Least Squares Problems.

We are interested in solving the problem

$$Ax = b \quad \text{when} \quad b \notin \operatorname{Im} A$$

In this case, we are trying to solve the problem in the sense

$$\min \|Ax-b\|_2 \quad \text{or "least squares".}$$

Suppose A has full column rank and A is $m \times n$. We start by ~~com~~ finding the point where $\nabla \|Ax-b\|_2$ vanishes.

Now

$$\|Ax-b\|_2 = (Ax-b)^T(Ax-b)$$

Thus we want $\nabla \|Ax-b\|_2 \cdot v = 0$ for all v.

$$\lim_{\epsilon \to 0} \frac{(A(x+\epsilon v)-b)^T(A(x+\epsilon v)-b) - (Ax-b)^T(Ax-b)}{\epsilon}$$

$$= \lim_{\epsilon \to 0} \frac{((Ax-b)+\epsilon Av)^T((Ax-b)+\epsilon Av) - (Ax-b)^T(Ax-b)}{\epsilon}$$

$$= \lim_{\epsilon \to 0} \frac{2\epsilon(Av)^T(Ax-b) + \epsilon^2(Av^T)Av}{\epsilon}$$

$$= 2(Av)^T(Ax-b)$$

$$= 2v^T(A^TAx - A^Tb) = 0.$$

Of course, this is true for all $\vec{v} \iff$

$$A^TAx = A^Tb$$

This $(n\times m)(m\times n) = n\times n$ system is called the <u>normal equations</u>. Note: this matrix (A^TA) is positive definite and hence nonsingular.

Why is this the global min? We complete the square: suppose x satisfies $A^T A x = A^T b$ and we write $x' = x + e$. Then

$$(Ax' - b)^T (Ax' - b) = (Ae + Ax - b)^T (Ae + Ax - b)$$

$$= Ae^T Ae + (Ax-b)^T (Ax-b) + 2(Ae)^T (Ax-b)$$

$$= \|Ae\|_2 + \|Ax - b\|_2 + \underline{2e(\cancel{A^T A x - A^T b})}$$

$$= \|Ae\|_2 + \|Ax - b\|_2$$

This is clearly minimized when $e = 0$.

What do we do to solve the normal equations? We first note that $m \geq n$ (~~[illegible]~~ overdetermined)

Now A^TA is spd, so we can use Cholesky decomposition $(\frac{1}{3}n^3 + O(n^2))$ flops. But just computing A^TA takes $\sim n^2m > n^3$ operations!

We also note that

$$\kappa(A^TA) = \kappa(A)^2$$

so we have ~~spa~~ potential stability problems for ill-conditioned matrices.

Now what do we do then? We first introduce a new matrix decomposition.

~~1~~

Theorem. If A is $m \times n$ with $m \geq n$ and A has full column rank, then $\exists$ a unique $m \times n$ orthogonal matrix Q and a unique $n \times n$ upper triangular R with positive diagonals so $A = QR$.

Proof. Consider the columns $A_1, \ldots, A_n$ of A. By assumption, they span an n dimensional subspace of $\mathbb{R}^m$.

Apply Gram-Schmidt. The resulting orthonormal vectors q_i are the cols of Q. Further

$$A_i = r_{ii} q_i + r_{i(i-1)} q_{i-1} + \cdots + r_{i1} q_1$$

by the Gram-Schmidt construction.

The r_{ij} are the entries in R. $\square$

How does this look as an algorithm?

for $i = 1$ to n

  $q_i = a_i$

  for $j = 1$ to $i-1$

    $r_{ji} = q_j^T a_i$ } Classical Gram-Schmidt

    $r_{ji} = q_j^T q_i$ } Modified Gram-Schmidt

    $q_i = q_i - r_{ji} q_j$

  end

  $r_{ii} = \|q_i\|_2$

  $q_i = q_i / r_{ii}$

end

the same?

homework.

The flop count here is

$$\sum_{i=1}^{n} i \cdot (3m) \approx \frac{3}{2} m n^2 .$$

Suppose we had such a decomposition. Then if x solves our problem,

$$
\begin{aligned}
A^T A x &= A^T b \\
x &= (A^T A)^{-1} A^T b \\
&= (R^T Q^T Q R)^{-1} R^T Q^T b \\
&= (R^T R)^{-1} R^T Q^T b \\
&= R^{-1} R^{-T} R^T Q^T b \\
&= R^{-1} Q^T b
\end{aligned}
$$

or

$Rx = Q^T b$, which we can solve by forward substitution.

Here is a third technique which will prove really useful.

