overdetermined Linear Systems:

- Applications where this type of problem arises.

1) Curve fitting:

Given data (t_j, f_j) $j=1,2,\ldots,m$ approximate the data with an $n < m-1$ degree polynomial

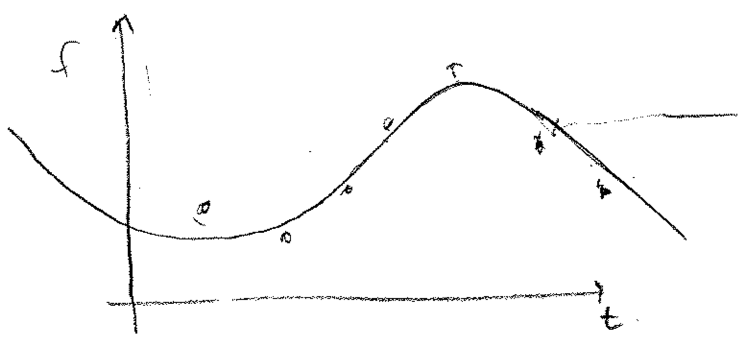

What do we mean by "approximate" $\Rightarrow$ minimize some error. It turns out that 2-norm is a good norm to minimize.

$\Rightarrow$ Least-squares problem.

Linear system. $p(t) = c_0 + c_1 t + \cdots + c_n t^n$.

Find coefficients.

$$\underbrace{\begin{bmatrix} 1 & t_1 & t_1^2 & \cdots & t_1^n \\ 1 & t_2 & t_2^2 & \cdots & t_2^n \\ \vdots & & & & \\ 1 & t_m & t_m^2 & \cdots & t_m^n \end{bmatrix}}_{A}{}_{m \times n+1} \underbrace{\begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_n \end{bmatrix}_{n+1}}_{X} = \underbrace{\begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_m \end{bmatrix}_m}_{b}$$

Over determined Linear Systems

Let $r = Ax - b$, then the least-squares problem is to find x that minimizes $\|r\|_2 = \|Ax-b\|_2$

i.e. $x^* = \min_{x \in \mathbb{R}^n} \|Ax-b\|$ is the solution

2) Statistical Modeling.

- One often wishes to estimate certain parameters x_j based on several observations, where the observations have some noise.

<u>Example:</u> Grades.

Suppose we wish to predict the final grade (b) of a student in beginning calculus based on their GPA ($v^{(1)}$), class standing ($v^{(2)}$), major ($v^{(3)}$), SAT Quantitative score ($v^{(4)}$), score on their first homework assignment ($v^{(5)}$).

Based on past data, construct a linear model of the form

$$b(v^{(1)}, v^{(2)}, \dots, v^{(5)}) = \sum_{j=1}^{5} c_j v^{(j)}$$

Now, we make several observations using the data from past students $v_k^{(1)}, v_k^{(2)}, \dots, v_k^{(5)}$ & b_k $k = 1, 2, \dots, m$ try to "fit" this data to our linear model: This means minimizing.

$$\begin{bmatrix} v_1^{(1)} & v_1^{(2)} & \cdots & v_1^{(5)} \\ v_2^{(1)} & v_2^{(2)} & \cdots & v_2^{(5)} \\ \vdots & \vdots & \ddots & \vdots \\ v_m^{(1)} & v_m^{(2)} & \cdots & v_m^{(5)} \end{bmatrix} \underbrace{\begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_5 \end{bmatrix}}_{x} - \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix} = 0 \implies \|Ax - b\|_2$$

Overdetermined Linear Systems.

Normal Equations.

