
Lecture Notes to Accompany

Scientific Computing
An Introductory Survey
Second Edition

by Michael T. Heath

Chapter 3

Linear Least Squares

Copyright © 2001. Reproduction permitted only for noncommercial, educational use in conjunction with the book.

1

Linear Least Squares

For linear problems, obtain *overdetermined* linear system $\mathbf{Ax} = \mathbf{b}$, with $m \times n$ matrix \mathbf{A} , $m > n$

Better written $\mathbf{Ax} \cong \mathbf{b}$, since equality usually not exactly satisfiable when $m > n$

Least squares solution \mathbf{x} minimizes squared Euclidean norm of residual vector $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$,

$$\min_{\mathbf{x}} \|\mathbf{r}\|_2^2 = \min_{\mathbf{x}} \|\mathbf{b} - \mathbf{Ax}\|_2^2$$

3

Method of Least Squares

Measurement errors inevitable in observational and experimental sciences

Errors smoothed out by averaging over many cases, i.e., taking more measurements than strictly necessary to determine parameters of system

Resulting system *overdetermined*, so usually no exact solution

Project higher dimensional data into lower dimensional space to suppress irrelevant detail

Projection most conveniently accomplished by method of *least squares*

2

Example: Data Fitting

Given m data points (t_i, y_i) , find n -vector \mathbf{x} of parameters that gives “best fit” to model function $f(t, \mathbf{x})$:

$$\min_{\mathbf{x}} \sum_{i=1}^m (y_i - f(t_i, \mathbf{x}))^2$$

Problem *linear* if function f linear in components of \mathbf{x} :

$$f(t, \mathbf{x}) = x_1\phi_1(t) + x_2\phi_2(t) + \cdots + x_n\phi_n(t),$$

where functions ϕ_j depend only on t

Written in matrix form as $\mathbf{Ax} \cong \mathbf{b}$, with $a_{ij} = \phi_j(t_i)$ and $b_i = y_i$

4

Example: Data Fitting

Polynomial fitting,

$$f(t, \mathbf{x}) = x_1 + x_2 t + x_3 t^2 + \dots + x_n t^{n-1},$$

is linear, since polynomial linear in coefficients, though nonlinear in independent variable t

Fitting sum of exponentials, with

$$f(t, \mathbf{x}) = x_1 e^{x_2 t} + \dots + x_{n-1} e^{x_n t},$$

is nonlinear problem

For now, consider only linear least squares problems

5

Example: Data Fitting

Fitting quadratic polynomial to five data points gives linear least squares problem

$$A\mathbf{x} = \begin{bmatrix} 1 & t_1 & t_1^2 \\ 1 & t_2 & t_2^2 \\ 1 & t_3 & t_3^2 \\ 1 & t_4 & t_4^2 \\ 1 & t_5 & t_5^2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \cong \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix} = \mathbf{b}$$

Matrix with columns (or rows) successive powers of independent variable called *Vandermonde* matrix

6

Example: Data Fitting

For data

$$\begin{array}{c|ccccc} t & -1.0 & -0.5 & 0.0 & 0.5 & 1.0 \\ y & 1.0 & 0.5 & 0.0 & 0.5 & 2.0 \end{array}$$

overdetermined 5×3 linear system is

$$A\mathbf{x} = \begin{bmatrix} 1 & -1.0 & 1.0 \\ 1 & -0.5 & 0.25 \\ 1 & 0.0 & 0.0 \\ 1 & 0.5 & 0.25 \\ 1 & 1.0 & 1.0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \cong \begin{bmatrix} 1.0 \\ 0.5 \\ 0.0 \\ 0.5 \\ 2.0 \end{bmatrix} = \mathbf{b}$$

Solution, which we will see later how to compute, is

$$\mathbf{x} = [0.086 \quad 0.40 \quad 1.4]^T,$$


so approximating polynomial is

$$p(t) = 0.086 + 0.4t + 1.4t^2$$

7

Example: Data Fitting

Resulting curve and original data points shown in graph:



8

Normal Equations

Least squares minimizes squared Euclidean norm

$$\|r\|_2^2 = r^T r$$

of residual vector

$$r = b - Ax$$

To minimize

$$\begin{aligned} \|r\|_2^2 &= r^T r = (b - Ax)^T (b - Ax) \\ &= b^T b - 2x^T A^T b + x^T A^T A x, \end{aligned}$$

take derivative with respect to x and set to o ,

$$2A^T A x - 2A^T b = o,$$

which reduces to $n \times n$ linear system

$$A^T A x = A^T b$$

known as system of *normal equations*

Existence and Uniqueness

Linear least squares problem $Ax \cong b$ always has solution

Solution unique if, and only if, columns of A linearly independent, i.e., $\text{rank}(A) = n$, where A is $m \times n$