Theorem. (SVD) Let A be any $m \times n$ matrix, with $m \geq n$. Then we can write

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

where U and V are orthogonal, U is $m \times n$, V is $n \times n$ and Σ is a diagonal matrix with entries

$$\sigma_{11} \geq \sigma_{22} \geq \dots \geq \sigma_{nn} \geq 0$$

The columns of U and V are called left and right <u>singular vectors</u> while the σ_i are <u>singular values</u>.

The claim here is simple and striking!

"Any matrix is diagonal in suitably chosen orthogonal coordinates on its range and domain"

Proof. We use induction on m and n, assuming the SVD exists for $(m-1)\times(n-1)$ matrices. We may assume $A \neq 0$.

Since $m \geq n$, the base case is $n=1$, m arbitrary (A is a column vector). We have

$$\underset{1\times m}{A} = \underset{1\times m}{U}\,\underset{1\times 1}{\Sigma}\,\underset{1\times 1}{V^T}$$

$$= \left(\frac{1}{\|A\|_2}A\right)\left[\|A\|_2\right]\left[1\right].$$

For the inductive step, choose v so $\|v\|_2 = 1$ and $\|A\|_2 = \|Av\|_2 > 0$. (This exists by definition of the matrix 2-norm.) Let $u = \frac{Av}{\|Av\|_2}$. Now complete u to an orthogonal basis of $\mathbb{R}^n$, forming an orthogonal matrix $U = [u, \tilde{U}]$.

We can do the same with $V = [v, \tilde{V}]$.

Now

$$U^T A V = \begin{bmatrix} u^T \\ \tilde{U}^T \end{bmatrix} A \begin{bmatrix} v & \tilde{V} \end{bmatrix}$$

$$= \begin{bmatrix} u^T A \\ \tilde{U}^T A \end{bmatrix} \begin{bmatrix} v & \tilde{V} \end{bmatrix}$$

$$= \begin{bmatrix} u^T A v & u^T A \tilde{V} \\ \tilde{U}^T A v & \tilde{U}^T A \tilde{V} \end{bmatrix}$$

We know

$$u^T A v = \left(\frac{Av}{\|Av\|_2} \right)^T Av = \frac{\|Av\|_2^2}{\|Av\|_2} = \|Av\|_2 = \|A\|_2$$

call this value σ. Now

$$\tilde{U}^T Av = \tilde{U}^T u \cdot \|Av\|_2 = 0$$

~~But~~ (u is orthogonal to ~~remaining~~ cols of $\tilde{U}$).

Now consider

$$u^T A \tilde{V}.$$

We claim $u^T A \tilde{V} = 0$. To see this, first observe

$$\|A\|_2 = \|U^T A V\|_2$$

since U and V are orthogonal matrices.

Now observe that $\|U^T A V\|_2 = \|(U^T A V)^T\|_2$.

Now

$$\|(U^T A V)^T e_1\|_2 = \|[1, \dots, 0]\ U^T A V\|_2$$

$$= \|[\sigma \quad u^T A \tilde{V}]\|_2 = \sqrt{\cancel{\sigma}\,\sigma^2 + \|u^T A \tilde{V}\|_2^2}$$

But this is equal to $\|A\|_2 = \sigma$ (tracing back through our chain of inequalities), so $\|u^T A \tilde{V}\|_2 = 0$ and $u^T A \tilde{V} = 0$.

We now know

$$U^T A V = \begin{bmatrix} \sigma & 0 \\ 0 & \tilde{U}^T A \tilde{V} \end{bmatrix} = \begin{bmatrix} \sigma & 0 \\ 0 & \tilde{A} \end{bmatrix}.$$

Applying inductive hypothesis to $\tilde{A}$, we can write $\tilde{A} = U_1 \Sigma_1 V_1^T$ where then

$$U^T A V = \begin{bmatrix} \sigma & 0 \\ 0 & U_1 \Sigma_1 V_1^T \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 \\ 0 & U_1 \end{bmatrix} \begin{bmatrix} \sigma & 0 \\ 0 & \Sigma_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & V_1 \end{bmatrix}^T$$

So

$$A = \left(U \begin{bmatrix} 1 & 0 \\ 0 & U_1 \end{bmatrix} \right) \left(\begin{bmatrix} \sigma & 0 \\ 0 & \Sigma_1 \end{bmatrix} \right) \left(\begin{bmatrix} 1 & 0 \\ 0 & V_1 \end{bmatrix}^T V^T \right)$$

which is the svd. $\square$

The SVD has a lot of useful properties which we can prove here.

Theorem. Let $A = U\Sigma V^T$ be the SVD of A where A is $m \times n$ with $m \geq n$.

