

Linear regression
Class 25, 18.05, Spring 2014
Jeremy Orloff and Jonathan Bloom

1 Learning Goals

1. Be able to use the method of least squares to fit a line to bivariate data.
2. Be able to give a formula for the total squared error when fitting any type of curve to data.
3. Be able to say the words homoscedasticity and heteroscedasticity.


2 Introduction

Suppose we have collected bivariate data (x_i, y_i) , $i = 1, \dots, n$. The goal of linear regression is to model the relationship between x and y by finding a function $y = f(x)$ that is a close fit to the data. The modeling assumptions we will use are that x_i is *not* random and that y_i is a function of x_i plus some random noise. With these assumptions x is called the *independent* or *predictor variable* and y is called the *dependent* or *response variable*.

Example 1. The cost of a first class stamp in cents over time is given in the following list.
.05 (1963) .06 (1968) .08 (1971) .10 (1974) .13 (1975) .15 (1978) .20 (1981) .22 (1985)
.25 (1988) .29 (1991) .32 (1995) .33 (1999) .34 (2001) .37 (2002) .39 (2006) .41 (2007)
.42 (2008) .44 (2009) .45 (2012) .46 (2013)

Using the R function `lm` we found the ‘least squares fit’ for a line to this data is $y = .026 + .8714x$, where x is the number of years since 1960 and y is in cents.

Using this result we ‘predict’ that in 2015 ($x = 55$) the cost of a stamp will be 48 cents (since $.025950 + .87138 \cdot 55 = 47.95$).




Stamp cost (cents) vs. time (years since 1960)
Red dot is predicted cost in 2015.

Note that none of the data points actually lie on the line. Rather this line has the ‘best fit’ with respect *all* the data, with a small error for each data point.

Example 2. Suppose we have n pairs of fathers and adult sons. Let x_i and y_i be the heights of the i^{th} father and son, respectively. The least squares line for this data could be used to predict the adult height of a young boy from that of his father.

Example 3. We are not limited to best fit lines. For all positive d , the method of least squares may be used to find a polynomial of degree d with the ‘best fit’ to the data. Here’s a figure showing the least squares fit of a parabola ($d = 2$).



3 Fitting a line using least squares

Suppose we have data (x_i, y_i) as above. The goal is to find the line

$$y = \beta_1 x + \beta_0$$

that ‘best fits’ the data. Our model says that each y_i is predicted by x_i up to some error ϵ_i :

$$y_i = \beta_1 x_i + \beta_0 + \epsilon_i.$$

So

$$\epsilon_i = y_i - \beta_1 x_i - \beta_0.$$

The method of least squares finds the values $\hat{\beta}_0$ and $\hat{\beta}_1$ of β_0 and β_1 that minimize the sum of the squared errors:

$$S(\beta_0, \beta_1) = \sum \epsilon_i^2 = \sum_i (y_i - \beta_1 x_i - \beta_0)^2.$$

Using calculus or linear algebra (details in the appendix), we find

$$\hat{\beta}_1 = \frac{s_{xy}}{s_{xx}} \quad \hat{\beta}_0 = \bar{y} - \hat{\beta}_1 \bar{x} \tag{1}$$

where

$$\bar{x} = \frac{1}{n} \sum x_i, \quad \bar{y} = \frac{1}{n} \sum y_i, \quad s_{xx} = \frac{1}{(n-1)} \sum (x_i - \bar{x})^2, \quad s_{xy} = \frac{1}{(n-1)} \sum (x_i - \bar{x})(y_i - \bar{y}).$$

Here \bar{x} is the sample mean of x , \bar{y} is the sample mean of y , s_{xx} is the sample variance of x , and s_{xy} is the sample covariance of x and y .