Given the overdetermined linear system.

$$\begin{bmatrix} \\ A \\ \\ \end{bmatrix}_{m\times n} \begin{bmatrix} x \end{bmatrix} = \begin{bmatrix} \\ b \\ \\ \end{bmatrix}.$$

Assume A has Full column rank
(i.e the basis formed by the columns of A has dimension n)

The x that minimizes $\|Ax-b\|_2$ is given by the solution of the linear system $A^TAx = A^Tb$

Theorem: If x satisfies $A^TAx = A^Tb$ then
$$\|Ax-b\|_2 \le \|Ay-b\|_2 \quad \forall y \in \mathbb{R}^n$$

Proof:

Let $r_x = Ax-b$ & $r_y = Ay-b$ for some $y \in \mathbb{R}^n$.

Note that since $A^TAx = A^Tb$

$$\Rightarrow A^TAx - A^Tb = A^T(Ax-b) = A^Tr_x = 0$$

similarly $r_x^TA = 0$

Now,

$$r_y = Ay-b = Ax-b+A(y-x) = r_x + A(y-x).$$

$$\Rightarrow \|r_y\|_2^2 = (r_x + A(y-x))^T(r_x + A(y-x)).$$

$$= \|r_x\|_2^2 + (y-x)^T\underbrace{A^Tr_x}_{0} + \underbrace{r_x^TA}_{0}(y-x) + (y-x)^TA^TA(y-x).$$

$$= \|r_x\|_2^2 + \|A(y-x)\|_2^2 \quad \Rightarrow \quad \|r_x\|_2 \le \|r_y\|_2$$

\#

Overdetermined Linear Systems.

- A^TA is symmetric and positive definite (prove this), thus we can use Cholesky's Method to solve

$$A^TAx = A^Tb.$$

What will be the total cost?

$$A_{m\times n} \Rightarrow \underbrace{A^T_{n\times m}A_{m\times n}}_{} = B_{n\times n} \Rightarrow n^2m \text{ Flops.}$$

$$B_{n\times n} \Rightarrow L^TL \Rightarrow \tfrac{1}{3}n^3 \text{ Flops}$$

Since $m > n$ the cost in solving $A^TAx = A^Tb$ with Cholesky is dominated by $O(mn^2)$ cost in forming A^TA.

QR Decomposition Solution to Least square problem.

- Recall that for a non-singular square matrix $A \in \mathbb{R}^{n\times n}$, we decompose A using Householder matrices to

$$A = QR$$

- We can also apply this idea to non-square A with full column rank.
- Let A be $m\times n$ with $m \geq n$

$$\underbrace{\begin{bmatrix} & H^{(n)} & \end{bmatrix} \begin{bmatrix} z \end{bmatrix}}_{\substack{m\times m \\ n\times m}} \begin{bmatrix} & A & \end{bmatrix}_{m\times n} = \begin{bmatrix} x & x & - & x \\ 0 & x & - & x \\ 0 & & & \\ \vdots & & & \\ 0 & x & - & x \end{bmatrix}$$

Overdetermined Linear Systems

Repeat process to get

$$\underbrace{\begin{bmatrix} \quad H^{(n)} \quad \end{bmatrix}\begin{bmatrix} \quad H^{(n-1)} \quad \end{bmatrix} \cdots \begin{bmatrix} \quad H^{(1)} \quad \end{bmatrix}_{m\times m}}_{\tilde{Q}^T} \begin{bmatrix} \quad A \quad \end{bmatrix}_{m\times n} = \underbrace{\overbrace{\begin{bmatrix} \times & & \\ & \ddots & \\ 0 & & \times \end{bmatrix}}^{n}\Big\}n}_{\tilde{R}\ \ m\times n}$$

$$\Rightarrow \begin{bmatrix} \quad A \quad \end{bmatrix}_{m\times n} = \begin{bmatrix} \overbrace{\quad \check{Q} \quad}^{n} & \quad \end{bmatrix}_{m} \underbrace{\begin{bmatrix} \overbrace{R}^{n} \\ 0 \end{bmatrix}}_{\hat{R}}$$

Q   Doesn't matter.   R

Let Q be the first n columns of $\check{Q}$ and R be the first n rows of $\tilde{R}$.

