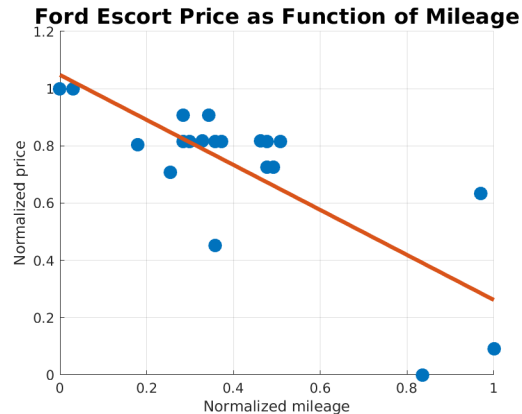


Optimization: Linear Regression

MATH2070: Numerical Methods in Scientific Computing I

Location: http://people.sc.fsu.edu/~jburkardt/classes/math2070_2019/optimization_linear_regression/optimization_linear_re



Find coefficients for an approximating formula to data.

Linear Regression Problem

Find the coefficients in a formula that will minimize the sum-of-squares error to given data.

1 Simple linear regression problem

In a previous class, we looked at a problem in which we wanted to fit $n = 23$ Ford Escort mileage (x) and price (y) data values with a simple linear formula:

$$y = b + m * x$$

We first normalized our data so that $0 \leq x, y \leq 1$. Then we used gradient descent to estimate the values of b and m :

```
1  bm0 = [ 0.5, 0.0 ];
2  r = 0.01;
3  dxtol = 0.001;
4  dftol = 0.001;
5  itmax = 1000;
6  [ bm, it ] = gradient_descent2 ( @(x)ford_f(x), @(x)ford_df(x), bm0, ...
7  r, dxtol, dftol, itmax );
```

Listing 1: Calling `gradient_descent2` for the Ford data.

This calculation came up with values of b and m for the formula:

$$y = 1.04802 - 0.796532 * x$$

Since we have 23 sets of data, and a straight line could only exactly match two such sets, we must expect that our formula is only an approximation. For each data index i , we therefore have an error:

$$e_i = y_i - 1.04802 + 0.796532 * x_i$$

We choose as our error function the sum of the squares of these individual errors:

$$E = \sum_{i=1}^{23} e_i^2$$

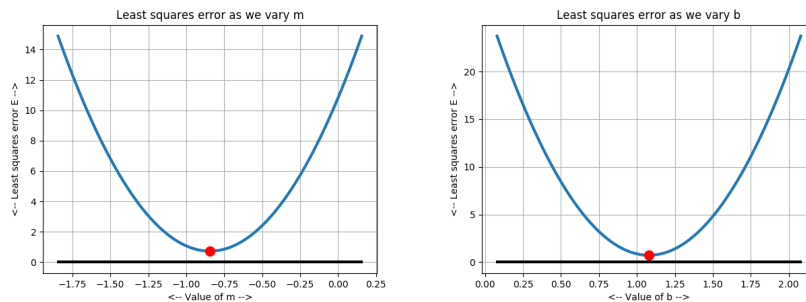
The fact that we say that b and m should be the “best” values is equivalent to the following statement:

*For all possible values of b and m in the formula $y = b + m * x$, the values we are seeking should result in the least possible value of the sum-of-squares error E .*

2 Features of the sum-of-squares error E

When we say that b and m are the best values, that means that choosing any other values will make the error worse, at least when measured using our sum-of-squares error.

To verify this, let’s look at what happens if we choose other values. A simple check is to hold one variable fixed, and vary the other. We get two plots, both of which look like parabolas, and both of which are minimized when we use the numbers we computed by gradient descent:



Our values of m and b are the best for minimizing the error.

The fact that both graphs are parabolas is not an accident. In fact, if we think of the error as a function $E(b, m)$ where both b and m are free to vary, then a plot would display a 3d shape that looks like a cup whose bottom is at the optimal values. This suggests that we could solve our linear regression problem in a different, simpler way.

The error function $E(b, m)$ is a kind of parabolic shape. It is a quadratic in each of the variables b and m . Calculus tells us that this function will be minimized when the partial derivatives $\frac{\partial E}{\partial b}$ and $\frac{\partial E}{\partial m}$ are zero.

Because E is a parabola, we have a linear system of equations to solve:

$$E = \sum_{i=1}^n (y_i - b - m * x_i)^2$$

$$\frac{\partial E}{\partial b} = -2 * \sum_{i=1}^n y_i - b - m * x_i$$

$$\frac{\partial E}{\partial m} = -2 * \sum_{i=1}^n (y_i - b - m * x_i) * x_i$$

which gives us the linear system:

$$\sum_{i=1}^n \begin{pmatrix} 1 & x_i \\ x_i & x_i^2 \end{pmatrix} \begin{pmatrix} b \\ m \end{pmatrix} = \sum_{i=1}^n \begin{pmatrix} y_i \\ x_i y_i \end{pmatrix}$$

To simplify things, consider the following matrix:

$$A = \begin{pmatrix} 1 & x_1 \\ 1 & x_2 \\ \dots & \dots \\ 1 & x_n \end{pmatrix}$$

Then our linear system actually has the form

$$A' A \begin{pmatrix} b \\ m \end{pmatrix} = A' y$$

This system is known as the **normal equations**. The matrix A is easy to define; it turns out that, in general, the matrix $A'A$ will be square and nonsingular, and hence we can solve the linear system using a standard linear equation solver.

