

CS 450 – Numerical Analysis

Chapter 3: Linear Least Squares †

Prof. Michael T. Heath

Department of Computer Science
University of Illinois at Urbana-Champaign
heath@illinois.edu

January 28, 2019

†Lecture slides based on the textbook *Scientific Computing: An Introductory Survey* by Michael T. Heath, copyright © 2018 by the Society for Industrial and Applied Mathematics. <http://www.siam.org/books/cl80>

Linear Least Squares

Method of Least Squares

- ▶ Measurement errors and other sources of random variation are inevitable in observational and experimental sciences
- ▶ Such variability can be smoothed out by averaging over many cases, e.g., taking more measurements than are strictly necessary to determine parameters of system
- ▶ Resulting system is *overdetermined*, so usually there is no exact solution
- ▶ In effect, higher dimensional data are projected onto lower dimensional space to suppress noise or irrelevant detail
- ▶ Such projection is most conveniently accomplished by method of *least squares*

Linear Least Squares

- ▶ For linear problems, we obtain *overdetermined* linear system $\mathbf{Ax} = \mathbf{b}$, with $m \times n$ matrix \mathbf{A} , $m > n$
- ▶ System is better written $\mathbf{Ax} \cong \mathbf{b}$, since equality is usually not exactly satisfiable when $m > n$
- ▶ *Least squares* solution \mathbf{x} minimizes squared Euclidean norm of residual vector $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$,

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

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 is *linear* if function f is 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

- ▶ Linear problem can be written in matrix form as $\mathbf{Ax} \cong \mathbf{b}$, with $a_{ij} = \phi_j(t_i)$ and $b_i = y_i$

Data Fitting

- ▶ Polynomial fitting

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

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

- ▶ Fitting sum of exponentials

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

is example of nonlinear problem

- ▶ For now, we will consider only linear least squares problems

Example: Data Fitting

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

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

- ▶ Matrix whose columns (or rows) are successive powers of independent variable is called *Vandermonde matrix*

Example, continued

- ▶ 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

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

- ▶ 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

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

Example, continued

- ▶ Resulting curve and original data points are shown in graph



[⟨ interactive example ⟩](#)

Existence, Uniqueness, and Conditioning

Existence and Uniqueness

- ▶ Linear least squares problem $\mathbf{Ax} \cong \mathbf{b}$ *always* has solution
- ▶ Solution is *unique* if, and only if, columns of \mathbf{A} are *linearly independent*, i.e., $\text{rank}(\mathbf{A}) = n$, where \mathbf{A} is $m \times n$
- ▶ If $\text{rank}(\mathbf{A}) < n$, then \mathbf{A} is *rank-deficient*, and solution of linear least squares problem is not unique
- ▶ For now, we assume \mathbf{A} has full column rank n

Normal Equations

- ▶ To minimize squared Euclidean norm of residual vector

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

take derivative with respect to \mathbf{x} and set it to $\mathbf{0}$,

$$2\mathbf{A}^T \mathbf{A} \mathbf{x} - 2\mathbf{A}^T \mathbf{b} = \mathbf{0}$$

which reduces to $n \times n$ linear system of *normal equations*

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

Orthogonality

- ▶ Vectors \mathbf{v}_1 and \mathbf{v}_2 are *orthogonal* if their inner product is zero, $\mathbf{v}_1^T \mathbf{v}_2 = 0$
- ▶ Space spanned by columns of $m \times n$ matrix \mathbf{A} , $\text{span}(\mathbf{A}) = \{\mathbf{Ax} : \mathbf{x} \in \mathbb{R}^n\}$, is of dimension at most n
- ▶ If $m > n$, \mathbf{b} generally does not lie in $\text{span}(\mathbf{A})$, so there is no exact solution to $\mathbf{Ax} = \mathbf{b}$
- ▶ Vector $\mathbf{y} = \mathbf{Ax}$ in $\text{span}(\mathbf{A})$ closest to \mathbf{b} in 2-norm occurs when residual $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$ is *orthogonal* to $\text{span}(\mathbf{A})$,


$$\mathbf{0} = \mathbf{A}^T \mathbf{r} = \mathbf{A}^T (\mathbf{b} - \mathbf{Ax})$$

again giving system of *normal equations*

$$\mathbf{A}^T \mathbf{Ax} = \mathbf{A}^T \mathbf{b}$$

Orthogonality, continued

- ▶ Geometric relationships among \mathbf{b} , \mathbf{r} , and $\text{span}(\mathbf{A})$ are shown in diagram