Example 4. Use least squares to fit a line to the following data: (0,1), (2,1), (3,4).


answer: In our case, $(x_1, y_1) = (0, 1)$, $(x_2, y_2) = (2, 1)$ and $(x_3, y_3) = (3, 4)$. So

$$\bar{x} = \frac{5}{3}, \quad \bar{y} = 2, \quad s_{xx} = \frac{14}{9}, \quad s_{xy} = \frac{4}{3}$$

Using the above formulas we get

$$\hat{\beta}_1 = \frac{6}{7}, \quad \hat{\beta}_0 = \frac{4}{7}.$$

So the least squares line has equation $y = \frac{4}{7} + \frac{6}{7}x$. This is shown as the green line in the following figure. We will discuss the blue parabola soon.



Least squares fit of a line (green) and a parabola (blue)


Simple linear regression: Fitting a line to bivariate data is called *simple linear regression*.

3.1 Residuals

A useful visual check of the linear regression model is to plot the residuals. For a line the model is

$$y_i = \hat{\beta}_1 x_i + \hat{\beta}_0 + \epsilon_i.$$

We think of $\hat{\beta}_1 x_i + \hat{\beta}_0$ as predicting or explaining y_i . The left-over term ϵ_i is called the *residual*, which we think of as random noise or measurement error. The data points should hover near the regression line. The residuals should look about the same across the range of x .




Data with regression line (left) and residuals (right). Note the homoscedasticity.

3.2 Homoscedasticity

An important assumption of the linear regression model is that the residuals ϵ_i have the same variance for all i . This is called *homoscedasticity*. You can see this is the case for both figures above. The data hovers in the band of fixed width around the regression line and at every x the residuals have about the same vertical spread.

Below is a figure showing *heteroscedastic* data. The vertical spread of the data increases as x increases. Before using least squares on this data we would need to transform the data to be homoscedastic.



Heteroscedastic Data

4 Linear regression for fitting polynomials

When we fit a line to data it is called *simple linear regression*. We can also use linear regression to fit polynomials to data. The use of the word linear in both cases may seem confusing. This is because the word ‘linear’ in linear regression does not refer to fitting a line. Rather it refers to the fact that the model is linear in the unknown parameters β_i (i.e. each β_i has exponent 1).

Example 5. Take the same data as in the Example 4 and use least squares to find the best fitting parabola to the data.


answer: A parabola has the formula $y = \beta_0 + \beta_1x + \beta_2x^2$. The squared error is

$$S(\beta_0, \beta_1, \beta_2) = \sum (y_i - (\beta_0 + \beta_1x_i + \beta_2x_i^2))^2.$$

After substituting the given values for each x_i and y_i , we can use calculus to find the triple $(\beta_0, \beta_1, \beta_2)$ that minimizes S . We find that the least squares parabola has equation

$$y = 1 - 2x + x^2.$$

Note that for 3 points the quadratic fit is perfect.



Least squares fit of a line (green) and a parabola (blue)

Example 6. The pairs (x_i, y_i) may give the the age and vocabulary size of a n children. Since we expect that young children acquire new words at an accelerating pace, we might guess that a higher order polynomial might best fit the data.

Example 7. (Transforming the data) Sometimes it is necessary to transform the data before using linear regression. For example, let's suppose the relationship is exponential, i.e. $y = ce^{ax}$. Then

$$\ln(y) = ax + \ln(c).$$

So we can use simple linear regression to obtain a model

$$\ln(y_i) = \hat{\beta}_0 + \hat{\beta}_1x_i$$

and then exponentiate to obtain the exponential model

$$y_i = e^{\hat{\beta}_0} e^{\hat{\beta}_1x_i}.$$


4.1 Overfitting

You can always achieve a better fit by using a higher order polynomial. For instance, given 6 data points (with distinct x_i) one can always find a fifth order polynomial that goes through all of them. This can result in what's called overfitting. That is, fitting the noise as well as the true relationship between x and y . An overfit model will fit the original data better but perform less well on predicting y for new values of x . Indeed, a primary challenge of statistical modeling is balancing model fit against model complexity.

Example 8. In the plot below, we fit polynomials of degree 1, 3, and 9 to bivariate data consisting of 10 data points. The degree 2 model (maroon) gives a significantly better fit

than the degree 1 model (blue). The degree 10 model (orange) gives fits the data exactly, but at a glance we would guess it is overfit. That is, we don't expect it to do a good job fitting the next data point we see.

