

Lecture 20

Least Squares Fitting: Noisy Data

Suppose we have n data points $\{(x_i, y_i)\}_{i=1}^n$. In the previous section we learned about interpolation, where the fitting function $f(x)$ is required to satisfy $y_i = f(x_i)$ for $i = 1, \dots, n$. Very often data has a significant amount of noise, so it makes more sense to only ask for $y_i \approx f(x_i)$. The least squares fitting method produces such an f . The next illustration shows the effects noise can have and how least squares is used.

Traffic flow model

Suppose you are interested in the time it takes to travel on a certain section of highway for the sake of planning. According to theory, assuming up to a moderate amount of traffic, the time should be approximately

$$T(x) = ax + b$$

where b is the travel time when there is no other traffic and x is the current number of cars on the road (in hundreds). To determine the coefficients a and b you could run several experiments which consist of driving the highway at different times of day and also estimating the number of cars on the road using a counter. Of course both of these measurements will contain *noise*, i.e. random fluctuations.

We could simulate such data in MATLAB as follows:

```
>> x = 1:.1:6;  
>> T = .1*x + 1;  
>> Tn = T + .1*randn(size(x));  
>> plot(x, Tn, ' .')
```

The data should look like it lies on a line, but with noise. Click on the **Tools** button and choose **Basic fitting**. Then choose a **linear** fit. The resulting line should go through the data in what looks like a very reasonable way. Click on **show equations**. Compare the equation with $T(x) = .1x + 1$. The coefficients should be pretty close considering the amount of noise in the plot. Next, try to fit the data with a spline. The result should be ugly. We can see from this example that **splines are not suited to noisy data**.

How does MATLAB obtain a very nice line to approximate noisy data? The answer is a very standard numerical/statistical method known as *least squares*.

Least squares

The least squares method requires you to first select what type of f to consider (sometimes called a model) and what parameters it depends on. For example, you might choose

$$f(x) = a_1 \exp(a_2 x) + a_0 \quad \text{or} \quad f(x) = a_2 x^2 + a_1 x + a_0.$$

Given data $\{(x_i, y_i)\}_{i=1}^n$, we can then determine the error at each point as $e_i = y_i - f(x_i)$. To make f reasonable, we wish to simultaneously minimize all the errors $\{e_1, e_2, \dots, e_n\}$. There are many possible ways one could go about this, but the standard one is to try to minimize the *sum of the squares* of the errors. That is, we denote by \mathcal{E} the sum of the squares

$$\mathcal{E} = \sum_{i=1}^n e_i^2 = \sum_{i=1}^n (y_i - f(x_i))^2 .$$

Since f depends on some parameters, \mathcal{E} is a function of those parameters. To minimize \mathcal{E} , we can compute the partial derivatives of it with respect to each of the parameters, set the results equal to zero, and solve the resulting system of equations.

Linear least squares

If we choose f to depend linearly on the parameters, then the resulting system of equations will be a linear system, which is easy to solve. This does not mean we have to choose f to be a line. For example, the function $f(x) = a_2x^2 + a_1x + a_0$ is a quadratic in x , but depends linearly on each of a_2 , a_1 , and a_0 . To illustrate the linear least squares process, we will go through it using this quadratic f .

The error is

$$\mathcal{E}(a_0, a_1, a_2) = \sum_{i=1}^n (y_i - (a_2x_i^2 + a_1x_i + a_0))^2 ,$$

which has partial derivatives

$$\begin{aligned} \frac{\partial \mathcal{E}}{\partial a_0} &= -2 \sum_{i=1}^n (y_i - (a_2x_i^2 + a_1x_i + a_0)) , \\ \frac{\partial \mathcal{E}}{\partial a_1} &= -2 \sum_{i=1}^n (y_i - (a_2x_i^2 + a_1x_i + a_0)) x_i , \quad \text{and} \\ \frac{\partial \mathcal{E}}{\partial a_2} &= -2 \sum_{i=1}^n (y_i - (a_2x_i^2 + a_1x_i + a_0)) x_i^2 . \end{aligned}$$

