Math 310 Linear Algebra Lecture: Outline on Least Squares Probelms

1. Least Squares Problems
   a. Consider an inconsistent system $A\mathbf{x} = \mathbf{b}$ with A $m \times n$ and $\mathbf{b}$ a vector in $\mathbb{R}^m$.
   b. A vector $\hat{\mathbf{x}}$ in $\mathbb{R}^n$ is called a least squares solution if it minimizes $\|\mathbf{b} - A\hat{\mathbf{x}}\|$
   c. That is the same as saying that $\mathbf{p} = A\hat{\mathbf{x}}$ is the closest element of col A to $\mathbf{b}$.
   d. Conclusion: $\hat{\mathbf{x}}$ is a least squares solution iff $A\hat{\mathbf{x}}$ is the projection of $\mathbf{b}$ on col A iff $A^T A\ \hat{\mathbf{x}} = A^T\mathbf{b}$.

   Special Case: The matrix A^TA is invertible iff the cols of A are independent. In this case, [illegible]$(A^TA)^{-1}A^T\mathbf{b}$ and $\mathbf{p} = A(A^TA)^{-1}A^T\mathbf{b}$.

2. Solving Least Squares Problems
   a. Abstract case: given equation $A\mathbf{x} = \mathbf{b}$.
   b. First consider whether system may be consistent. If you are using software, find the rref of $[A \mid \mathbf{b}]$
   c. If it is inconsistent, find the rref of $[A^T A \mid A^T\mathbf{b}]$ to solve $A^T A\ \hat{\mathbf{x}} = A^T\mathbf{b}$.
   d. If there is a unique solution, it is $\hat{\mathbf{x}}$. If there is not a unique solution, any solution can be used. One convention is to take the solution of minimal length.
   e. The *error* incurred by the least squares solution is the distance between what we want ($\mathbf{b}$) and what we actually get ($A\hat{\mathbf{x}}$). That is, *error* = $\|\mathbf{b} - A\hat{\mathbf{x}}\|$. This is sometimes called the *residual*.
   f. Examples: See examples 1 and 2 in section 6.5

3. Application 1: McDonalds Data – estimating calories as a function of fat, carbs, and protein.
   a. The table at right shows for a list of menu items at McDonalds the number of grams of protein, carbohydrates, fat, and the number of cals.
   b. Question: is it possible to predict the number of cals for a menu item as a function of protein, carb, and fat content?
   c. Linearity Assumption: we will try to find an approximation equation of the form

      $\text{kcal} = x \cdot \text{protein} + y \cdot \text{carb} + z \cdot \text{protein}$

      where x, y, z, are constants. The goal is to find the best possible constants.
   d. Matrix formulation: If we want the equation above to hold exactly for every menu item, that is the same as saying, that the columns of the table satisfy

      $x \cdot \text{col } 1 + y \cdot \text{col } 2 + z \cdot \text{col } 3 = \text{ col } 4$

      Define the first three columns of the table to be the matrix A. Define the last column to be the vector $\mathbf{b}$. Then the preceding equation is $A\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \mathbf{b}$. This is not a consistent system, so we cannot satisfy it exactly. But we can find the best solution in the least squares sense.
   e. Using Freemat, we enter the data for A and $\mathbf{b}$. First we try to solve the exact system by rref-ing $[A \mid \mathbf{b}]$. We find a row [0 0 0 | 1] in the rref, indicating an inconsistent system.

Protein	Carb	Fat	kcal
19	31	15	327
5	30	5	186
8	94	10	500
9	0	19	206
13	3	13	180
2	14	7	125
26	41	33	563
15	30	14	307
12	30	10	255
24	33	22	424
30	32	31	524
14	37	25	432
3	26	12	220
2	29	14	253
2	32	14	260
4	49	11	308
10	66	9	383
9	62	9	362
9	60	8	352
7	46	11	310
7	53	10	328
7	46	9	289

f. Next we rref the augmented matrix $[A^T A \mid A^T \mathbf{b}]$. (Remember that in freemat we can get the transpose of *A* by entering *A*'.) Here are the results

```
                          3863    8873    4329   89497
[a'*a a'*b])        =     8873   41648   11195  301943
                          4329   11195    5569  111785


                          1.0000         0         0    3.8551
rref([a'*a a'*b]) =            0    1.0000         0    3.9999
                               0         0    1.0000    9.0354
```

