is equal to 3.

To find the desired linear combination we need to solve:

$$x\begin{pmatrix}1\\1\\0\end{pmatrix}+y\begin{pmatrix}0\\2\\3\end{pmatrix}+z\begin{pmatrix}1\\2\\3\end{pmatrix}=\begin{pmatrix}3\\6\\6\end{pmatrix}$$

or

$$\begin{pmatrix}1&0&1\\1&2&2\\0&3&3\end{pmatrix}\begin{pmatrix}x\\y\\z\end{pmatrix}=\begin{pmatrix}3\\6\\6\end{pmatrix}$$

Gaussian elimination could proceed as follows (the sequence of steps is not unique of course): first divide the third row by 3

$$\begin{array}{cccc}1&0&1&3\\1&2&2&6\\0&1&1&2\end{array}$$

$$\begin{array}{cccc}1&0&1&3\\0&2&1&3\\0&1&1&2\end{array}$$

$$\begin{array}{cccc}1&0&1&3\\0&0&-1&-1\\0&1&1&2\end{array}$$

$$\begin{array}{cccc}1&0&1&3\\0&0&1&1\\0&1&0&1\end{array}$$

$$\begin{array}{cccc}1&0&1&3\\0&1&0&1\\0&0&1&1\end{array}$$

Thus we have $z = y = 1$ and $x + z = 3$, which implies that $x = 2$. So, the solution is $(2, 1, 1)$ and you can verify that

$$2\begin{pmatrix}1\\1\\0\end{pmatrix}+1\begin{pmatrix}0\\2\\3\end{pmatrix}+1\begin{pmatrix}1\\2\\3\end{pmatrix}=\begin{pmatrix}3\\6\\6\end{pmatrix}$$

3.10 Least Squares

In this section we will consider the problem of solving $A\mathbf{x} = \mathbf{y}$ when no solution exists! I.e., we consider what happens when there is no vector that satisfies the equations exactly. This sort of situation occurs all the time in science and engineering. Often we make repeated measurements which, because of noise, for example, are not exactly consistent. Suppose we make n measurements of some quantity x. Let x_i denote the i-th measurement. You can think of this as n equations with 1 unknown:

$$\begin{pmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} x = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{pmatrix}$$

Obviously unless all the x_i are the same, there cannot be a value of x which satisfies all the equations simultaneously. Being practical people we could, at least for this simple problem, ignore all the linear algebra and simply assert that we want to find the value of x which minimizes the sum of squared errors:

$$\min_x \sum_{i=1}^{n} (x - x_i)^2 .$$

Differentiating this equation with respect to x and setting the result equal to zero gives:

$$x_{\mathrm{ls}} = \frac{1}{n} \sum_{i=1}^{n} x_i$$

where we have used x_{ls} to denote the *least squares* value of x. In other words the value of x that minimizes the sum of squares of the errors is just the mean of the data.

In more complicated situations (with n equations and m unknowns) it's not quite so obvious how to proceed. Let's return to the basic problem of solving $A\mathbf{x} = \mathbf{y}$. If $\mathbf{y}$ were in the column space of A, then there would exist a vector $\mathbf{x}$ such that $A\mathbf{x} = \mathbf{y}$. On the other hand, if $\mathbf{y}$ is not in the column space of A a reasonable strategy is to try to find an approximate solution from within the column space. In other words, find a linear combination of the columns of A that is as close as possible in a least squares sense to the data. Let's call this approximate solution $\mathbf{x_{ls}}$. Since $A\mathbf{x_{ls}}$ is, by definition, confined to the column space of A then $A\mathbf{x_{ls}} - \mathbf{y}$ (the error in fitting the data) must be in the orthogonal complement of the column space. The orthogonal complement of the column space is the left null space, so $A\mathbf{x_{ls}} - \mathbf{y}$ must get mapped into zero by A^T:

$$A^T \left(A\mathbf{x_{ls}} - \mathbf{y} \right) = 0$$

or

$$A^T A\mathbf{x_{ls}} = A^T \mathbf{y}$$

These are called the *normal equations.* Now we saw in the last chapter that the outer product of a vector or matrix with itself defined a projection operator onto the subspace spanned by the vector (or columns of the matrix). If we look again at the normal equations and assume for the moment that the matrix A^TA is invertible, then the least squares solution is:

$$\mathbf{x_{ls}} = (A^TA)^{-1}A^T\mathbf{y}$$

The matrix $(A^TA)^{-1}A^T$ is an example of what is called a *generalized inverse* of A. In the even that A is not invertible in the usual sense, this provides a reasonable generalization (not the only one) of the ordinary inverse.

Now A applied to the least squares solution is the approximation to the data from within the column space. So $A\mathbf{x_{ls}}$ is precisely the projection of the data $\mathbf{y}$ onto the column space:

$$A\mathbf{x_{ls}} = A(A^TA)^{-1}A^T\mathbf{y}.$$

Before when we did orthogonal projections, the projecting vectors/matrices were orthogonal, so A^TA term would have been the identity, but the outer product structure in $A\mathbf{x_{ls}}$ is evident.