In fact, we generated this data using a quadratic model, so the degree 2 model will tend to perform best fitting new data points.



4.2 R function `lm`

As you would expect we don't actually do linear regression by hand. Computationally, linear regression reduces to solving simultaneous equations, i.e. to matrix calculations. The R function `lm` can be used to fit any order polynomial to data. (`lm` stands for *linear model*). We will explore this in the next studio class. In fact `lm` can fit many types of functions besides polynomials, as you can explore using R help or google.

5 Multiple linear regression

Data is not always bivariate. It can be trivariate or even of some higher dimension. Suppose we have data in the form of tuples

$$(y_i, x_{1,i}, x_{2,i}, \dots, x_{m,i})$$

We can analyze this in a manner very similar to linear regression on bivariate data. That is, we can use least squares to fit the model

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_m x_m.$$

Here each x_j is a predictor variable and y is the response variable. For example, we might be interested in how a fish population varies with measured levels of several pollutants, or we might want to predict the adult height of a son based on the height of the mother and the height of the father.

6 Least squares as a statistical model

The linear regression model for fitting a line says that the value y_i in the pair (x_i, y_i) is drawn from a random variable

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i$$

where the ‘error’ terms ε_i are independent random variables with mean 0 and standard deviation σ . The standard assumption is that the ε_i are i.i.d. with distribution $N(0, \sigma^2)$. In any case, the mean of Y_i is given by:

$$E(Y_i) = \beta_0 + \beta_1 x_i + E(\varepsilon_i) = \beta_0 + \beta_1 x_i.$$

From this perspective, the least squares method chooses the values of β_0 and β_1 which minimize the sample variance about the line.

In fact, the least square estimate $(\hat{\beta}_0, \hat{\beta}_1)$ coincides with the maximum likelihood estimate for the parameters (β_0, β_1) ; that is, among all possible coefficients, $(\hat{\beta}_0, \hat{\beta}_1)$ are the ones that make the observed data most probable.

7 Regression to the mean

The reason for the term ‘regression’ is that the predicted response variable y will tend to be ‘closer’ to (i.e., regress to) its mean than the predictor variable x is to its mean. Here closer is in quotes because we have to control for the scale (i.e. standard deviation) of each variable. The way we control for scale is to first standardize each variable.

$$u_i = \frac{x_i - \bar{x}}{\sqrt{s_{xx}}}, \quad v_i = \frac{y_i - \bar{y}}{\sqrt{s_{yy}}}.$$

Standardization changes the mean to 0 and variance to 1:

$$\bar{u} = \bar{v} = 0, \quad s_{uu} = s_{vv} = 1.$$

The algebraic properties of covariance show

$$s_{uv} = \frac{s_{xy}}{\sqrt{s_{xx}s_{yy}}} = \rho,$$

the correlation coefficient. Thus the least squares fit to $v = \beta_0 + \beta_1 u$ has

$$\hat{\beta}_1 = \frac{s_{uv}}{s_{uu}} = \rho \quad \text{and} \quad \hat{\beta}_0 = \bar{v} - \hat{\beta}_1 \bar{u} = 0.$$

So the least squares line is $v = \rho u$. Since ρ is the correlation coefficient, it is between -1 and 1. Let’s assume it is positive and less than 1 (i.e., x and y are positively but not perfectly correlated). Then the formula $v = \rho u$ means that if u is positive then the predicted value of v is less than u . That is, v is closer to 0 than u . Equivalently,

$$\frac{y - \bar{y}}{\sqrt{s_{yy}}} < \frac{x - \bar{x}}{\sqrt{s_{xx}}}$$

i.e., y regresses to \bar{y} . Notice how the standardization takes care of controlling the scale.

Consider the extreme case of 0 correlation between x and y . Then, no matter what the x value, the predicted value of y is always \bar{y} . That is, y has regressed all the way to its mean. Note also that the regression line always goes through the point (\bar{x}, \bar{y}) .