Orthogonal Projectors

- ▶ Matrix \mathbf{P} is *orthogonal projector* if it is *idempotent* ($\mathbf{P}^2 = \mathbf{P}$) and *symmetric* ($\mathbf{P}^T = \mathbf{P}$)
- ▶ Orthogonal projector onto orthogonal complement $\text{span}(\mathbf{P})^\perp$ is given by $\mathbf{P}_\perp = \mathbf{I} - \mathbf{P}$
- ▶ For any vector \mathbf{v} ,

$$\mathbf{v} = (\mathbf{P} + (\mathbf{I} - \mathbf{P})) \mathbf{v} = \mathbf{P}\mathbf{v} + \mathbf{P}_\perp \mathbf{v}$$

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

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

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

$$\mathbf{b} = \mathbf{P}\mathbf{b} + \mathbf{P}_\perp \mathbf{b} = \mathbf{A}\mathbf{x} + (\mathbf{b} - \mathbf{A}\mathbf{x}) = \mathbf{y} + \mathbf{r}$$

Pseudoinverse and Condition Number

- ▶ Nonsquare $m \times n$ matrix \mathbf{A} has no inverse in usual sense
- ▶ If $\text{rank}(\mathbf{A}) = n$, *pseudoinverse* is defined by

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

and condition number by

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

- ▶ By convention, $\text{cond}(\mathbf{A}) = \infty$ if $\text{rank}(\mathbf{A}) < n$
- ▶ 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 $\mathbf{A}\mathbf{x} \cong \mathbf{b}$ is given by $\mathbf{x} = \mathbf{A}^+ \mathbf{b}$

Sensitivity and Conditioning

- ▶ Sensitivity of least squares solution to $\mathbf{Ax} \cong \mathbf{b}$ depends on \mathbf{b} as well as \mathbf{A}
- ▶ Define angle θ between \mathbf{b} and $\mathbf{y} = \mathbf{Ax}$ by

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

- ▶ Bound on perturbation $\Delta\mathbf{x}$ in solution \mathbf{x} due to perturbation $\Delta\mathbf{b}$ in \mathbf{b} is given by

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

Sensitivity and Conditioning, contnued

- ▶ Similarly, for perturbation \mathbf{E} in matrix \mathbf{A} ,

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

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

Solving Linear Least Squares Problems

Normal Equations Method

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

$$\mathbf{A}^T \mathbf{A} = \mathbf{L} \mathbf{L}^T$$

can be used to obtain solution \mathbf{x} to system of normal equations

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

which has same solution as linear least squares problem $\mathbf{A} \mathbf{x} \cong \mathbf{b}$

- ▶ Normal equations method involves transformations

rectangular \longrightarrow square \longrightarrow triangular

that preserve least squares solution in principle, but may not be satisfactory in finite-precision arithmetic

Example: Normal Equations Method

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

$$\begin{aligned}
 \mathbf{A}^T \mathbf{A} &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ -1.0 & -0.5 & 0.0 & 0.5 & 1.0 \\ 1.0 & 0.25 & 0.0 & 0.25 & 1.0 \end{bmatrix} \begin{bmatrix} 1 & -1.0 & 1.0 \\ 1 & -0.5 & 0.25 \\ 1 & 0.0 & 0.0 \\ 1 & 0.5 & 0.25 \\ 1 & 1.0 & 1.0 \end{bmatrix} \\
 &= \begin{bmatrix} 5.0 & 0.0 & 2.5 \\ 0.0 & 2.5 & 0.0 \\ 2.5 & 0.0 & 2.125 \end{bmatrix}, \\
 \mathbf{A}^T \mathbf{b} &= \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ -1.0 & -0.5 & 0.0 & 0.5 & 1.0 \\ 1.0 & 0.25 & 0.0 & 0.25 & 1.0 \end{bmatrix} \begin{bmatrix} 1.0 \\ 0.5 \\ 0.0 \\ 0.5 \\ 2.0 \end{bmatrix} = \begin{bmatrix} 4.0 \\ 1.0 \\ 3.25 \end{bmatrix}
 \end{aligned}$$