Setting these equal to zero and rearranging yields the linear system

$$\begin{pmatrix} \sum_{i=1}^n 1 & \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 \\ \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 & \sum_{i=1}^n x_i^3 \\ \sum_{i=1}^n x_i^2 & \sum_{i=1}^n x_i^3 & \sum_{i=1}^n x_i^4 \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^n y_i \\ \sum_{i=1}^n y_i x_i \\ \sum_{i=1}^n y_i x_i^2 \end{pmatrix} .$$

The entries in the matrix are determined from simple formulas using the data. Since this is a linear system, it is easily solved. The process is quick and easily automated, which is one reason it is very standard. MATLAB's basic fitting tool uses this process to obtain a d degree polynomial fit whenever the number of data points is more than $d + 1$.

Drag coefficients

Drag due to air resistance is proportional to the square of the velocity, i.e. $d = kv^2$. In a wind tunnel experiment the velocity v can be varied by setting the speed of the fan and the drag can be measured directly (it is the force on the object). In this and every experiment some random noise will occur. The following sequence of commands replicates the data one might receive from a wind tunnel:

```
>> v = 0:1:60;
>> d = .1234*v.^2;
>> dn = d + .4*v.*randn(size(v));
>> plot(v,dn,'*')
```

The plot should look like a quadratic, but with some noise. Using the tools menu, add a quadratic fit and enable the “show equations” option. What is the coefficient of x^2 ? How close is it to 0.1234?

Note that whenever you select a polynomial in MATLAB with a degree less than $n - 1$ MATLAB will produce a least squares fit.

You will notice that the quadratic fit includes both a constant and linear term. We know from the physical situation that these should not be there; they are remnants of noise and the fitting process. Since we know the curve should be kv^2 , we can do better by employing that knowledge. For instance, we know that the graph of d versus v^2 should be a straight line. We can produce this easily:

```
>> z = v.^2;
>> plot(z,dn,'*')
```

By changing the independent variable from v to $z = v^2$ we produced a plot that looks like a line with noise. Add a linear fit. What is the linear coefficient? This should be closer to 0.1234 than using a quadratic fit.

The second fit still has a constant term, which we know should not be there. If there was no noise, then at every data point we would have $k = d/v^2$. We can express this as a linear system $z'*k = dn'$, which is badly overdetermined since there are more equations than unknowns. Since there is noise, each point will give a different estimate for k . In other words, the overdetermined linear system is also inconsistent. When MATLAB encounters such systems, it automatically gives a least squares solution of the matrix problem, i.e. one that minimizes the sum of the squared errors, which is exactly what we want. To get the least squares estimate for k , do

```
>> k = z'\dn'
```

This will produce a number close to .1234.

Note that this is an application where we have physical knowledge. In this situation extrapolation would be meaningful. For instance we could use the plot to find the predicted drag at 80 mph.

Exercises

- 20.1 Find two tables of data in an engineering textbook or engineering website. Plot each (with '*' at data points) and decide if the data is best represented by a polynomial interpolant, spline interpolant, or least squares fit polynomial. Label the axes and include a title. Turn in the best plot of each. Indicate the source and meaning of the data.
- 20.2 The following table contains information from a chemistry experiment in which the concentration of a product was measured at one minute intervals:

Time	0	1	2	3	4	5	6	7
Concentration	3.033	3.306	3.672	3.929	4.123	4.282	4.399	4.527

Plot this data. Suppose it is known that this chemical reaction should follow the law: $c = a - b \exp(-0.2t)$. Following the example in the notes about the drag coefficients, change one of the variables so that the law is a linear function. Then plot the new variables and use the linear fit option to estimate a and b . What will be the eventual concentration? Finally, plot the graph of $a - b \exp(-0.2t)$ on the interval $[0,15]$ (as a solid curve), along with the data from the table (using '*').