If $\text{rank}(A) < n$, then A is *rank-deficient*, and solution of linear least squares problem is not unique

For now, assume A has full column rank n

9

10

Orthogonality

Vectors v_1 and v_2 are *orthogonal* if their inner product is zero, $v_1^T v_2 = 0$

Space spanned by columns of $m \times n$ matrix A , $\text{span}(A) = \{Ax : x \in \mathbb{R}^n\}$, is of dimension at most n

If $m > n$, b generally does not lie in $\text{span}(A)$, so no exact solution to $Ax = b$

Vector $y = Ax$ in $\text{span}(A)$ closest to b in 2-norm occurs when residual $r = b - Ax$ orthogonal to $\text{span}(A)$

Thus,


$$o = A^T r = A^T (b - Ax),$$

or

$$A^T A x = A^T b$$

11

Orthogonality, continued



12

Orthogonal Projectors

Matrix P is *orthogonal projector* if *idempotent* ($P^2 = P$) and *symmetric* ($P^T = P$)

Orthogonal projector onto orthogonal complement $\text{span}(P)^\perp$ given by $P_\perp = I - P$

For any vector v ,

$$v = (P + (I - P))v = Pv + P_\perp v$$

For least squares problem $Ax \cong b$, if $\text{rank}(A) = n$, then

$$P = A(A^T A)^{-1} A^T$$

is orthogonal projector onto $\text{span}(A)$, and

$$b = Pb + P_\perp b = Ax + (b - Ax) = y + r$$

13

Sensitivity and Conditioning

Sensitivity of least squares solution to $Ax \cong b$ depends on b as well as A

Define angle θ between b and $y = Ax$ by

$$\cos(\theta) = \frac{\|y\|_2}{\|b\|_2} = \frac{\|Ax\|_2}{\|b\|_2}$$

(see previous drawing)

Bound on perturbation Δx in solution x due to perturbation Δb in b given by

$$\frac{\|\Delta x\|_2}{\|x\|_2} \leq \text{cond}(A) \frac{1}{\cos(\theta)} \frac{\|\Delta b\|_2}{\|b\|_2}$$

15

Pseudoinverse and Condition Number

Nonsquare $m \times n$ matrix A has no inverse in usual sense

If $\text{rank}(A) = n$, *pseudoinverse* defined by

$$A^+ = (A^T A)^{-1} A^T$$

and

$$\text{cond}(A) = \|A\|_2 \cdot \|A^+\|_2$$

By convention, $\text{cond}(A) = \infty$ if $\text{rank}(A) < n$

Just as condition number of square matrix measures closeness to singularity, condition number of rectangular matrix measures closeness to rank deficiency

Least squares solution of $Ax \cong b$ is given by $x = A^+ b$

14

Sensitivity and Conditioning, cont.

Similarly, for perturbation E in matrix A ,

$$\frac{\|\Delta x\|_2}{\|x\|_2} \lesssim ([\text{cond}(A)]^2 \tan(\theta) + \text{cond}(A)) \frac{\|E\|_2}{\|A\|_2}$$

Condition number of least squares solution about $\text{cond}(A)$ if residual small, but can be squared or arbitrarily worse for large residual

16

Normal Equations Method

If $m \times n$ matrix A has rank n , then symmetric $n \times n$ matrix $A^T A$ is positive definite, so its Cholesky factorization

$$A^T A = LL^T,$$

can be used to obtain solution x to system of normal equations

$$A^T A x = A^T b,$$

which has same solution as linear least squares problem $Ax \cong b$

Normal equations method involves transformations

rectangular \rightarrow square \rightarrow triangular

17

Example: Normal Equations Method

For polynomial data-fitting example given previously, normal equations method gives

$$A^T A = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ -1.0 & -0.5 & 0.0 & 0.5 & 1.0 \\ 1.0 & 0.25 & 0.0 & 0.25 & 1.0 \end{bmatrix}$$

$$\begin{bmatrix} 1 & -1.0 & 1.0 \\ 1 & -0.5 & 0.25 \\ 1 & 0.0 & 0.0 \\ 1 & 0.5 & 0.25 \\ 1 & 1.0 & 1.0 \end{bmatrix} = \begin{bmatrix} 5.0 & 0.0 & 2.5 \\ 0.0 & 2.5 & 0.0 \\ 2.5 & 0.0 & 2.125 \end{bmatrix},$$

$$A^T b =$$

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ -1.0 & -0.5 & 0.0 & 0.5 & 1.0 \\ 1.0 & 0.25 & 0.0 & 0.25 & 1.0 \end{bmatrix} \begin{bmatrix} 1.0 \\ 0.5 \\ 0.0 \\ 0.5 \\ 2.0 \end{bmatrix} = \begin{bmatrix} 4.0 \\ 1.0 \\ 3.25 \end{bmatrix}$$