1. If A is symmetric with eigenvalues λ_i and orthonormal eigenvectors u_i, then

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

where $\sigma_i = |\lambda_i|$ and $v_i = \text{sign}(\lambda_i) u_i$ is an SVD of A (here we need the convention $\text{sign}(0) = 1$).

2. The eigenvalues of A^TA are σ_i^2. The right singular vectors v_i are the corresponding eigenvectors.

3. The eigenvectors of AA^T $(m \times m)$ are the σ_i^2 and m-n zeros. The left singular vectors u_i are orthogonal eigenvectors for the σ_i.

4. If $H = \begin{bmatrix} 0 & A^T \\ A & 0 \end{bmatrix}$ where A is square and $U\Sigma V^T$ is the SVD of A. The eigenvalues of H are $\pm\sigma_i$ with unit eigenvectors $\frac{1}{\sqrt{2}}\begin{bmatrix} v_i \\ u_i \end{bmatrix}$.

5. If A has full rank, the solution to

$$\min_x \|Ax-b\|_2 \text{ is } x = V\Sigma^{-1}U^Tb.$$

6. $\|A\|_2 = \sigma_1$. If A is square and nonsingular, $\|A^{-1}\|_2 = 1/\sigma_n$ and

$$\|A\|_2 \|A^{-1}\|_2 = \sigma_1/\sigma_n.$$

7. Suppose some of the σ_i are 0, so $\sigma_1 \geq \ldots \geq \sigma_r > 0$, $\sigma_{r+1} = \ldots = \sigma_n = 0$. Then rank $A = r$ and

$$\operatorname{Ker} A = \operatorname{span}(v_{r+1}, \ldots, v_n)$$

while

$$\operatorname{Im} A = \operatorname{span}(u_1, \ldots, u_r).$$

8. Let S^{n-1} be the unit sphere in $\mathbb{R}^n$. Then the image $A S^{n-1}$ is an ellipsoid centered at the origin with axes $\sigma_i u_i$.

9. If $V = [v_1, \ldots, v_n]$ and $U = [u_1, \ldots, u_n]$ so $A = U \Sigma V^T = \sum_{i=1}^{n} \sigma_i u_i v_i^T$, where each $u_i v_i^T$ is a rank 1 matrix. Then a rank k matrix closest to A (in $\| \ \|_2$) is $A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^T$.

Further

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

We can also write A_k as $U\Sigma_k V^T$ where $\Sigma_k = \begin{bmatrix} \sigma_1 & & & & & \\ & \ddots & & & & \\ & & \sigma_k & & & \\ & & & 0 & & \\ & & & & \ddots & \\ & & & & & 0 \end{bmatrix}$.

Proof.

1. Suppose $A = U\Sigma V^T$. Then suppose

$$Ax = U\Sigma V^T x = \text{the SVD of } A,$$

Now if $x = v_i$, then since $V^T = V^{-1}$,

$$\begin{aligned} A v_i &= U\Sigma V^{-1} v_i \\ &= U\Sigma V^{-1} V e_i \\ &= U\Sigma e_i \\ &= U \sigma_i e_i = \sigma_i u_i. \end{aligned}$$

So if u_i is an eigenvector of A with eigen~~vector~~ value λ_i then if U, V are as in claim

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

$$= U \Sigma \begin{pmatrix} \operatorname{sign}\lambda_1 & & \\ & \ddots & \\ & & \operatorname{sign}\lambda_n \end{pmatrix} U^T u_i$$

$$= U \Sigma \begin{pmatrix} \operatorname{sign}\lambda_1 & & \\ & \ddots & \\ & & \operatorname{sign}\lambda_n \end{pmatrix} e_i$$

$$= U \lambda_i e_i = \lambda_i u_i = A u_i.$$

Thus $U \Sigma V^T = A$. Now U, V are orthogonal, Σ diagonal, positive, as desired.

2. We check

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

$$= V \Sigma^2 V^T$$

as before, v_i are eigenvectors w/evals σ_i^2.

3. Choose $\tilde{U}$ so $[U \ \tilde{U}]$ is ~~(nxn)~~ $(m \times m)$ and orthogonal. Then

$$\begin{aligned} AA^T &= U\Sigma V^T V \Sigma U^T \\ &= U \Sigma^2 U^T \\ &= \begin{bmatrix} U & \tilde{U} \end{bmatrix} \begin{bmatrix} \Sigma^2 & 0 \\ 0 & 0 \end{bmatrix} [U \ \tilde{U}]^T \end{aligned}$$

But as before, this means $U_i, \tilde{U}_i$ are the eigenvectors of AA^T.

4. Homework.

5. $\|Ax - b\|_2^2 = \| U\Sigma V^T x - b\|_2^2$.

