Least squares curve fitting is a data analysis tool to extract trends in data sets. Given a data set (x_i, y_i), $i = 1, \ldots, n$, the goal is to find the coefficients of a simple function $f(x)$ (often a polynomial) that follows a trend in the data. The *least squares* method minimizes the sum of the squares of the difference between the given y_i data and the fit function $f(x_i)$. Figure 1 shows two curve fits (dashed red curves) to data represented by blue dots.

Fitting a Line to Data

- To fit a line to data, the fit function is

$$f(x) = mx + b \tag{1}$$

  where m is the slope and b.

- The goal of curve fitting is to find m and b so that Equation (1) does the best job of matching the given (x_i, y_i) data.

- Given a data set (x_i, y_i), $i = 1, \ldots, n$, the least squares method gives

$$m = \frac{n \sum x_i y_i - \sum x_i \sum y_i}{n \sum x_i^2 - \left(\sum x_i\right)^2} \tag{2}$$

$$b = \frac{\sum y_i - m \sum x_i}{n} \tag{3}$$

  where all summations are for $i = 1, \ldots, n$.

- The R^2 metric indicates how well the fit function matches the data. The value of R^2 is of the quality of the fit

$$R^2 = \frac{\sum (\hat{y}_i - \bar{y})^2}{\sum (y_i - \bar{y})^2} = 1 - \frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2} \tag{4}$$

  where $\hat{y}_i = f(x_i)$ is the fit function evaluated at the given data, and $\bar{y}$ is the average of the y_i values. For any data set, $0 \leq R^2 \leq 1$.

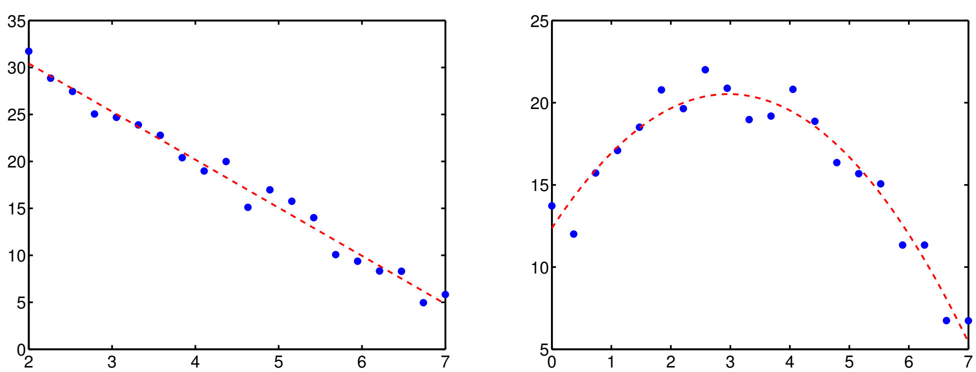

Figure 1: Linear (left) and quadratic (right) curve fits to data.

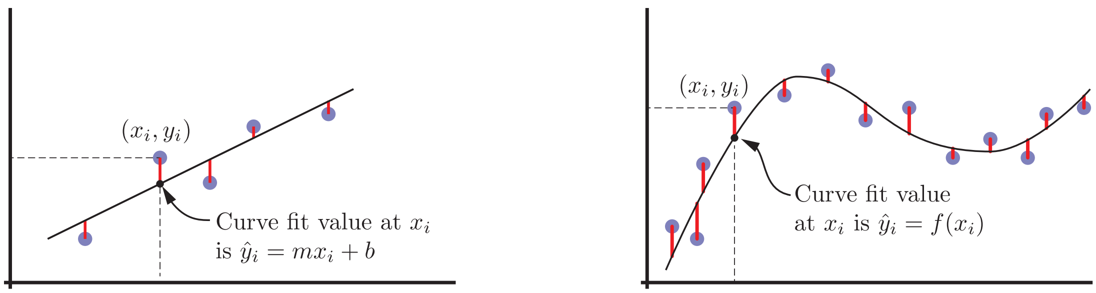

Figure 2: Geometry of a least squares fit: linear fit (left) and arbitrary $f(x)$ (right).

Geometry of a Line Fit

The left side of Figure 2 shows the relationship between the original (x_i, y_i) data and a least squares line fit. For a given data set, the values of m and b in the fit function, $f(x) = mx + b$, are computed with Equation (2) and Equation (3).

The fit function can be evaluated at any x, which is just one advantage of obtaining the fit. We use $\hat{y}_i = f(x_i)$ to designate the value of the fit function at the discrete x_i values.

The right side of Figure 2 shows a least squares fit to an arbitrary $f(x)$, which could be a polynomial or some other function. As with the line fit, the fit function $f(x)$ can be evaluated at any x. The $\hat{y}_i$ values are the fit function evaluated at the x_i from the given data set.

For both images in Figure 2, the vertical red bars are the vertical (y-direction) difference between the data and the fit function. The mathematical computations of the least squares method obtains the $f(x)$ that minimizes $(\hat{y}_i - y_i)^2$, i.e., the sum of the lengths of the red lines in Figure 2.

Fitting a Polynomial to Data

In some situations, a line does not adequately follow the trend in the data. The plot in the right hand side of Figure! 1 shows an example of a data set (blue dots) that cannot be fit by a line.

- To fit a polynomial to data, the fit function is

  $$f(x) = c_{N+1}x^N - +c_N x^{N-1} + \cdots + c_2 x + c_1 \tag{5}$$

  where c_j, $j = 1, \ldots, N+1$ are the coefficients of the degree N polynomial.

  Examples:

  $$f(x) = c_2 x + c_1 \qquad \text{linear fit} \tag{6}$$

  $$f(x) = c_3 x^2 + c_2 x + c_1 \qquad \text{quadratic fit} \tag{7}$$

  $$f(x) = c_4 x^3 + c_3 x^2 + c_4 x + c_1 \qquad \text{cubic fit} \tag{8}$$

- The goal of curve fitting is to find values of the coefficients c_j, $j = 1, \ldots, N+1$ so that Equation (5) does the best job of matching the given (x_i, y_i) data.

- Given a data set (x_i, y_i), $i = 1, \ldots, n$, formulas for finding the c_j in Equation (5) are cumbersome for $N > 1$. Instead of writing explicit formulas, we use an *algorithm*, to compute the

c_j. The algorithm can be elegantly expressed in terms of linear algebra, but the details are beyond the scope of this brief document.

- The R^2 metric for a polynomial curve fit is computed with Equation (4). In other words, Equation (4) is the general expression for R^2.

General Recommendations

- It is easy to make least squares curve fits with software like Excel, MATLAB, *R* and python. Avoid the temptation to use higher and higher degree polynomials just to increase the value of R^2. In other words

  Use the lowest degree polynomial that does an adequate job of fitting the data.

- R^2 is an indicator of how well the fit function matches the data, but there are cases where high values of R^2 can be misleading

  Always plot the fit function with the original data to inspect the quality of the fit.

- To systematically compare curve fitting options, e.g., comparing fits with different degree polynomials, compute and plot the residuals of each fit. The residual is

  $$r_i = f(x_i) - y_i \tag{9}$$

  which is the difference (measured along the y axis) between the fit function and the given data[1] If the residual shows a pattern or trend, then a different fit function is warranted. When the plot of the residual looks random, further refinement of the fit, e.g., to high degree polynomial is not advantageous, even though a higher degree polynomial will further reduce R^2.

[1] The vertical red bars in Figure 2 have a length equal to $|r_i|$.