18

Example Continued

Cholesky factorization of symmetric positive definite matrix gives

$$A^T A = \begin{bmatrix} 5.0 & 0.0 & 2.5 \\ 0.0 & 2.5 & 0.0 \\ 2.5 & 0.0 & 2.125 \end{bmatrix} =$$

$$\begin{bmatrix} 2.236 & 0 & 0 \\ 0 & 1.581 & 0 \\ 1.118 & 0 & 0.935 \end{bmatrix} \begin{bmatrix} 2.236 & 0 & 1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & 0.935 \end{bmatrix} = LL^T$$

19

Example Continued

Solving lower triangular system $Lz = A^T b$ by forward-substitution gives

$$z = \begin{bmatrix} 1.789 \\ 0.632 \\ 1.336 \end{bmatrix}$$

Solving upper triangular system $L^T x = z$ by back-substitution gives least squares solution

$$x = \begin{bmatrix} 0.086 \\ 0.400 \\ 1.429 \end{bmatrix}$$

20

Shortcomings of Normal Equations

Information can be lost in forming $A^T A$ and $A^T b$

For example, take

$$A = \begin{bmatrix} 1 & 1 \\ \epsilon & 0 \\ 0 & \epsilon \end{bmatrix},$$

where ϵ is positive number smaller than $\sqrt{\epsilon_{\text{mach}}}$

Then in floating-point arithmetic

$$A^T A = \begin{bmatrix} 1 + \epsilon^2 & 1 \\ 1 & 1 + \epsilon^2 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix},$$

which is singular

Sensitivity of solution also worsened, since

$$\text{cond}(A^T A) = [\text{cond}(A)]^2$$

21

Augmented System Method, continued

Introducing scaling parameter α gives system

$$\begin{bmatrix} \alpha I & A \\ A^T & O \end{bmatrix} \begin{bmatrix} r/\alpha \\ x \end{bmatrix} = \begin{bmatrix} b \\ o \end{bmatrix},$$

which allows control over relative weights of two subsystems in choosing pivots

Reasonable rule of thumb

$$\alpha = \max_{i,j} |a_{ij}| / 1000$$

Augmented system sometimes useful, but far from ideal in work and storage required

23

Augmented System Method

Definition of residual and orthogonality requirement give $(m+n) \times (m+n)$ augmented system

$$\begin{bmatrix} I & A \\ A^T & O \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ o \end{bmatrix}$$

System not positive definite, larger than original, and requires storing two copies of A

But allows greater freedom in choosing pivots in computing LDL^T or LU factorization

22

Orthogonal Transformations

Seek alternative method that avoids numerical difficulties of normal equations

Need numerically robust transformation that produces easier problem

What kind of transformation leaves least squares solution unchanged?

Square matrix Q is *orthogonal* if $Q^T Q = I$

Preserves Euclidean norm, since

$$\|Qv\|_2^2 = (Qv)^T Qv = v^T Q^T Qv = v^T v = \|v\|_2^2$$

Multiplying both sides of least squares problem by orthogonal matrix does not change solution

24

Triangular Least Squares Problems

As with square linear systems, suitable target in simplifying least squares problems is triangular form

Upper triangular overdetermined ($m > n$) least squares problem has form

$$\begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x} \cong \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix},$$

with \mathbf{R} $n \times n$ upper triangular and \mathbf{b} partitioned similarly

Residual is

$$\|\mathbf{r}\|_2^2 = \|\mathbf{b}_1 - \mathbf{R}\mathbf{x}\|_2^2 + \|\mathbf{b}_2\|_2^2$$