3 Example: Normal equations for the projectile data

In class, we saw a problem involving the motion of a projectile, with time in seconds and height in feet:

	t	h
1		
2		
3	0.0	1100.0
4	1.0	1080.0
5	2.0	1040.0
6	3.0	960.0
7	4.0	840.0

Listing 2: The projectile data.

We are seeking a model for this data of the form

$$h = c_0 + c_1 t + c_2 t^2$$

We have three unknowns, c_0 , c_1 and c_2 , and we have 5 data values, so our matrix A will have shape 5×3 . The first column is all 1's, the second column is the t values, and the third will be the values of t^2 . Our right hand side is the h values.

To set up our linear system, we need to compute $A'A$ as our system matrix, and $A' * y$ as our system right hand side. Here we go:

```

1  [ t, h ] = projectile_data ( );
2  n = length ( t );
3
4  A = zeros ( n, 3 );
5  A(:,1) = ones(n,1);
6  A(:,2) = t;
7  A(:,3) = t.^2;
8
9  AtA = A' * A;
10 Aty = A' * h;
11
12 c = AtA \ Aty;

```

Listing 3: projectile_normal.py

```

1 projectile_normal:
2 Use normal equations to seek minimizer
3 of sum-of-squares error for projectile data.
4
5 c = ( 1097.714285714286, 4.571428571427918, -17.14285714285697)
6 f(c) = 45.71428571428567
7 f'(x,y) = ( 3.19744e-13, 1.51346e-12, 2.05347e-11 )

```

Listing 4: Output from projectile_normal.py

These values almost correspond to the values in Professor Layton’s notes, page 132:

```

1 c2 = ( 1097.714286, 4.573428573, -17.14285714 )

```

Listing 5: Projectile coefficients from Professor Layton.

Presumably, both solutions are “wrong” (because they are approximations). But is one better than the other? Actually, we can answer this question because $f(c_2)=45.71440572704612$, so the values we just computed with the normal equations are (slightly) better!

4 Exercise: The China Data

The file *china_data.txt*, contains 12 pairs *year* (our *x* variable) and *income* (our *y* variable). We are looking for a relationship of the form

$$income = b + m * year$$

that allows us to “explain” the relationship between these variables, and to predict income values are other years. Our task is to set up the linear regression system, solve it, and report the solution.

Download the data file *china_data.txt* and then try the following commands:

```

1  data = load ( 'china_data.txt' );
2  year = data(:,1);
3  income = data(:,2);
4  n = length ( year )
5
6  A = zeros ( n, 2 );
7  A(:,1) = 1.0;
8  A(:,2) = year;
9
10 AtA = A' * A;
11 Aty = A' * income;
12
13 bm = AtA \ Aty;
14

```

```
15 fprintf ( 'Best formula is income = %g + %g *year\n', bm(1), bm(2) );
```

Listing 6: Set up for China data problem

5 Example: Least squares solver for the projectile data

You may have noticed that we want to approximate the solution to the rectangular set of equations $Ax = y$, and that to do so using the normal equations, we always have to compute the transpose matrix A' and multiply through to get a square system. In a better world, you might think that some function would automatically take care of those details for you.

In fact, there is a better approach entirely, called a linear least squares solver. This lets you specify your problem using the original data, and automatically works out a way to solve the system so that the sum-of-squares error is minimized. The linear least squares solver is designed to solve **any** rectangular system of linear equations $Ax = b$, choosing the solution that minimizes the sum-of-squares of the error.

MATLAB's backslash operator solves a linear system $A * x = b$ via the command `x=A\b`.

However, this operator is not restricted to a square matrix. If the matrix is rectangular, MATLAB applies a least squares technique to produce a result. If the system is overdetermined, the result is the best in the sense that it minimizes the sum-of-squares error. If the system is underdetermined, there are multiple solutions, and MATLAB returns the solution with least L2 norm.

Knowing this, we can redo the projectile problem:

```
1 function projectile_least_squares ( )
2
3     [ t, h ] = projectile_data ( );
4     n = length ( t );
5
6     A = zeros ( n, 3 );
7     A(:,1) = ones(n,1);
8     A(:,2) = t;
9     A(:,3) = t.^2;
10
11     c = A \ h;
```

Listing 7: projectile_least_squares.m

```
1 projectile_least_squares:
2 Use a least squares solver to seek minimizer
3 of sum-of-squares error for projectile data.
4
5 c = ( 1097.714285714286, 4.571428571428354, -17.1428571428571)
6 f(c) = 45.71428571428572
7 f'(c) = (3.24718e-12, 5.35039e-12, 3.20739e-11)
```

Listing 8: Output from projectile_least_squares.m

We get almost identical optimal values for c that the normal equations approach gave us.