$$\Rightarrow \quad A = QR.$$

<u>Theorem</u>: let $A \in \mathbb{R}^{m\times n}$ have full column rank then the x that minimizes $\|Ax-b\|_2$ is given by the

$$x = R^{-1}(Q^T b)$$

<u>Proof</u>: From before, we know that the if x is a solution of $A^TAx = b$ then x minimizes $\|Ax-b\|_2$

Let $A = QR$

$$\Rightarrow A^TAx = (QR)^T(QR)x = R^TQ^TQRx = R^TRx$$

$$\Rightarrow R^TRx = A^Tb = R^TQ^Tb \quad \Rightarrow \quad Rx = Q^Tb.$$

$$\Rightarrow \quad x = R^{-1}(Q^Tb) \qquad \#$$

Overdetermined Linear Systems

- Let's look at how the QR solution to $A^TAx = A^Tb$ is computed in practice.

Two procedures.

1) If you know the right-hand side, simply augment A with b and use Householder matrices to reduce.

Like Gaussian Elimination, write

$$Ax = b \Rightarrow \left[\begin{array}{c|c} A & b \end{array}\right]_{m\times(n+1)}$$

Throw Householder matrices on the left to reduce the augmented matrix

$$\Rightarrow \underbrace{\left[H^{(n)}\right]}_{m\times m}\left[H^{(n-1)}\right]\cdots\left[H^{(1)}\right]\left[\begin{array}{c|c} A & b\end{array}\right] = \left[\begin{array}{c|c} R & g_1 \\ \hline 0 & g_2\end{array}\right]$$

Solve

$$\left[R\right]\left[x\right] = \left[g_1\right]$$ with back substitution.

Total operation count to leading order is $2n^2m - \frac{2}{3}n^3$

This is about twice the cost of the normal equations-cholesky approach for $m \gg n$ and the same if $m = n$.

2) If you don't know the RHS then we do the following steps.

i) Reduce A to upper triangular form using Householder matrices

$$\Rightarrow H^{(n)}H^{(n-1)}\cdots H^{(1)}A = \overbrace{\begin{bmatrix} R \\ 0 \end{bmatrix}}^{n}\Big\}n$$

Overdetermined Linear Systems

2) (cont.)

But, do not compute Q. Simply store each $\omega^{(k)}$, $k=1,\dots,n$, that went into determining $\tilde{H}^{(k)}$

Recall that $\tilde{H}^{(k)} = I - 2\omega^{(k)}\omega^{(k)T}$ where

$$\underbrace{\begin{bmatrix} \\ \omega^{(1)} \\ \\ \end{bmatrix}}_{m} \underbrace{\begin{bmatrix} \\ \omega^{(2)} \\ \end{bmatrix}}_{m-1} \cdots \underbrace{\begin{bmatrix} \omega^{(n)} \end{bmatrix}}_{m-n+1}$$

(These can be stored in the columns of $\begin{bmatrix} R \\ 0 \end{bmatrix}$)

ii) When we are given a RHS $b \in \mathbb{R}^m$ compute $Q^T b$ using the loop

$$\begin{aligned} &\text{for } j=1 \text{ to } n \\ &\quad b(j:m) = b(j:m) - 2\omega^{(j)}\left(\omega^{(j)T} b(j:m)\right) \\ &\text{end} \end{aligned}$$

This is equivalent to computing

$$Q^T b = H^{(n)} H^{(n-1)} \cdots H^{(1)} b = \tilde{b}$$

But, without the cost of computing Q directly.

iii) Solve $Rx = \tilde{b}$ using back substitution.

Total cost for first solve is again $2n^2m - \frac{2}{3}n^3$, each subsequent solve is just $\mathcal{O}(n^2)$

- While QR solution to the least-squares problem is more expensive, it is also more stable. This is due to the fact that we never compute A^TA and we use orthogonal matrices in the computation.