25

QR Factorization

Given $m \times n$ matrix \mathbf{A} , with $m > n$, we seek $m \times m$ orthogonal matrix \mathbf{Q} such that

$$\mathbf{A} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix},$$

with \mathbf{R} $n \times n$ and upper triangular

Linear least squares problem $\mathbf{A}\mathbf{x} \cong \mathbf{b}$ transformed into triangular least squares problem

$$\mathbf{Q}^T \mathbf{A}\mathbf{x} = \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x} \cong \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{bmatrix} = \mathbf{Q}^T \mathbf{b},$$

which has same solution, since $\|\mathbf{r}\|_2^2 =$

$$\begin{aligned} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 &= \|\mathbf{b} - \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x}\|_2^2 = \|\mathbf{Q}^T \mathbf{b} - \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x}\|_2^2 \\ &= \|\mathbf{c}_1 - \mathbf{R}\mathbf{x}\|_2^2 + \|\mathbf{c}_2\|_2^2 \end{aligned}$$

because orthogonal transformation preserves Euclidean norm

27

Triangular Least Squares Problems, cont.

Have no control over second term, $\|\mathbf{b}_2\|_2^2$, but first term becomes zero if \mathbf{x} satisfies triangular system

$$\mathbf{R}\mathbf{x} = \mathbf{b}_1,$$

which can be solved by back-substitution

Resulting \mathbf{x} is least squares solution, and minimum sum of squares is

$$\|\mathbf{r}\|_2^2 = \|\mathbf{b}_2\|_2^2$$

So strategy is to transform general least squares problem to triangular form using orthogonal transformation

26

Orthogonal Bases

Partition $m \times m$ orthogonal matrix $\mathbf{Q} = [\mathbf{Q}_1 \ \mathbf{Q}_2]$, with \mathbf{Q}_1 $m \times n$

Then

$$\mathbf{A} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} = [\mathbf{Q}_1 \ \mathbf{Q}_2] \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} = \mathbf{Q}_1 \mathbf{R}$$

is *reduced* QR factorization of \mathbf{A}

Columns of \mathbf{Q}_1 are *orthonormal basis* for $\text{span}(\mathbf{A})$, and columns of \mathbf{Q}_2 are orthonormal basis for $\text{span}(\mathbf{A})^\perp$

$\mathbf{Q}_1 \mathbf{Q}_1^T$ is orthogonal projector onto $\text{span}(\mathbf{A})$

Solution to least squares problem $\mathbf{A}\mathbf{x} \cong \mathbf{b}$ given by solution to square system

$$\mathbf{Q}_1^T \mathbf{A}\mathbf{x} = \mathbf{R}\mathbf{x} = \mathbf{c}_1 = \mathbf{Q}_1^T \mathbf{b}$$

28

QR Factorization

To compute QR factorization of $m \times n$ matrix \mathbf{A} , with $m > n$, annihilate subdiagonal entries of successive columns of \mathbf{A} , eventually reaching upper triangular form

Similar to LU factorization by Gaussian elimination, but uses orthogonal transformations instead of elementary elimination matrices

Possible methods include

- Householder transformations
- Givens rotations
- Gram-Schmidt orthogonalization

29

Example: Householder Transformation

Let $\mathbf{a} = [2 \ 1 \ 2]^T$

By foregoing recipe,

$$\mathbf{v} = \mathbf{a} - \alpha \mathbf{e}_1 = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} - \alpha \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} - \begin{bmatrix} \alpha \\ 0 \\ 0 \end{bmatrix},$$

where $\alpha = \pm \|\mathbf{a}\|_2 = \pm 3$

Since a_1 positive, choosing negative sign for α avoids cancellation

$$\text{Thus, } \mathbf{v} = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} - \begin{bmatrix} -3 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 5 \\ 1 \\ 2 \end{bmatrix}$$

To confirm that transformation works,

$$\mathbf{H}\mathbf{a} = \mathbf{a} - 2 \frac{\mathbf{v}^T \mathbf{a}}{\mathbf{v}^T \mathbf{v}} \mathbf{v} = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} - 2 \frac{15}{30} \begin{bmatrix} 5 \\ 1 \\ 2 \end{bmatrix} = \begin{bmatrix} -3 \\ 0 \\ 0 \end{bmatrix}$$

31

Householder Transformations

Householder transformation has form

$$\mathbf{H} = \mathbf{I} - 2 \frac{\mathbf{v}\mathbf{v}^T}{\mathbf{v}^T \mathbf{v}}$$

for nonzero vector \mathbf{v}

$\mathbf{H} = \mathbf{H}^T = \mathbf{H}^{-1}$, so \mathbf{H} is orthogonal and symmetric

Given vector \mathbf{a} , choose \mathbf{v} so that

$$\mathbf{H}\mathbf{a} = \begin{bmatrix} \alpha \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \alpha \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \alpha \mathbf{e}_1$$