   Now A has full rank, so Σ does as well, and Σ is invertible. So let $[U \ \tilde{U}]$ be square $(m \times m)$ and orthogonal as above.

Now

$$\| U \Sigma V^T x - b \|_2^2 = \left\| \begin{bmatrix} U^T \\ \tilde{U}^T \end{bmatrix} (U \Sigma V^T x - b) \right\|_2^2$$

$$= \left\| \begin{bmatrix} \Sigma V^T x - U^T b \\ - \tilde{U}^T b \end{bmatrix} \right\|_2^2$$

$$= \| \Sigma V^T x - U^T b \|_2^2 + \| \tilde{U}^T b \|_2^2$$

~~This is minimized~~

We can't change $\tilde{U}^T b$ (which is just the norm of projection of b to the subspace orthogonal to the column space of U).

But we can minimize this by choosing x so

$$\Sigma V^T x - U^T b = 0$$

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

7. Choose an $m \times (m-n)$ matrix $\tilde{U}$ so that $[U\ \tilde{U}]$ is square, orthogonal as before. Call this $\hat{U}$. Now $\hat{U}, V$ are nonsingular and so

$$\text{rank } A = \text{rank } \hat{U} A V$$

$$= \text{rank } [U\ \tilde{U}]\, U \Sigma V^T V$$

$$= \text{rank} \begin{bmatrix} \Sigma_1 \\ \tilde{U} U \Sigma_1 \end{bmatrix} = \text{rank} \begin{bmatrix} \Sigma_1 \\ 0 \end{bmatrix}$$

$$= \text{rank } \Sigma_1 = r.$$

Further, ~~Ker~~ if $\hat{\Sigma} = \begin{bmatrix} \Sigma_1 \\ 0 \end{bmatrix}$ then $\text{Ker } \hat{\Sigma}$ is clearly the subspace spanned by $e_{r+1}, \ldots, e_n$. These are the images of $v_{r+1}, \ldots, v_n$ under V^T, so that must be $\text{Ker } A$.

Now the image of A is then

$U \operatorname{span}(e_1, \ldots, e_r) = \operatorname{span}(u_1, \ldots, u_r)$.

8. We may as well write S^{n-1} as

$$\{ \textstyle\sum a_i v_i \mid \sum a_i^2 = 1 \}$$

since the matrix V is orthogonal. This maps by Σ to an ellipsoid of ~~the~~ axes of length σ_i. Multiplying by U rotates each axis $\sigma_i e_i$ to $\sigma_i u_i$, as claimed.

9. A_K certainly has rank K. Suppose B is another rank K matrix. Now

$$\operatorname{span}(v_1, \ldots, v_{K+1}) \cap \operatorname{Ker} B$$

has dimension at least 1, so let h

be a unit vector in intersection.

Now

$$\|A-B\|_2^2 \geq \|(A-B)h\|_2^2$$

$$= \|Ah\|_2^2$$

$$= \|U\Sigma V^T h\|_2^2$$

$$= \|\Sigma(V^T h)\|_2^2$$

We know $h \in \mathrm{span}(v_1, \dots, v_{k+1})$ so $V^T h$ is in $\mathrm{span}(e_1, \dots, e_{k+1})$. Each coordinate of $V^T h$ gets scaled by some σ_i with $i \in 1, \dots, k+1$, by ordering of σ_i, we have all these at least σ_{k+1}. So

$$\|\Sigma(V^T h)\|_2^2 \geq \sigma_{k+1}^2 \|V^T h\|_2^2 = \sigma_{k+1}^2.$$

The min is clearly acheived if $V^T h = e_{k+1}$.

We now check

$$\|A - A_k\|_2 = \left\| \sum_{i=k+1}^{n} \sigma_i u_i v_i^T \right\|_2$$

$$= \left\| U \begin{bmatrix} 0 & & & & \\ & \ddots & & & \\ & & 0 & & \\ & & & \sigma_{k+1} & \\ & & & & \ddots & \\ & & & & & \sigma_n \end{bmatrix} V^T \right\|_2$$

$$= \sigma_{k+1}. \qquad \square$$

We now make a definition.

Definition. Suppose A is $m \times n$ with $m \geq n$ and A has full rank. Further, suppose $A = QR = U\Sigma V^T$, as QR and SVD decompositions. We define the <u>Moore-Penrose pseudoinverse</u>

$$A^+ \equiv (A^T A)^{-1} A^T = V \Sigma^{-1} U^T$$

$$= R^{-1} Q^T.$$

If $m < n$, $A^+ \equiv A^T (AA^T)^{-1}$.

We can tie all our least squares methods together with the formula

$$x \text{ solves } \min_x \|Ax-b\|_2 \iff x = A^+ b.$$