Overdetermined Linear Systems

Singular Value Decomposition (SVD)

We now discuss a method that can be used to effectively solve rank deficient least squares problem. The SVD.

- While this is one application of the SVD, it is by no means the only. It is used extensively in image processing, to represent compact linear operators efficiently, and statistics to name a few.

  See Trefethen and Bau: "Numerical Linear Algebra" for an overview of its application.

- Theorem: Let $A \in \mathbb{R}^{m \times n}$, then $\exists$ orthogonal matrices $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ s.t.

$$U \Sigma V^T = A$$

where $\Sigma \in \mathbb{R}^{m \times n}$ and

$$\Sigma = \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & 0 \\ & & \ddots & \\ & & & \sigma_n \\ & 0 & & \end{bmatrix} \qquad \sigma_1 \geq \sigma_2 \geq \sigma_3 \geq \cdots \geq \sigma_n \geq 0$$

We call σ_j the singular values of A, the columns of U. $u_1, u_2, \ldots, u_m$ the left singular vectors and the columns of V, $v_1, v_2, \ldots, v_n$ the right singular vectors.

Proof: (see Atkinson p. 478)

Overdetermined Linear Systems

- Some important properties of the SVD of a matrix.

1) If A is square and symmetric with eigenvalues $\lambda_1, \dots, \lambda_n$, then $\sigma_j = |\lambda_j|$ and $u_1, \dots, u_n$ are the corresponding eigenvectors of A.

   Proof: Definition

2) σ_j^2 are the eigenvalues of the symmetric matrix $\underset{n \times n}{A^TA} \in R^{n \times n}$ and V_j are the corresponding eigenvectors

   Proof: Let $A = U\Sigma V^T$

$$\Rightarrow A^TA = V\Sigma^T U^T U \Sigma V^T = V\Sigma^T\Sigma V^T = V\begin{bmatrix} \sigma_1^2 & & 0 \\ & \ddots & \\ 0 & & \sigma_n^2 \end{bmatrix}V^T$$

$\Rightarrow$ Principal axis theorem implies σ_j^2 are eigenvalues and V eigenvectors

3) The eigenvalues of the symmetric matrix $AA^T \in \mathbb{R}^{m \times m}$ are σ_j^2 $j = 1, \dots, n, n+1, \dots, m$ where $\sigma_{n+1, \dots, m}^2 = 0$ and u_j are the corresponding eigenvectors

   Proof:

$$AA^T = (U\Sigma V^T)(V\Sigma^T U^T)$$
$$= U\Sigma\Sigma^T U^T = U\begin{bmatrix} \sigma_1^2 & & & \\ & \ddots & & \\ & & \sigma_n^2 & \\ & & & 0 \\ & & & & \ddots \\ & & & & & 0 \end{bmatrix}U^T$$

$\Rightarrow$ Principal axis theorem.

4) $U_j \in \mathbb{R}^m$ $j = 1, \dots, m$ are an orthonormal basis for $\mathbb{R}^m$

   $V_j \in \mathbb{R}^n$ $j = 1, \dots, n$ are an orthonormal basis for $\mathbb{R}^n$.

Proof: Follows from the fact that they are eigenvectors of A^TA and AA^T, respectively.

Overdetermined Linear Systems.

A maps vectors in $\mathbb{R}^m$ to vectors in $\mathbb{R}^n$

5) If $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > \sigma_{r+1} = \cdots = \sigma_n = 0$ then the rank of A is r. Furthermore, the range space of A (i.e. all the subspace of vectors of the form 'Aw for all w)' is $\text{span}\{u_1, \ldots, u_r\}$ and the null space $N(A)$ is $\text{Span}\{v_{r+1}, \ldots, v_n\}$. Similarly the range space of A^T is $\text{span}\{v_1, \ldots, v_r\}$ and the left null space $N(A^T)$ is $\text{span}\{u_{r+1}, \ldots, u_n\}$.