The generalized inverse projects the data onto the column space of A.

A few observations:

- When A is invertible (square, full rank) $A(A^TA)^{-1}A^T = AA^{-1}(A^T)^{-1}A^T = I$, so every vector projects to itself.
- A^TA has the same null space as A. Proof: clearly if $A\mathbf{x} = 0$, then $A^TA\mathbf{x} = 0$. Going the other way, suppose $A^TA\mathbf{x} = 0$. Then $\mathbf{x}^TA^TA\mathbf{x} = 0$. But this can also be written as $(A\mathbf{x}, A\mathbf{x}) = \|A\mathbf{x}\|^2 = 0$. By the properties of the norm, $\|A\mathbf{x}\|^2 = 0 \Rightarrow A\mathbf{x} = 0$.
- As a corollary of this, if A has linearly independent columns (i.e., the rank $r = m$) then A^TA is invertible.

Finally, it's not too difficult to show that the normal equations can also be derived by directly minimizing the following function:

$$\|A\mathbf{x} - \mathbf{y}\|^2 = (A\mathbf{x} - \mathbf{y}, A\mathbf{x} - \mathbf{y}).$$

This is just the sum of the squared errors, but for n simultaneous equations in m unknowns. You can either write this vector function out explicitly in terms of its components and use ordinary calculus, or you can actually differentiate the expression with respect to the vector $\mathbf{x}$ and set the result equal to zero. So for instance, since

$$(A\mathbf{x}, A\mathbf{x}) = (A^TA\mathbf{x}, \mathbf{x}) = (\mathbf{x}, A^TA\mathbf{x})$$

differentiating $(A\mathbf{x}, A\mathbf{x})$ with respect to $\mathbf{x}$ yields $2A^TAx$, one factor coming from each factor of $\mathbf{x}$. The details will be left as an exercise.

3.10.1 Examples of Least Squares

Let us return to the problem we started above:

$$\begin{pmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} x = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{pmatrix}$$

Ignoring linear algebra and just going for a least squares value of the parameter x we came up with:

$$x_{\mathrm{ls}} = \frac{1}{n} \sum_{i=1}^{n} x_i.$$

Let's make sure we get the same thing using the generalized inverse approach. Now, $A^T A$ is just

$$(1, 1, 1, ..., 1) \begin{pmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} = n.$$

So the generalized inverse of A is

$$(A^T A)^{-1} A^T = \frac{1}{n} (1, 1, 1, ..., 1).$$

Hence the generalized inverse solution is:

$$\frac{1}{n} (1, 1, 1, ..., 1) \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{pmatrix} = \frac{1}{n} \sum_{i=1}^{n} x_i$$

as we knew already.

Consider a more interesting example

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}$$

Thus $x + y = \alpha$, $y = \beta$ and $2y = \gamma$. So, for example, if $\alpha = 1$, and $\beta = \gamma = 0$, then $x = 1$, $y = 0$ is a solution. In that case the right hand side is in the column space of A.

But now suppose the right hand side is $\alpha = \beta = 0$ and $\gamma = 1$. It is not hard to see that the column vector $(0, 1, 1)^T$ is not in the column space of A. (Show this as an exercise.) So what do we do? We solve the normal equations. Here are the steps. We want to solve (in the least squares sense) the following system:

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

So first compute

$$A^T A = \begin{pmatrix} 1 & 1 \\ 1 & 6 \end{pmatrix}.$$

The inverse of this matrix is

$$\left(A^T A\right)^{-1} = \frac{1}{5} \begin{pmatrix} 6 & -1 \\ -1 & 1 \end{pmatrix}.$$

So the generalized inverse solution (i.e., the least squares solution) is

$$\mathbf{x_{ls}} = \begin{pmatrix} 1 & -1/5 & -2/5 \\ 1 & 1/5 & 2/5 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} -2/5 \\ 2/5 \end{pmatrix}.$$

The interpretation of this solution is that it satisfies the first equation exactly (since $x + y = 0$) and it does an *average* job of satisfying the second and third equations. Least squares tends to average inconsistent information.

3.11 Eigenvalues and Eigenvectors

Recall that in Chapter 1 we showed that the equations of motion for two coupled masses are

$$m_1 \ddot{x}_1 = -k_1 x_1 - k_2(x_1 - x_2).$$

$$m_2 \ddot{x}_2 = -k_3 x_2 - k_2(x_2 - x_1).$$

or, restricting ourselves to the case in which $m_1 = m_2 = m$ and $k_1 = k_2 = k$

$$\begin{aligned} \ddot{x}_1 &= -\frac{k}{m} x_1 - \frac{k}{m}(x_1 - x_2) \\ &= -\omega_0^2 x_1 - \omega_0^2 (x_1 - x_2) \\ &= -2\omega_0^2 x_1 + \omega_0^2 x_2. \end{aligned} \tag{3.11.1}$$