CS6964: Notes On Linear Systems

1 Linear Systems

Systems of equations that are *linear in the unknowns* are said to be *linear systems*

For instance

$$\begin{aligned} ax_1 + bx_2 &= c \\ dx_1 + ex_2 &= f \end{aligned}$$

gives 2 equations and 2 unknowns.

More generally we have

$$\begin{aligned} a_{11}x_1 + \ldots + a_{1N}x_N &= b_1 \\ a_{21}x_1 + \ldots + a_{2N}x_N &= b_2 \\ \ldots &\quad \ldots \\ a_{M1}x_1 + \ldots + a_{MN}x_N &= b_M \end{aligned}$$

This can be written in matrix representation as

$$A\mathbf{x} = \mathbf{b} \tag{1}$$

where

$$\begin{bmatrix} a_{11} & a_{12} & \ldots & a_{1N} \\ a_{21} & a_{22} & \ldots & a_{2N} \\ & & \ldots & \\ a_{M1} & a_{M2} & \ldots & a_{MN} \end{bmatrix} \tag{2}$$

and

$$b = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_M \end{pmatrix} \tag{3}$$

The solution is given by multiplying boths sides by A^{-1}.

$$A^{-1}A\mathbf{x} = I\mathbf{x} = \mathbf{x} = A^{-1}\mathbf{b} \tag{4}$$

This is the same answer you would get if you isolated one variable at a time and then substituted back into the other equation. We will be concerned with solving linear equations with uknowns $\mathbf{x}$ under the conditions

- When $M < N$ or $M = N$ and the equations are degenerate or *singular*.

  Degeneracy happens when the equations are not linearly independent. The determinant of a square matrix is zero if and only if it's singular. The number of linearly independent equations is called the rank of A. If the rank is less than the size, the matrix is said to be *not of full rank*. Because of this, there is a set of nonzero vectors that produce zero output. These vectors can be added to any solution, and it's still a solution. Therefore the solution is not unique.

  The space spanned by the set of vectors that produce zero is called the *nullity* of A. We would like to know a single solution and the characterize the nullity of A.

- When $M = N$ and A is of full rank. We would like to find the unique solution $\mathbf{x}$ in a way that is efficient and accurate. In some cases the matrix A can be *nearly* singular. In such cases it can be impossible to compute the solution (even though it exists) because of numerical errors.

- When $M > N$, the system is *over determined*. In this case we would like to find the *best* compromise solution. Often, the best solution is defined in the sense of least squares. That is, minimize:

  $$(A\mathbf{x} - \mathbf{b})^2 \tag{5}$$

  There are many ways to do this. One way is to convert the equations into another linear $N \times N$ problem that solves this least squares. This is

  $$A^T A\mathbf{x} = A^T\mathbf{b} \tag{6}$$

  These are called the normal equations of the initial problem. The solution is

  $$\mathbf{x} = (A^T A)^{-1} A^T\mathbf{b} \tag{7}$$

  The matrix $(A^T A)^{-1} A^T$ is called the psuedo inverse. Notice, sometimes $(A^T A)^{-1}$ can be solved with a method for well-posed linear systems, but can be singular or nearly so. Therefore, the normal equations for overconstrained systems should be used with caution.

- In some cases we would like to solve

  $$A\mathbf{x} = \mathbf{b}_i \tag{8}$$

  for many different $\mathbf{b}_i$'s

2 Mechanisms for Solving Linear Systems

For well-posed systems (i.e. square matrix of full rank) there are several mechanisms for solving. Note that explicitly computing the inverse A^{-1}, is generally not recommended.

Gaussian Elimination: Somewhat robust. Can detect singularities. Not very efficient.

LU Decomposition: Somewhat robust. Efficient. Can solve for many **b**'s. Gives the determinant directly.

Iterative techniques: Can be slow. Are very accurate. Often use to clean up the accumulated round-off error associated with these other methods.

3 Singular Value Decomposition

A very powerful tool in numerical linear algrebra is *singular value decomposition.*

$$\begin{pmatrix} & & \\ & A & \\ & & \end{pmatrix} = UWV^T = \begin{pmatrix} & & \\ & U & \\ & & \end{pmatrix} \begin{pmatrix} w_1 & & & 0 \\ & w_2 & & \\ & & \dots & \\ & & \dots & \\ 0 & & & w_N \end{pmatrix} \begin{pmatrix} & V^T & \end{pmatrix} \tag{9}$$

Where the matrices U and V are orthogonal in their columns (V in rows as well). The singular values and their associated columns are defined to within a scalar factor, and therefore we usually assume the column vectors of U and V are normalized.

$$\begin{pmatrix} & U^T & \end{pmatrix} \begin{pmatrix} & & \\ & U & \\ & & \end{pmatrix} = \begin{pmatrix} & V^T & \end{pmatrix} \begin{pmatrix} & V & \end{pmatrix} \tag{10}$$

Such a decompostion is always possible, regardless of the matrix, and there are stable numerical algorithms given in "Numerical Recipes". The decomposition is unique, up to a swapping of rows in all matrices.

Why would we want to do a SVD?

3.1 SVD of Square Matrix

If A is square (say $N \times N$) then U V and W are all the same size.

Because U and V are orthogonal, inverses are the transpose.

Because W is diagonal, inverse is reciprical of diagonal elements.

