolute Error Between Two Vectors

The following code shows how to calculate the mean absolute error between a vector of observed values and a vector of predicted values:

```
library(Metrics)

#define observed and predicted values
observed <- c(12, 13, 14, 15, 15, 22, 27, 29, 29, 30, 32)
predicted <- c(11, 13, 14, 14, 16, 19, 24, 30, 32, 36, 30)

#calculate mean absolute error between vectors
mae(observed, predicted)

1.909091
```

The mean absolute error (MAE) turns out to be 1.909.

This tells us that the average absolute difference between the observed values and the predicted values is 1.909.

Example 2: Calculate Mean Absolute Error for a Regression Model

The following code shows how to fit a in R and then calculate the mean absolute error between the

predictions made by the model and the actual observed response values:

```
library(Metrics)
```

```
#create data
```

```
df <- data.frame(x1=c(1, 3, 3, 4, 4, 6, 6, 8, 9, 3),  
x2=c(7, 7, 4, 10, 13, 12, 17, 19, 20, 34),  
y=c(17, 18, 19, 20, 24, 28, 25, 29, 30, 32))
```

```
#view first six rows of data
```

```
head(df)
```

```
x1 x2 y  
1 1 7 17  
2 3 7 18  
3 3 4 19  
4 4 10 20  
5 4 13 24  
6 6 12 28
```

```
#fit regression model
```

```
model <- lm(y~x1+x2, data=df)
```

```
#calculate MAE between predicted values and observed values
```

```
mae(df$y, predict(model))
```

```
1.238241
```

The mean absolute error (MAE) turns out to be 1.238.

This tells us that the average absolute difference between the observed values and the predicted values is 1.238.

In general, the lower the value for the MAE the better a model is able to fit a dataset. When comparing two different models, we can compare the MAE of each model to know which one offers a better fit to a dataset.