Example, continued

- ▶ Cholesky factorization of symmetric positive definite matrix $\mathbf{A}^T \mathbf{A}$ gives

$$\begin{aligned} \mathbf{A}^T \mathbf{A} &= \begin{bmatrix} 5.0 & 0.0 & 2.5 \\ 0.0 & 2.5 & 0.0 \\ 2.5 & 0.0 & 2.125 \end{bmatrix} \\ &= \begin{bmatrix} 2.236 & 0 & 0 \\ 0 & 1.581 & 0 \\ 1.118 & 0 & 0.935 \end{bmatrix} \begin{bmatrix} 2.236 & 0 & 1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & 0.935 \end{bmatrix} = \mathbf{L}\mathbf{L}^T \end{aligned}$$

- ▶ Solving lower triangular system $\mathbf{L}\mathbf{z} = \mathbf{A}^T \mathbf{b}$ by forward-substitution gives $\mathbf{z} = [1.789 \quad 0.632 \quad 1.336]^T$
- ▶ Solving upper triangular system $\mathbf{L}^T \mathbf{x} = \mathbf{z}$ by back-substitution gives $\mathbf{x} = [0.086 \quad 0.400 \quad 1.429]^T$

Shortcomings of Normal Equations

- ▶ Information can be lost in forming $\mathbf{A}^T \mathbf{A}$ and $\mathbf{A}^T \mathbf{b}$
- ▶ For example, take

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

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

- ▶ Then in floating-point arithmetic

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

which is singular

- ▶ Sensitivity of solution is also worsened, since

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

Augmented System Method

- ▶ Definition of residual together with orthogonality requirement give $(m + n) \times (m + n)$ augmented system

$$\begin{bmatrix} I & \mathbf{A} \\ \mathbf{A}^T & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{r} \\ \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}$$

- ▶ Augmented system is not positive definite, is larger than original system, and requires storing two copies of \mathbf{A}
- ▶ But it allows greater freedom in choosing pivots in computing \mathbf{LDL}^T or \mathbf{LU} factorization

Augmented System Method, continued

- ▶ Introducing scaling parameter α gives system

$$\begin{bmatrix} \alpha I & \mathbf{A} \\ \mathbf{A}^T & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{r}/\alpha \\ \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}$$

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

- ▶ Reasonable rule of thumb is to take

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

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

Orthogonalization Methods

Orthogonal Transformations

- ▶ We seek alternative method that avoids numerical difficulties of normal equations
- ▶ We need numerically robust transformation that produces easier problem without changing solution
- ▶ What kind of transformation leaves least squares solution unchanged?
- ▶ Square matrix Q is *orthogonal* if $Q^T Q = I$
- ▶ Multiplication of vector by orthogonal matrix preserves Euclidean norm

$$\|Q\mathbf{v}\|_2^2 = (Q\mathbf{v})^T Q\mathbf{v} = \mathbf{v}^T Q^T Q\mathbf{v} = \mathbf{v}^T \mathbf{v} = \|\mathbf{v}\|_2^2$$

- ▶ Thus, multiplying both sides of least squares problem by orthogonal matrix does not change its solution

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

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

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

- ▶ Residual is

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

Triangular Least Squares Problems, continued

- ▶ We have no control over second term, $\|\mathbf{b}_2\|_2^2$, but first term becomes zero if \mathbf{x} satisfies $n \times n$ triangular system

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

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 our strategy is to transform general least squares problem to triangular form using orthogonal transformation so that least squares solution is preserved

QR Factorization

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

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

where \mathbf{R} is $n \times n$ and upper triangular

- ▶ Linear least squares problem $\mathbf{Ax} \cong \mathbf{b}$ is then transformed into triangular least squares problem

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

which has same solution, since

$$\|\mathbf{r}\|_2^2 = \|\mathbf{b} - \mathbf{Ax}\|_2^2 = \|\mathbf{b} - \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x}\|_2^2 = \|\mathbf{Q}^T \mathbf{b} - \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x}\|_2^2$$