Proof: See Trefethen & Bau

6) If A is square then $\|A\|_2 = \sigma_1$ and $\|A^{-1}\|_2 = \frac{1}{\sigma_n}$

$\Rightarrow \text{cond}(A)_2 = \dfrac{\sigma_1}{\sigma_n} \Rightarrow$ Ratio of the largest to the smallest singular value

Proof: $\|A\|_2 = \underbrace{\|U^T A V\|_2}_{\text{since unitary}} = \|\Sigma\|_2 = \sigma_1$

Similar argument for $\|A^{-1}\|_2$.

SVD solution to the least squares problem

- Consider the over determined linear system

$$Ax = b \quad \text{where } A \in \mathbb{R}^{m \times n},\ x \in \mathbb{R}^n,\ b \in \mathbb{R}^m$$

- First assume that A has full column rank. (i.e. rank$(A) = n$)

Let $A = \left[\ U\ \right]_{m\times m} \underbrace{\begin{bmatrix} \sigma_1 & & 0 \\ & \ddots & \\ 0 & & \sigma_n \end{bmatrix}_{m\times n}}_{\Sigma} \left[\ V^T\ \right]_{n\times n}$

Define $\hat{\Sigma} = \begin{bmatrix} \sigma_1 & & 0 \\ & \ddots & \\ 0 & & \sigma_n \end{bmatrix}_{n\times n}$

Overdetermined Linear Systems

- Theorem: If A has full column rank then the x that minimizes $\|Ax-b\|_2$ is given by.

$$x = V\hat{\Sigma}^{-1}U^T b$$

Proof! From normal equations, we have.

$$A^TAx = (V\hat{\Sigma}U^T)(U\hat{\Sigma}V^T)x = V\hat{\Sigma}\hat{\Sigma}V^Tx$$

$$A^Tb = V\hat{\Sigma}U^Tb$$

$\Rightarrow$ The solution to $A^TAx = A^Tb$ is given by.

$$x = V\hat{\Sigma}^{-1}U^Tb$$

#

since $\hat{\Sigma}V^Tx = U^Tb$

- In reality we simply solve the linear systems

$$\hat{\Sigma}z = U^Tb \quad \& \quad x = Vz$$

Moore-Penrose Pseudo-inverse for rectangular matrix
$\Rightarrow A^+ = V\hat{\Sigma}^{-1}U^T$

- Rank Deficient overdetermined linear systems . $Ax=b$

- If A is rank deficient, i.e. $\text{rank}(A) = r < n$, then the singular values of A satisfy. $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > \sigma_{r+1} = \cdots = \sigma_n = 0$

- If $\text{rank}(A) = r < n$ then there is an $n-r$ dimensional set of vectors x that minimize $\|Ax-b\|_2$.

- The SVD solution to a rank deficient problem allows us to find any of the solutions to the rank deficient least squares problem. However, the main bennefit is that it allows us to compute the unique vector x that minimizes $\|Ax-b\|_2$ and has minimum norm $\|x\|_2$.

* Let T be the space of solutions that minimize $\|Ax-b\|_2$, then find $x^* \in T$ subject to the constraint $\|x^*\| \leq y \quad \forall y \in T$

Overdetermined Linear Systems

For rank deficient problem with $\text{rank}(A) = r < n$ the SVD of A is given as

$$A = \begin{bmatrix} & & \\ & U & \\ & & \end{bmatrix} \begin{bmatrix} \sigma_1 & & & & 0 \\ & \ddots & & & \\ & & \sigma_r & & \\ & & & 0 & \\ 0 & & & & \ddots \; 0 \end{bmatrix} \begin{bmatrix} V^T \end{bmatrix}$$