So the inverse of A is

$$A^{-1} = V\,[\mathrm{diag}(1/w_j)]\,U^T \tag{11}$$

This won't work if one of the w_i's is zero, or even if one is very small (so small that it's value is dominated by roundoff error). If more than one w_i's is zero, then the number of zero w_i's gives the nullity of A.

SVD gives a mechanism for finding the inverse, and some clear indications of what's wrong when it fails. We can see this as follows. Any **x** that is composed of columns of V

with corresponding w_i's that are zero gives $A\mathbf{x} = 0$. We know this because if multiply A by column of V, call it $\mathbf{v}_i$, we get

$$A\mathbf{v}_i = UWV^T\mathbf{v}_i = UW\,(0,\ldots,1,\ldots 0) = U\,(0,\ldots,w_i,\ldots 0) = w_i\mathbf{u}_i, \tag{12}$$

where is zero when $w_i = 0$. Thus the null space is spanned by the vectors (columns of V^T) associated with zero w_is is the null space. Any point in this space when multiplied by A returns 0. The space spanned by the vectors associate with the nonzero elements of W is called the *range* of A. The dimensionality of that space is the rank. The rank plus the nullity is equal to the size of A, which is N. Therefore, the nullity of A is the dimension of its null space.

Solving Singular Problems

If A is singular, one might want to single out a solution that is somehow better than the others. In this case want might want the smallest solution, i.e. the smallest length $\mathbf{x}^2$.

One way to do that is to use the diagonals, but replace $1/w_j$ by zero where w_j is zero. I.e. use zero whereever $1/w_j$ blows up.

The solution is then

$$\mathbf{x} = V\left[\mathrm{diag}(1/w_j)\right]\left(U^T\mathbf{b}\right) \tag{13}$$

This gives the shortest solution to the problem.

Proof

Consider a solution that is modified by a vector $\mathbf{x}'$ could it be shorter by some other solution? What happens when we add a vector $\mathbf{x}'$ that is in the null space.

$$\begin{aligned} |\mathbf{x}+\mathbf{x}'| &= \left|VW^{-1}U^T\mathbf{b} + \mathbf{x}'\right| \\ &= \left|W^{-1}U^T\mathbf{b} + V^T\mathbf{x}'\right| \end{aligned}$$

But, the first term has nonzero elements only in those places where the $w_j \neq 0$. The second term, because it's in the null space, has nonzero elements only in those places where $w_j = 0$. Therefore the two terms are orthogonal vectors, and their sum must be greater length than either part.

Solving Overconstrained Problems

We can solve over constrained problems and get the least squared solution.

The solution strategy is the same

$$\mathbf{x} = V\left[\mathrm{diag}(1/w_j)\right]\left(U^T\mathbf{b}\right), \tag{14}$$

but it is quaranteed to minimize the square of the residual

$$\epsilon = A\mathbf{x} - \mathbf{b}. \tag{15}$$

Proof

Consider the solution given above, and modify it by adding some arbitrary vector $\mathbf{x}'$. Let $\mathbf{b}' = A\mathbf{x}'$. Clearly $\mathbf{b}'$ is in the range of A.

$$\begin{aligned} |A\mathbf{x} - \mathbf{b} + \mathbf{b}'| &= \left|\left(UWV^T\right)\left(VW^{-1}U^T\mathbf{b}\right) - \mathbf{b} + \mathbf{b}'\right| \\ &= \left|\left(UWW^{-1}U^T - 1\right)\mathbf{b} + \mathbf{b}'\right| \\ &= \left|U\left[\left(WW^{-1} - 1\right)U^T\mathbf{b} + U^T\mathbf{b}'\right]\right| \\ &= \left|\left(WW^{-1} - 1\right)U^T\mathbf{b} + U^T\mathbf{b}'\right| \end{aligned}$$

However $(WW^{-1} - 1)$ is a diagonal matrix with nonzero element only where $w_j = 0$. Because $\mathbf{b}'$ lies in the range of A, U^T has nonzero elements only where $w_j \neq 0$. Therefore these terms are orthongal vectors, and the minimum of their sum is obtained when $\mathbf{b}' = 0$.

3.2 Solving Homogeneous Problems

Consider a homogeneous problem of the form

$$A\mathbf{x} = 0. \tag{16}$$

If it is overconstrained, we would like to solve the associated least squares problem, which minimizes

$$\min\left[||\epsilon||^2\right] \text{ where } \epsilon = A\mathbf{x}. \tag{17}$$

The solution of such systems is not unique—it is defined up to a scalar value. Therefore, we usually look for the solution to the constrained problem, i.e. $||\mathbf{x}|| = 1$. I.e. we look for the unit length solution that produces the smallest residual.

The solution is given by SVD to be the column vector of V which corresponds to the singular value in W which has the smallest magnitude. How do we know this?

Proof

Consider the input $\mathbf{s}$ which is this vector, and assume it is position i. The output is

$$UWV^T\mathbf{s} = UW\left(0, \ldots, 1, \ldots 0\right) = U\left(0, \ldots, w_i, \ldots 0\right) = w_i\mathbf{u}_i. \tag{18}$$

Notice that $\mathbf{u}_i$ is a column of U, and it has length 1. This is the shortest output possible from this matrix, because any other unit length input would include weighted sums of $\mathbf{u}_i$s, with weights that must be larger than w_i, because it is the smallest singular value.