What is the interpretation of the intercept in a regression model and can you provide some examples?

Authored by

stats writer

May 2, 2024

The intercept in a regression model is the value of the dependent variable when all independent variables are equal to zero. It represents the baseline or starting point of the regression line and indicates the expected value of the dependent variable in the absence of any other factors.

For example, in a simple linear regression model predicting a person's salary (dependent variable) based on their years of experience (independent variable), the intercept would represent the expected salary of someone with zero years of experience.

Similarly, in a multiple regression model predicting housing prices (dependent variable) based on variables such as location, square footage, and number of bedrooms (independent variables), the intercept would represent the expected price of a house with zero square footage, zero bedrooms and located at the origin point (0,0) on the map.

In summary, the intercept in a regression model provides valuable insights into the starting point of the relationship between the dependent and independent variables and helps to interpret the results of the regression analysis.

Interpret the Intercept in a Regression Model (With Examples)

The intercept (sometimes called the "constant") in a regression model represents the mean value of the response variable when all of the predictor variables in the model are equal to zero.

This tutorial explains how to interpret the intercept value in both simple linear regression and multiple linear regression models.

Interpreting the Intercept in Simple Linear Regression

A simple linear regression model takes the following

form:

? = β0 + β1(x)

where:

?: The predicted value for the response variableβ0: The mean value of the response variable when x = 0β1: The average change in the response variable for a one unit increase in xx: The value for the predictor variable

In some cases, it makes sense to interpret the value for the intercept in a simple linear regression model but not always. The following examples illustrate this.

Example 1: Intercept Makes Sense to Interpret

Suppose we'd like to fit a simple linear regression model using *hours studied* as a predictor variable and *exam score* as the response variable.

We collect this data for 50 students in a certain college course and fit the following regression model:

Exam score = 65.4 + 2.67(hours)

The value for the intercept term in this model is 65.4. This means the average exam score is 65.4 when the number of hours studied is equal to zero.

This makes sense to interpret since it's plausible for a student to study for zero hours in preparation for an exam.

Example 2: Intercept Does Not Make Sense to Interpret

Suppose we'd like to fit a simple linear regression model using *weight* (in pounds) as a predictor variable and *height* (in inches) as the response variable.

We collect this data for 50 individuals and fit the following regression model:

The value for the intercept term in this model is 22.3. This would mean the average height of a person is 22.3 inches when their weight is equal to zero.

This does not make sense to interpret since it's not possible for a person to weigh zero pounds.

However, we still need to keep the intercept term in the model in order to use the model to make predictions.

The intercept just doesn't have any meaningful interpretation for this model.

Interpreting the Intercept in Multiple Linear Regression

A multiple linear regression model takes the following form:

? = β0 + β1(x1) + β2(x2) + β3(x3) + … + βk(xk)

where:

?: The predicted value for the response variableβ0: The mean value of the response variable when all predictor variables are zeroβj: The average change in the response variable for a one unit increase in the jth predictor variable, assuming all other predictor variables are held constantxj: The value for the jth predictor variable

Similar to simple linear regression, it makes sense to interpret the value for the intercept in a multiple linear regression model sometimes but not always. The following examples illustrate this.

Example 1: Intercept Makes Sense to Interpret

Suppose we'd like to fit a multiple linear regression model using *hours studied* and *prep exams taken* as the predictor variables and *exam score* as the response variable.

We collect this data for 50 students in a certain college course and fit the following regression model:

Exam score = 58.4 + 2.23(hours) + 1.34(# prep exams)

The value for the intercept term in this model is 58.4. This means the average exam score is 58.4 when the number of hours studied and the number of prep exams taken are both equal to zero.

This makes sense to interpret since it's plausible for a student to study for zero hours and take zero prep exams before the actual exam.

Example 2: Intercept Does Not Make Sense to Interpret

Suppose we'd like to fit a multiple linear regression model using *square footage* and *number of bedrooms* as predictor variables and *selling price* as the response variable.

We collect this data for 100 houses in a certain city and fit the following regression model:

Price = 87,244 + 3.44(square footage) + 843.45(# bedrooms)

The value for the intercept term in this model is 87,244. This would mean the average selling price of a house is $87,244 when the square footage and number of bedrooms in a house are both equal to zero.

This does not make sense to interpret since it's not possible for a house to have zero square footage and zero bedrooms.

However, we still need to keep the intercept term in the model in order to use it to make predictions. The intercept just doesn't have any meaningful interpretation for this model.