Substituting into formula for \mathbf{H} , can take

$$\mathbf{v} = \mathbf{a} - \alpha \mathbf{e}_1$$

and $\alpha = \pm \|\mathbf{a}\|_2$, with sign chosen to avoid cancellation

30

Householder QR Factorization

To compute QR factorization of \mathbf{A} , use Householder transformations to annihilate subdiagonal entries of each successive column

Each Householder transformation applied to entire matrix, but does not affect prior columns, so zeros preserved

In applying Householder transformation \mathbf{H} to arbitrary vector \mathbf{u} ,

$$\mathbf{H}\mathbf{u} = \left(\mathbf{I} - 2 \frac{\mathbf{v}\mathbf{v}^T}{\mathbf{v}^T \mathbf{v}} \right) \mathbf{u} = \mathbf{u} - \left(2 \frac{\mathbf{v}^T \mathbf{u}}{\mathbf{v}^T \mathbf{v}} \right) \mathbf{v},$$

which is much cheaper than general matrix-vector multiplication and requires only vector \mathbf{v} , not full matrix \mathbf{H}

32

Householder QR Factorization, cont.

Process just described produces factorization

$$H_n \cdots H_1 A = \begin{bmatrix} R \\ O \end{bmatrix}$$

with R $n \times n$ and upper triangular

If $Q = H_1 \cdots H_n$, then

$$A = Q \begin{bmatrix} R \\ O \end{bmatrix}$$

To preserve solution of linear least squares problem, right-hand-side b transformed by same sequence of Householder transformations

Then solve triangular least squares problem

$$\begin{bmatrix} R \\ O \end{bmatrix} x \cong Q^T b$$

for solution x of original least squares problem

33

Householder QR Factorization, cont.

For solving linear least squares problem, product Q of Householder transformations need not be formed explicitly

R can be stored in upper triangle of array initially containing A

Householder vectors v can be stored in (now zero) lower triangular portion of A (almost)

Householder transformations most easily applied in this form anyway

34

Example: Householder QR Factorization

For polynomial data-fitting example given previously, with

$$A = \begin{bmatrix} 1 & -1.0 & 1.0 \\ 1 & -0.5 & 0.25 \\ 1 & 0.0 & 0.0 \\ 1 & 0.5 & 0.25 \\ 1 & 1.0 & 1.0 \end{bmatrix}, \quad b = \begin{bmatrix} 1.0 \\ 0.5 \\ 0.0 \\ 0.5 \\ 2.0 \end{bmatrix},$$

Householder vector v_1 for annihilating sub-diagonal entries of first column of A is

$$v_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} -2.236 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 3.236 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

35

Example Continued

Applying resulting Householder transformation H_1 yields transformed matrix and right-hand side

$$H_1 A = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & -0.191 & -0.405 \\ 0 & 0.309 & -0.655 \\ 0 & 0.809 & -0.405 \\ 0 & 1.309 & 0.345 \end{bmatrix}$$

$$H_1 b = \begin{bmatrix} -1.789 \\ -0.362 \\ -0.862 \\ -0.362 \\ 1.138 \end{bmatrix}$$

36

Example Continued

Householder vector v_2 for annihilating sub-diagonal entries of second column of H_1A is

$$v_2 = \begin{bmatrix} 0 \\ -0.191 \\ 0.309 \\ 0.809 \\ 1.309 \end{bmatrix} - \begin{bmatrix} 0 \\ 1.581 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ -1.772 \\ 0.309 \\ 0.809 \\ 1.309 \end{bmatrix}$$

Applying resulting Householder transformation H_2 yields

$$H_2H_1A = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & -0.725 \\ 0 & 0 & -0.589 \\ 0 & 0 & 0.047 \end{bmatrix}$$

37

Example Continued

$$H_2H_1b = \begin{bmatrix} -1.789 \\ 0.632 \\ -1.035 \\ -0.816 \\ 0.404 \end{bmatrix}$$

Householder vector v_3 for annihilating sub-diagonal entries of third column of H_2H_1A is

$$v_3 = \begin{bmatrix} 0 \\ 0 \\ -0.725 \\ -0.589 \\ 0.047 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 0.935 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -1.660 \\ -0.589 \\ 0.047 \end{bmatrix}$$

38

Example Continued

Applying resulting Householder transformation H_3 yields

$$H_3H_2H_1A = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & 0.935 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$H_3H_2H_1b = \begin{bmatrix} -1.789 \\ 0.632 \\ 1.336 \\ 0.026 \\ 0.337 \end{bmatrix}$$