g. So the least squares solution is given by $x = 3.8551$, $y = 3.9999$, $z = 9.0354$. That is, we have found $\hat{\mathbf{x}} = [3.8551 \quad 3.9999 \quad 9.0354]^T$. How good a prediction do we get for the kcal data using these coefficients? We can compare the actual data, **b**, with the predicted values, $A\hat{\mathbf{x}}$. The results, shown as rows:

```
.01 * [b a*xhat]' =

Columns 1 to 11

   3.2700  1.8600   5.0000  2.0600   1.8000  1.2500   5.6300  3.0700   2.5500  4.2400  5.2400
   3.3277  1.8445   4.9718  2.0637   1.7958  1.2696   5.6239  3.0432   2.5661  4.2330  5.2374

Columns 12 to 22

   4.3200   2.2000  2.5300  2.6000   3.0800  3.8300   3.6200  3.5200  3.1000  3.2800  2.8900
   4.2785   2.2399  2.5020  2.6220   3.1080  3.8386   3.6400  3.4697  3.1037  3.2933  2.9230
```

h. Alternatively, compare **b** with the vector of differences: $\mathbf{b} - A\hat{\mathbf{x}}$.

```
.01 * [ b (b - a*xhat)]' =

Columns 1 to 11

   3.2700  1.8600  5.0000   2.0600  1.8000   1.2500  5.6300  3.0700   2.5500  4.2400   5.2400
  -0.0577  0.0155  0.0282  -0.0037  0.0042  -0.0196  0.0061  0.0268  -0.0161  0.0070   0.0026

Columns 12 to 22

   4.3200   2.2000  2.5300   2.6000   3.0800   3.8300   3.6200  3.5200   3.1000  3.2800  2.8900
   0.0415  -0.0399  0.0280  -0.0220  -0.0280  -0.0086  -0.0200  0.0503  -0.0037 -0.0133 -0.0330
```

i. As an overall measure of the errors, we could use the so called root-mean-square – meaning the squareroot of the average of the squares of the errors. That would be found by squaring each difference from the error vector, taking the average of those, and then taking the squareroot. And that turns out to be equal to $\frac{1}{\sqrt{m}}\|\mathbf{b} - A\hat{\mathbf{x}}\| = \frac{1}{\sqrt{22}}\sqrt{(\mathbf{b} - A\hat{\mathbf{x}}) \cdot (\mathbf{b} - A\hat{\mathbf{x}})}$.

Freemat gives this value as `sqrt((b-a*xhat)'*(b-a*xhat)/22) = 2.6545`. This says roughly that on the average, we get an error of about 2.5 for each predicted value.

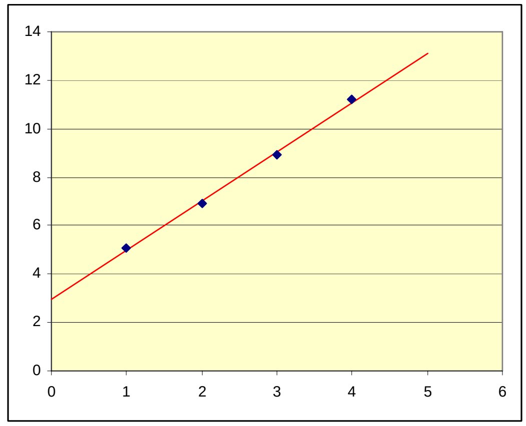

4. Application 2: Fitting a line to a set of points
   a. What line in the plane comes closest to going through the points (1,5.1), (2,6.9), (3,8.9), (4,11.2)?
   b. If the equation is expressed in the form $y = mx + b$, what we have to find are the constants m and b. Those are the unknowns.
   c. All four points would be exactly on our line if the following equations all hold:
      $5.1 = 1m + 1b$
      $6.9 = 2m + 1b$
      $8.9 = 3m + 1b$
      $11.2 = 4m + 1b$
   d. That gives us a linear system in m and b: $\begin{bmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 1 \\ 4 & 1 \end{bmatrix}\begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} 5.1 \\ 6.9 \\ 8.9 \\ 11.2 \end{bmatrix}$. We define the matrix on the left to be A and the column vector on the right to be $\mathbf{b}$. To look for solutions we rref $[A \mid \mathbf{b}]$. The system is inconsistent.
   e. Find the least squares solution. Note that the columns of A are evidently independent, so we know A^TA is invertible. The solution to the least squares problem is therefore $(A^T A)^{-1} A^T\mathbf{b} = \begin{bmatrix} .2 & -.5 \\ -.5 & 1.5 \end{bmatrix}\begin{bmatrix} 90.4 \\ 32.1 \end{bmatrix} = \begin{bmatrix} 2.03 \\ 2.95 \end{bmatrix}$. That is, the best line (in the least squares sense) is $y = 2.03x + 2.95$.
   f. What does this minimize? That is, what is the interpretation of the distance from $\mathbf{b}$ to $A\hat{\mathbf{x}}$ in this example? Each component of the error vector is of the form $y - 2.03x - 2.95$, for one of our original data points (x, y). That in turn is the vertical distance from the data point to line we have selected. So, the error measure here is the squareroot of the sum of the squares of the vertical distances from the data points to the approximating line. That is the classic definition of the *least squares line*.