Let $\underbrace{U_1 = \begin{bmatrix} u_1 & u_2 & \cdots & u_r \end{bmatrix}}_{\substack{\text{First } r \text{ columns of } U \\ m\times r}}$, $\underbrace{V_1 = \begin{bmatrix} v_1 & v_2 & \cdots & v_r \end{bmatrix}}_{\substack{\text{First } r \text{ columns of } V \\ n\times r}}$ $\Sigma_1 = \underbrace{\begin{bmatrix} \sigma_1 & & 0 \\ & \ddots & \\ 0 & & \sigma_r \end{bmatrix}}_{r\times r}$

Theorem: The solution x that minimizes $\|Ax-b\|_2$ and has minimal norm $\|x\|_2$ is given by

$$x = V_1 \Sigma_1^{-1} U_1^T b$$

Proof: Use the result that all solutions that minimizes $\|Ax-b\|_2$ can be written as $x = V_1\Sigma_1^{-1}U_1^T b + V_2 z$, $z \in \mathbb{R}^{n-r}$

where $V_2 = \begin{bmatrix} v_{r+1} & \cdots & v_n \end{bmatrix}$

$\Rightarrow \|x\|_2 = \|V_1\Sigma_1^{-1}U_1^T b\|_2^2 + \|V_2 z\|_2^2 \Rightarrow$ minimized when $z = 0$.

- Note that even if A has full rank, $\rightarrow r=n$ we can still use the above procedure to create "rank k" (for $k<n$) approximation to the least-squares solution. This has important applications to image processing

Overdetermined Linear Systems

- Application of SVD to image compression.

- An $m \times n$ image can be thought of as an $m \times n$ matrix with each entry interpreted as the intensity of the corresponding pixel.

- For large images it is often expensive to store or transmit all the $m \cdot n$ entries of the matrix. Instead we would like to compress the image so that fewer numbers need to be stored, from which we can still reconstruct an accurate version of the original image. (Demmel)

- The SVD is one tool for doing this.

- Let $A = m \times n$ matrix representing the image.

  then $$A = \underset{m \times m}{U} \, \underset{m \times n}{\Sigma} \, \underset{n \times n}{V}^T = \sum_{j=1}^{n} \sigma_j u_j v_j^T$$

- We can thus form the best rank-k approximation of A by

$$A_k = \sum_{j=1}^{k} \sigma_j u_j v_j^T$$

MATLAB:

```
[u,s,v]=svd(A);
k=3;
B=0;
for j=1:k
    B=B+s(j,j)*u(:,j)*v(:,j)';
end
```

i.e, A_k is the rank k matrix that minimizes $\|A - A_k\|_2$.

Note that with this decomposition it is only necessary to store u_j & $\sigma_j v_j$ $j = 1, \ldots, k$. This amounts to storing

$m \cdot k + n \cdot k$ numbers instead of $m \cdot n$.

Overdetermined Linear Systems

- The error in the compressed image is given by

$$\|A - A_k\|_2 = \sigma_{k+1}$$

- Two measures that we are interested in looking at for a compressed image are, relative error and the compression ratio these are given by

$$\text{Rel. Error} = \sigma_{k+1}/\sigma_k$$

$$\text{Compression Ratio} = \frac{(m+n)\cdot k}{m \cdot n}$$

We would like to balance these two values

- If the singular values of A drop off rapidly then A can be compressed quite significantly and the relative error will be quite small. This typically occurs for images with a lot of "regularity".

- Computing the SVD

  - Bidiagonalization of A via orthogonal matrices

$$\Rightarrow U^{(n)}U^{(n-1)}\cdots U^{(1)} A V^{(1)} V^{(2)} \cdots V^{(n)} = \begin{bmatrix} x & x & & & 0 \\ x & x & x & & \\ & x & x & x & \\ & & \ddots & \ddots & x \\ & & & x & x \\ & 0 & & & \\ & & & & \end{bmatrix}$$

Use something similar to QR iteration to "diagonalize"

To compute U, Σ, and V requires $4m^2n + 8mn^2 + 9n^3$ FLOPS

- See Golub and Van Loan p. 251 - 254