Now solve upper triangular system $Rx = c_1$ by back-substitution to obtain

$$x = [0.086 \quad 0.400 \quad 1.429]^T$$

39

Givens Rotations

Givens rotations introduce zeros one at a time

Given vector $[a_1 \ a_2]^T$, choose scalars c and s so that

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \begin{bmatrix} \alpha \\ 0 \end{bmatrix},$$

with $c^2 + s^2 = 1$, or equivalently, $\alpha = \sqrt{a_1^2 + a_2^2}$

Previous equation can be rewritten

$$\begin{bmatrix} a_1 & a_2 \\ a_2 & -a_1 \end{bmatrix} \begin{bmatrix} c \\ s \end{bmatrix} = \begin{bmatrix} \alpha \\ 0 \end{bmatrix}$$

Gaussian elimination yields triangular system

$$\begin{bmatrix} a_1 & a_2 \\ 0 & -a_1 - a_2^2/a_1 \end{bmatrix} \begin{bmatrix} c \\ s \end{bmatrix} = \begin{bmatrix} \alpha \\ -\alpha a_2/a_1 \end{bmatrix}$$

40

Example: Givens Rotation

Let $a = [4 \ 3]^T$

Computing cosine and sine,

$$c = \frac{a_1}{\sqrt{a_1^2 + a_2^2}} = \frac{4}{5} = 0.8$$

$$s = \frac{a_2}{\sqrt{a_1^2 + a_2^2}} = \frac{3}{5} = 0.6$$

Rotation given by

$$G = \begin{bmatrix} c & s \\ -s & c \end{bmatrix} = \begin{bmatrix} 0.8 & 0.6 \\ -0.6 & 0.8 \end{bmatrix}$$

To confirm that rotation works,

$$Ga = \begin{bmatrix} 0.8 & 0.6 \\ -0.6 & 0.8 \end{bmatrix} \begin{bmatrix} 4 \\ 3 \end{bmatrix} = \begin{bmatrix} 5 \\ 0 \end{bmatrix}$$

Givens Rotations, continued

Back-substitution then gives

$$s = \frac{\alpha a_2}{a_1^2 + a_2^2}, \quad c = \frac{\alpha a_1}{a_1^2 + a_2^2}$$

Finally, $c^2 + s^2 = 1$, or $\alpha = \sqrt{a_1^2 + a_2^2}$, implies

$$c = \frac{a_1}{\sqrt{a_1^2 + a_2^2}}, \quad s = \frac{a_2}{\sqrt{a_1^2 + a_2^2}}$$

41

42

Givens QR Factorization

To annihilate selected component of vector in n dimensions, rotate target component with another component

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & c & 0 & s & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & -s & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \end{bmatrix} = \begin{bmatrix} a_1 \\ \alpha \\ a_3 \\ 0 \\ a_5 \end{bmatrix}$$

Reduce matrix to upper triangular form using sequence of Givens rotations


Each rotation orthogonal, so their product is orthogonal, producing QR factorization

Straightforward implementation of Givens method requires about 50% more work than Householder method, and also requires more storage, since each rotation requires two numbers, c and s , to define it

43

Gram-Schmidt Orthogonalization

Given vectors a_1 and a_2 , can determine orthonormal vectors q_1 and q_2 with same span by orthogonalizing one vector against other:



```

for  $k = 1$  to  $n$ 
     $q_k = a_k$ 
    for  $j = 1$  to  $k - 1$ 
         $r_{jk} = q_j^T a_k$ 
         $q_k = q_k - r_{jk} q_j$ 
    end
     $r_{kk} = \|q_k\|_2$ 
     $q_k = q_k / r_{kk}$ 
end
    
```

44

Modified Gram-Schmidt

Classical Gram-Schmidt procedure often suffers loss of orthogonality in finite-precision

Also, separate storage is required for \mathbf{A} , \mathbf{Q} , and \mathbf{R} , since original \mathbf{a}_k needed in inner loop, so \mathbf{q}_k cannot overwrite columns of \mathbf{A}

Both deficiencies improved by *modified* Gram-Schmidt procedure, with each vector orthogonalized in turn against all *subsequent* vectors so \mathbf{q}_k can overwrite \mathbf{a}_k :

```
for k = 1 to n
  rkk = ||ak||2
  qk = ak/rkk
  for j = k + 1 to n
    rkj = qkTaj
    aj = aj - rkjqk
  end
end
end
```