5. Application 3: Fitting a polynomial
   a. Say you have n data points (x_j, y_j) and you want to fit a cubic polynomial as closely as possible. The cubic will have an equation of the form
      $y = ax^3 + bx^2 + cx + d$
      and we want to choose the constants a, b, c, d.
   b. For the curve to go exactly through the given points, substituting each data point into this equation would have to produce a true statement. That means we would have the following system of equations in the unknown coefficients:

$$y_1 = ax_1^3 + bx_1^2 + cx_1 + d$$

$$y_2 = ax_2^3 + bx_2^2 + cx_2 + d$$

$$\vdots$$

$$y_n = ax_n^3 + bx_n^2 + cx_n + d$$

where each subscripted quantity is a numerical value from our data, and the unknowns are the variables *a*, *b*, *c*, *d*. If this system is consistent, there is a cubic that goes through all of the points exactly. But if there is no such cubic, the system has to be inconsistent. In this case we can find the least squares solution.

c. The matrix *A* is similar to the one we found for a straight line. It has a column of all the cubes of the *x*'s, another for all the squares, another for the *x*'s themselves, and finally, a column of all 1's. Such a matrix always has independent columns when the *x* values are distinct, though I won't give a proof here.

d. The least squares process is the same as always – the solution is given by $(A^T A)^{-1} A^T \mathbf{b}$. The error that is being minimized is again the sum of the squared distances from the data points to the curve, just as it was in the least squares line example.

6. Application 4: Fitting a sine curve
   a. In the homework, there is a problem that asks you to fit a sine curve to a set of points.
   b. As in the prior two examples, we have an equation with some unspecified coefficients, and have to find the coefficients
   c. Here, the equation is $y = a + b\sin\left(\frac{\pi}{6}x\right) + c\cos\left(\frac{\pi}{6}x\right)$.
   d. The *x* values are the whole numbers 1 through 12, so the sines and cosines that will appear in the equations are for the angles that are multiples of $\pi/6 = 30^\circ$.
   e. Do this problem using exact values for the trig functions, using the values in the table below, and extending the table to 12 columns. When you form $A^T A$, you will find a major simplification occurs, because the columns are mutually orthogonal. That is the tip of an iceberg of colossal significance in math and applications.

	$\pi/6$	$2\pi/6 = \pi/3$	$3\pi/6 = \pi/2$	$4\pi/6$	$5\pi/6$	$6\pi/6$
Sine	1/2	$\sqrt{3}/2$	1	$\sqrt{3}/2$	1/2	0
Cosine	$\sqrt{3}/2$	1/2	0	$-1/2$	$-\sqrt{3}/2$	-1

Least Square Example:

EXAMPLE 1 Find a least-squares solution of the inconsistent system $A\mathbf{x} = \mathbf{b}$ for

$$A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$$

```
Using FreeMat:  (freemat responses shown in red)
--> A = [4 0; 0 2; 1 1]
A =
 4 0
 0 2
 1 1
--> b = [2;0;11]
b =
  2
  0
 11
--> rref([A b])
ans =
 1 0 0           This shows that the system
 0 1 0           is inconsistent
 0 0 1
--> LSaugmat = [A'*A A'*b]     This could also be computed as
LSaugmat =                     A'*[A b]
 17  1 19
  1  5 11
--> rref(LSaugmat)
ans =
 1 0 1
 0 1 2
--> xhat = [1;2]
xhat =
 1
 2
--> [A*xhat b  b-A*xhat]
ans =
  4  2 -2
  4  0 -4
  3 11  8
--> errorvec = b - A*xhat
errorvec =
 -2
 -4
  8
--> abserror = sqrt(error'*error)
abserror =
    9.1652
```