Orthogonal Bases

- ▶ If we partition $m \times m$ orthogonal matrix $\mathbf{Q} = [\mathbf{Q}_1 \ \mathbf{Q}_2]$, where \mathbf{Q}_1 is $m \times n$, 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 called *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 $\mathbf{A}\mathbf{x} \cong \mathbf{b}$ is given by solution to square system

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

Computing QR Factorization

- ▶ To compute QR factorization of $m \times n$ matrix \mathbf{A} , with $m > n$, we annihilate subdiagonal entries of successive columns of \mathbf{A} , eventually reaching upper triangular form
- ▶ Similar to LU factorization by Gaussian elimination, but use orthogonal transformations instead of elementary elimination matrices
- ▶ Available methods include
 - ▶ Householder transformations
 - ▶ Givens rotations
 - ▶ Gram-Schmidt orthogonalization

Householder QR Factorization

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} is orthogonal and symmetric: $\mathbf{H} = \mathbf{H}^T = \mathbf{H}^{-1}$
- ▶ Given vector \mathbf{a} , we want to 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} , we can take

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

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

Example: Householder Transformation

- ▶ If $\mathbf{a} = [2 \ 1 \ 2]^T$, then we take

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

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

- ▶ Since a_1 is positive, we choose negative sign for α to avoid cancellation, so $\mathbf{v} = \begin{bmatrix} 2 \\ 1 \\ 2 \end{bmatrix} - \begin{bmatrix} -3 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 5 \\ 1 \\ 2 \end{bmatrix}$
- ▶ To confirm that transformation works,

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

⟨ interactive example ⟩

Householder QR Factorization

- ▶ To compute QR factorization of \mathbf{A} , use Householder transformations to annihilate subdiagonal entries of each successive column
- ▶ Each Householder transformation is applied to entire matrix, but does not affect prior columns, so zeros are 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}

Householder QR Factorization, continued

- ▶ Process just described produces factorization

$$\mathbf{H}_n \cdots \mathbf{H}_1 \mathbf{A} = \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix}$$

where \mathbf{R} is $n \times n$ and upper triangular

- ▶ If $\mathbf{Q} = \mathbf{H}_1 \cdots \mathbf{H}_n$, then $\mathbf{A} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix}$
- ▶ To preserve solution of linear least squares problem, right-hand side \mathbf{b} is transformed by same sequence of Householder transformations
- ▶ Then solve triangular least squares problem $\begin{bmatrix} \mathbf{R} \\ \mathbf{O} \end{bmatrix} \mathbf{x} \cong \mathbf{Q}^T \mathbf{b}$

Householder QR Factorization, continued

- ▶ For solving linear least squares problem, product \mathbf{Q} of Householder transformations need not be formed explicitly
- ▶ \mathbf{R} can be stored in upper triangle of array initially containing \mathbf{A}
- ▶ Householder vectors \mathbf{v} can be stored in (now zero) lower triangular portion of \mathbf{A} (almost)
- ▶ Householder transformations most easily applied in this form anyway

Example: Householder QR Factorization

- ▶ For polynomial data-fitting example given previously, with

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

- ▶ Householder vector \mathbf{v}_1 for annihilating subdiagonal entries of first column of \mathbf{A} is

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} -2.236 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 3.236 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

Example, continued

- ▶ Applying resulting Householder transformation \mathbf{H}_1 yields transformed matrix and right-hand side

$$\mathbf{H}_1 \mathbf{A} = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & -0.191 & -0.405 \\ 0 & 0.309 & -0.655 \\ 0 & 0.809 & -0.405 \\ 0 & 1.309 & 0.345 \end{bmatrix}, \quad \mathbf{H}_1 \mathbf{b} = \begin{bmatrix} -1.789 \\ -0.362 \\ -0.862 \\ -0.362 \\ 1.138 \end{bmatrix}$$