45

Rank Deficiency

If $\text{rank}(\mathbf{A}) < n$, then QR factorization still exists, but yields singular upper triangular factor \mathbf{R} , and multiple vectors \mathbf{x} give minimum residual norm

Common practice selects minimum residual solution \mathbf{x} having smallest norm

Can be computed by QR factorization with column pivoting or by singular value decomposition (SVD)

Rank of matrix often not clear cut in practice, so relative tolerance used to determine rank

46

Example: Near Rank Deficiency

Consider 3×2 matrix

$$\mathbf{A} = \begin{bmatrix} 0.641 & 0.242 \\ 0.321 & 0.121 \\ 0.962 & 0.363 \end{bmatrix}$$

Computing QR factorization,

$$\mathbf{R} = \begin{bmatrix} 1.1997 & 0.4527 \\ 0 & 0.0002 \end{bmatrix}$$

\mathbf{R} extremely close to singular (exactly singular to 3-digit accuracy of problem statement)

If \mathbf{R} used to solve linear least squares problem, result is highly sensitive to perturbations in right-hand side

For practical purposes, $\text{rank}(\mathbf{A}) = 1$ rather than 2, because columns nearly linearly dependent

47

QR with Column Pivoting

Instead of processing columns in natural order, select for reduction at each stage column of remaining unreduced submatrix having maximum Euclidean norm

If $\text{rank}(\mathbf{A}) = k < n$, then after k steps, norms of remaining unreduced columns will be zero (or "negligible" in finite-precision arithmetic) below row k

Yields orthogonal factorization of form

$$\mathbf{Q}^T \mathbf{A} \mathbf{P} = \begin{bmatrix} \mathbf{R} & \mathbf{S} \\ \mathbf{O} & \mathbf{O} \end{bmatrix},$$

with \mathbf{R} $k \times k$, upper triangular, and nonsingular, and permutation matrix \mathbf{P} performing column interchanges

48

QR with Column Pivoting, cont.

Basic solution to least squares problem $\mathbf{Ax} \cong \mathbf{b}$ can now be computed by solving triangular system $\mathbf{Rz} = \mathbf{c}_1$, where \mathbf{c}_1 contains first k components of $\mathbf{Q}^T \mathbf{b}$, and then taking

$$\mathbf{x} = \mathbf{P} \begin{bmatrix} \mathbf{z} \\ \mathbf{o} \end{bmatrix}$$

Minimum-norm solution can be computed, if desired, at expense of additional processing to annihilate \mathbf{S}

$\text{rank}(\mathbf{A})$ usually unknown, so rank determined by monitoring norms of remaining unreduced columns and terminating factorization when maximum value falls below tolerance

49

Example: SVD

SVD of

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \\ 10 & 11 & 12 \end{bmatrix}$$

given by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T =$

$$\begin{bmatrix} .141 & .825 & -.420 & -.351 \\ .344 & .426 & .298 & .782 \\ .547 & .0278 & .664 & -.509 \\ .750 & -.371 & -.542 & .0790 \end{bmatrix}$$

$$\begin{bmatrix} 25.5 & 0 & 0 \\ 0 & 1.29 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} .504 & .574 & .644 \\ -.761 & -.057 & .646 \\ .408 & -.816 & .408 \end{bmatrix}$$

51

Singular Value Decomposition

Singular value decomposition (SVD) of $m \times n$ matrix \mathbf{A} has form

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T,$$

where \mathbf{U} is $m \times m$ orthogonal matrix, \mathbf{V} is $n \times n$ orthogonal matrix, and $\mathbf{\Sigma}$ is $m \times n$ diagonal matrix, with

$$\sigma_{ij} = \begin{cases} 0 & \text{for } i \neq j \\ \sigma_i \geq 0 & \text{for } i = j \end{cases}$$

Diagonal entries σ_i , called *singular values* of \mathbf{A} , usually ordered so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$

Columns \mathbf{u}_i of \mathbf{U} and \mathbf{v}_i of \mathbf{V} called left and right *singular vectors*

50

Applications of SVD

Minimum norm solution to $\mathbf{Ax} \cong \mathbf{b}$:

$$\mathbf{x} = \sum_{\sigma_i \neq 0} \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i$$

For ill-conditioned or rank deficient problems, "small" singular values can be dropped from summation to stabilize solution

Euclidean matrix norm:

$$\|\mathbf{A}\|_2 = \max_{\mathbf{x} \neq \mathbf{o}} \frac{\|\mathbf{Ax}\|_2}{\|\mathbf{x}\|_2} = \sigma_{\max}$$