6 Exercise: Do Younger Players Get the Money?

The file `basketball_data.txt` contains 30 records. Each record stores values of 5 quantities for a player: number, height (centimeters), weight(pounds), sponsorship(\$), and age(years). It seems like young players are getting the most sponsorship dollars. To see if this is true, let's explore the possibility of a linear relationship of the form:

$$\textit{sponsorship} = b + m * \textit{age}$$

where age will play the role of x and sponsorship the role of y .

Download the data file *basketball_data.txt* and then try the following commands:

```
1 data = load ( 'basketball_data.txt' );
2 sponsor = data (:,4);
3 age = data (:,5);
4 n = length ( age );
5
6 A = zeros ( n, 2 );
7 A (:,1) = ones (n,1);
8 A (:,2) = age;
9
10 c = A \ sponsor;
11
12 fprintf ( 1, 'Best formula is sponsor = %g + %g * age\n', c(1), c(2) );
```

Listing 9: Set up for basketball data problem

Does your model suggest that sponsorship decreases with age?

7 Example: Studying several variables

Linear regression can be used for more complicated formulas than just $y = b + m * x$. We can ask for a good quadratic model, such as $y = c_0 + c_1 * x + c_2 * x^2$. If we have several x factors that might affect the value of the y quantity, we can ask for a formula such as $y = c_0 + c_1 * x_1 + c_2 * x_2$. Let us consider an example of this problem.

The file *insurance_data.txt* includes 1338 records about medical insurance. Each record lists, for a given person, 7 factors: **age**, **sex**, **bmi**, **kid**, **smo**, **reg**, **bil**. We seek a linear formula for the medical charges **bil**, based on the values of the first five quantities. (We ignore **reg**, which records the region of the country.)

$$bil \approx c_0 + c_1 age + c_2 sex + c_3 bmi + c_4 kid + c_5 smo$$

```
1 data = load ( 'insurance_data.txt' );
2
3 age = data (:,1);
4 sex = data (:,2);
5 bmi = data (:,3);
6 kid = data (:,4);
7 smo = data (:,5);
8 bil = data (:,7);
9
10 n = length ( age );
11
12 A = zeros ( n, 6 );
13 A (:,1) = ones (n,1);
14 A (:,2) = age;
15 A (:,3) = sex;
16 A (:,4) = bmi;
17 A (:,5) = kid;
18 A (:,6) = smo;
19
20 c = A \ bil;
21
22 e = sum ( ( bil - c(1) - c(2) * age - c(3) * sex - c(4) * bmi - c(5) * kid - c(6) * smo )
23         .^2 );
24
25 fprintf ( 1, ' Sum-of-squares error = %g\n', e );
```

Listing 10: Analyze insurance data.

```

1 insurance_least_squares
2   c1      = -12052.5
3   c2 (age) = 257.735
4   c3 (sex) = -128.64
5   c4 (bmi) = 322.364
6   c5 (kid) = 474.411
7   c6 (smo) = 23823.4

```

Listing 11: Insurance data model coefficients.

The magnitude and sign of each coefficient tells us something about the importance of each factor. We can also experiment by looking at how good this model is for predicting a medical bill. For a patient with data age=19, sex=1, bmi=24.6, kid=1, smo=0, bil=1837.23, our formula predicts a bill of \$1,120.43, off by \$716.80. A patient with age=34, sex=0, bmi=31.92, kid=1, smo=1 had a bill of \$37,701, and our formula predicted \$31,298.20. If you have any experience of medical billing, it should be no surprise that the data is rather difficult to model accurately, but at least we have a ball park estimate.

8 Computing Assignment #8

The file *homes_data.txt* stores 50 records of home sales. Each record stores 9 separate variables.

- 1: **sel**, the selling price (\$)
- 2: **ask**, asking price (\$)
- 3: **liv**, living area (square feet)
- 4: **rms**, rooms (#)
- 5: **bed**, bedrooms (#)
- 6: **bat**, bathrooms (#)
- 7: **age**, age (years)
- 8: **lot**, lot size (acres)
- 9: **tax**, taxes (\$)

We'd like to be able to predict the selling price based on the values of *liv*, *rms*, *age*, and *lot*, that is, a formula

$$sel \approx c_1 + c_2 liv + c_3 rms + c_4 age + c_5 lot$$

Write a program called *hw8.m* which

1. reads the data file,
2. computes the coefficients,
3. prints each coefficient,
4. prints the predicted selling price of a house with *liv* = 1800, *rms* = 8, *age* = 25, *lot* = 0.5

Turn in: your file *hw8.m* by Friday, October 18.