Example 9. Regression to the mean is important in longitudinal studies. Rice (*Mathematical Statistics and Data Analysis*) gives the following example. Suppose children are given an IQ test at age 4 and another at age 5 we expect the results will be positively correlated. The above analysis says that, on average, those kids who do poorly on the first test will tend to show improvement (i.e. regress to the mean) on the second test. Thus, a useless intervention might be misinterpreted as useful since it seems to improve scores.

Example 10. Another example with practical consequences is reward and punishment. Imagine a school where high performance on an exam is rewarded and low performance is punished. Regression to the mean tells us that (on average) the high performing students will do slightly worse on the next exam and the low performing students will do slightly better. An unsophisticated view of the data will make it seem that punishment improved performance and reward actually hurt performance. There are real consequences if those in authority act on this idea.

8 Appendix

8.1 Proof of the formula for least square fit of a line

The most straightforward proof is to use calculus. The sum of the squared errors is

$$S(\beta_0, \beta_1) = \sum_{i=1}^n (y_i - \beta_1 x_i - \beta_0)^2.$$

Taking partial derivatives (and remembering that x_i and y_i are the data, hence constant)

$$\begin{aligned} \frac{\partial S}{\partial \beta_0} &= \sum_{i=1}^n -2(y_i - \beta_1 x_i - \beta_0) = 0 \\ \frac{\partial S}{\partial \beta_1} &= \sum_{i=1}^n -2x_i(y_i - \beta_1 x_i - \beta_0) = 0 \end{aligned}$$

Summing this up we get two linear equations in the unknowns β_0 and β_1 :

$$\begin{aligned} \left(\sum x_i\right) \beta_1 + n\beta_0 &= \sum y_i \\ \left(\sum x_i^2\right) \beta_1 + \left(\sum x_i\right) \beta_0 &= \sum x_i y_i \end{aligned}$$

Solving for β_1 and β_0 gives the formulas in Equation (1).

A sneakier approach which avoids calculus is to standardize the data, find the best fit line, and then unstandardize. We omit the details.

For a deluge of applications across disciplines see:

http://en.wikipedia.org/wiki/Linear_regression#Applications_of_linear_regression

8.2 Measuring the fit

Once one computes the regression coefficients, it is important to check how well the regression model fits the data (i.e., how closely the best fit line tracks the data). A common but crude ‘goodness of fit’ measure is the *coefficient of determination*, denoted R^2 . We’ll need some notation to define it. The *total sum of squares* is given by:

$$\text{TSS} = \sum (y_i - \bar{y})^2.$$

The *residual sum of squares* is given by the sum of the squares of the residuals. When fitting a line, this is:

$$\text{RSS} = \sum (y_i - \hat{\beta}_0 - \hat{\beta}_1 x_i)^2.$$

The RSS is the “unexplained” portion of the total sum of squares, i.e. unexplained by the regression equation. The difference $\text{TSS} - \text{RSS}$ is the “explained” portion of the total sum of squares. The *coefficient of determination* R^2 is the ratio of the “explained” portion to the total sum of squares:

$$R^2 = \frac{\text{TSS} - \text{RSS}}{\text{TSS}}.$$

In other words, R^2 measures the proportion of the variability of the data that is accounted for by the regression model. A value close to 1 indicated a good fit, while a value close to 0 indicates a poor fit. In the case of simple linear regression, R^2 is simply the square of the correlation coefficient between the observed values y_i and the predicted values $\beta_0 + \beta_1 x_i$.

Example 11. In the overfitting example (8), the values of R^2 are:

degree	R^2
1	0.3968
2	0.9455
9	1.0000

Notice the goodness of fit measure increases as n increases. The fit is better, but the model also becomes more complex, since it takes more coefficients to describe higher order polynomials.

MIT OpenCourseWare
<http://ocw.mit.edu>

18.05 Introduction to Probability and Statistics
Spring 2014

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.