Euclidean condition number of matrix:

$$\text{cond}(\mathbf{A}) = \sigma_{\max} / \sigma_{\min}$$

Rank of matrix: number of nonzero, or non-negligible, singular values

52

Pseudoinverse

Define pseudoinverse of scalar σ to be $1/\sigma$ if $\sigma \neq 0$, zero otherwise

Define pseudoinverse of (possibly rectangular) diagonal matrix by transposing and taking scalar pseudoinverse of each entry

Then pseudoinverse of general real $m \times n$ matrix \mathbf{A} given by

$$\mathbf{A}^+ = \mathbf{V}\Sigma^+\mathbf{U}^T$$

Pseudoinverse always exists whether or not matrix is square or has full rank

If \mathbf{A} is square and nonsingular, then $\mathbf{A}^+ = \mathbf{A}^{-1}$

In all cases, minimum-norm solution to $\mathbf{A}\mathbf{x} \cong \mathbf{b}$ is given by $\mathbf{A}^+\mathbf{b}$

53

Lower-Rank Matrix Approximation

Another way to write SVD:

$$\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^T = \sigma_1\mathbf{E}_1 + \sigma_2\mathbf{E}_2 + \cdots + \sigma_n\mathbf{E}_n,$$

with $\mathbf{E}_i = \mathbf{u}_i\mathbf{v}_i^T$

\mathbf{E}_i has rank 1 and can be stored using only $m + n$ storage locations

Product $\mathbf{E}_i\mathbf{x}$ can be formed using only $m + n$ multiplications

Condensed approximation to \mathbf{A} obtained by omitting from summation terms corresponding to small singular values

Approximation using k largest singular values is closest matrix of rank k to \mathbf{A}

Approximation is useful in image processing, data compression, information retrieval, cryptography, etc.

55

Orthogonal Bases

Columns of \mathbf{U} corresponding to nonzero singular values form orthonormal basis for $\text{span}(\mathbf{A})$

Remaining columns of \mathbf{U} form orthonormal basis for orthogonal complement $\text{span}(\mathbf{A})^\perp$

Columns of \mathbf{V} corresponding to zero singular values form orthonormal basis for null space of \mathbf{A}

Remaining columns of \mathbf{V} form orthonormal basis for orthogonal complement of null space

54

Total Least Squares

Ordinary least squares applicable when right hand side \mathbf{b} subject to random error but matrix \mathbf{A} known accurately

When all data, including \mathbf{A} , subject to error, then total least squares more appropriate

Total least squares minimizes orthogonal distances, rather than vertical distances, between model and data

Total least squares solution can be computed from SVD of $[\mathbf{A}, \mathbf{b}]$

56

Comparison of Methods

Forming normal equations matrix $\mathbf{A}^T \mathbf{A}$ requires about $n^2 m / 2$ multiplications, and solving resulting symmetric linear system requires about $n^3 / 6$ multiplications

Solving least squares problem using Householder QR factorization requires about $mn^2 - n^3 / 3$ multiplications

If $m \approx n$, two methods require about same amount of work

If $m \gg n$, Householder QR requires about twice as much work as normal equations

Cost of SVD proportional to $mn^2 + n^3$, with proportionality constant ranging from 4 to 10, depending on algorithm used

57

Comparison of Methods, continued

Householder is more accurate and more broadly applicable than normal equations

These advantages may not be worth additional cost, however, when problem is sufficiently well conditioned that normal equations provide adequate accuracy

For rank-deficient or nearly rank-deficient problem, Householder with column pivoting can produce useful solution when normal equations method fails outright

SVD is even more robust and reliable than Householder, but substantially more expensive

59

Comparison of Methods, continued

Normal equations method produces solution whose relative error is proportional to $[\text{cond}(\mathbf{A})]^2$

Required Cholesky factorization can be expected to break down if $\text{cond}(\mathbf{A}) \approx 1/\sqrt{\epsilon_{\text{mach}}}$ or worse

Householder method produces solution whose relative error is proportional to

$$\text{cond}(\mathbf{A}) + \|\mathbf{r}\|_2 [\text{cond}(\mathbf{A})]^2,$$

which is best possible, since this is inherent sensitivity of solution to least squares problem

Householder method can be expected to break down (in back-substitution phase) only if

$$\text{cond}(\mathbf{A}) \approx 1/\epsilon_{\text{mach}}$$

or worse

58