

3.2: The Linear Model Function

We use regression models to predict a system's behavior by extrapolating from previously measured output values when the system is tested with known input parameter values. The simplest regression model is a straight line. It has the mathematical form:

$$y = a_0 + a_1x_1$$

where x_1 is the input to the system, a_0 is the y-intercept of the line, a_1 is the slope, and y is the output value the model predicts.

R provides the function `lm()` that generates a linear model from the data contained in a data frame. For this one-factor model, R computes the values of a_0 and a_1 using the method of least squares. This method finds the line that most closely fits the measured data by minimizing the distances between the line and the individual data points. For the data frame `int00.dat`, we compute the model as follows:

```
> attach(int00.dat)
> int00.lm <- lm(perf ~ clock)
```

The first line in this example attaches the `int00.dat` data frame to the current workspace. The next line calls the `lm()` function and assigns the resulting *linear model object* to the variable `int00.lm`. We use the suffix `.lm` to emphasize that this variable contains a linear model. The argument in the `lm()` function, `(perf ~ clock)`, says that we want to find a model where the predictor `clock` explains the output `perf`.

Typing the variable's name, `int00.lm`, by itself causes R to print the argument with which the function `lm()` was called, along with the computed coefficients for the regression model.

```
> int00.lm
Call:
lm(formula = perf ~ clock)
Coefficients:
(Intercept)      clock
 51.7871      0.5863
```

In this case, the y-intercept is $a_0 = 51.7871$ and the slope is $a_1 = 0.5863$. Thus, the final regression model is:

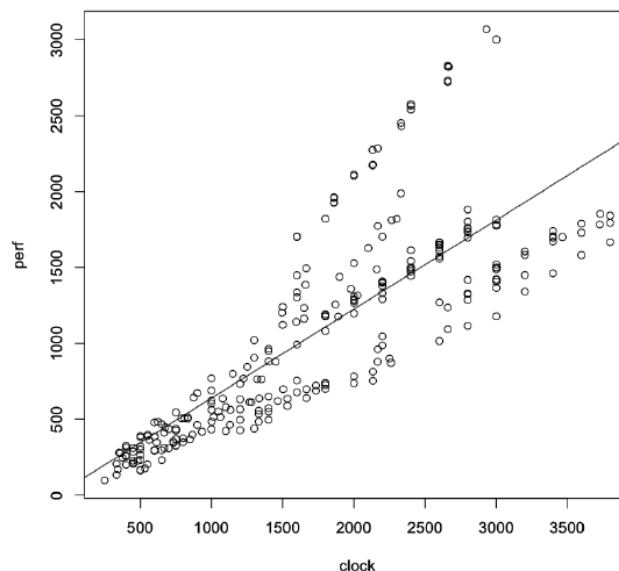


Figure 3.1.

$$\text{perf} = 51.7871 + 0.5863 * \text{clock}.$$

The following code plots the original data along with the fitted line, as shown in Figure 3.2. The function `abline()` is short for *(a,b)-line*. It plots a line on the active plot window, using the slope and intercept of the linear model given in its argument.

```
> plot(clock,perf)
> abline(int00.lm)
```

This page titled [3.2: The Linear Model Function](#) is shared under a [CC BY-NC 4.0](#) license and was authored, remixed, and/or curated by [David Lilja](#) ([University of Minnesota Libraries Publishing](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.