- ▶ Householder vector \mathbf{v}_2 for annihilating subdiagonal entries of second column of $\mathbf{H}_1 \mathbf{A}$ is

$$\mathbf{v}_2 = \begin{bmatrix} 0 \\ -0.191 \\ 0.309 \\ 0.809 \\ 1.309 \end{bmatrix} - \begin{bmatrix} 0 \\ 1.581 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ -1.772 \\ 0.309 \\ 0.809 \\ 1.309 \end{bmatrix}$$

Example, continued

- ▶ Applying resulting Householder transformation \mathbf{H}_2 yields

$$\mathbf{H}_2\mathbf{H}_1\mathbf{A} = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & -0.725 \\ 0 & 0 & -0.589 \\ 0 & 0 & 0.047 \end{bmatrix}, \quad \mathbf{H}_2\mathbf{H}_1\mathbf{b} = \begin{bmatrix} -1.789 \\ 0.632 \\ -1.035 \\ -0.816 \\ 0.404 \end{bmatrix}$$

- ▶ Householder vector \mathbf{v}_3 for annihilating subdiagonal entries of third column of $\mathbf{H}_2\mathbf{H}_1\mathbf{A}$ is

$$\mathbf{v}_3 = \begin{bmatrix} 0 \\ 0 \\ -0.725 \\ -0.589 \\ 0.047 \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 0.935 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -1.660 \\ -0.589 \\ 0.047 \end{bmatrix}$$

Example, continued

- ▶ Applying resulting Householder transformation \mathbf{H}_3 yields

$$\mathbf{H}_3 \mathbf{H}_2 \mathbf{H}_1 \mathbf{A} = \begin{bmatrix} -2.236 & 0 & -1.118 \\ 0 & 1.581 & 0 \\ 0 & 0 & 0.935 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{H}_3 \mathbf{H}_2 \mathbf{H}_1 \mathbf{b} = \begin{bmatrix} -1.789 \\ 0.632 \\ 1.336 \\ 0.026 \\ 0.337 \end{bmatrix}$$

- ▶ Now solve upper triangular system $\mathbf{R}\mathbf{x} = \mathbf{c}_1$ by back-substitution to obtain $\mathbf{x} = [0.086 \quad 0.400 \quad 1.429]^T$

⟨ interactive example ⟩

Givens QR Factorization

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

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

Givens Rotations, continued

- ▶ Back-substitution then gives

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

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

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

Example: Givens Rotation

- ▶ Let $\mathbf{a} = [4 \ 3]^T$
- ▶ To annihilate second entry we compute cosine and sine

$$c = \frac{a_1}{\sqrt{a_1^2 + a_2^2}} = \frac{4}{5} = 0.8 \quad \text{and} \quad s = \frac{a_2}{\sqrt{a_1^2 + a_2^2}} = \frac{3}{5} = 0.6$$

- ▶ Rotation is then given by

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

- ▶ To confirm that rotation works,

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

Givens QR Factorization

- ▶ More generally, to annihilate selected component of vector in n dimensions, rotate target component with another component

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

- ▶ By systematically annihilating successive entries, we can reduce matrix to upper triangular form using sequence of Givens rotations
- ▶ Each rotation is orthogonal, so their product is orthogonal, producing QR factorization

Givens 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 specify it
- ▶ These disadvantages can be overcome, but more complicated implementation is required
- ▶ Givens can be advantageous for computing QR factorization when many entries of matrix are already zero, since those annihilations can then be skipped

[⟨ interactive example ⟩](#)

Gram-Schmidt QR Factorization

Gram-Schmidt Orthogonalization

- ▶ Given vectors \mathbf{a}_1 and \mathbf{a}_2 , we seek orthonormal vectors \mathbf{q}_1 and \mathbf{q}_2 having same span
- ▶ This can be accomplished by subtracting from second vector its projection onto first vector and normalizing both resulting vectors, as shown in diagram



[⟨ interactive example ⟩](#)

Gram-Schmidt Orthogonalization

- ▶ Process can be extended to any number of vectors $\mathbf{a}_1, \dots, \mathbf{a}_k$, orthogonalizing each successive vector against all preceding ones, giving *classical Gram-Schmidt* procedure

```

for  $k = 1$  to  $n$ 
     $\mathbf{q}_k = \mathbf{a}_k$ 
    for  $j = 1$  to  $k - 1$ 
         $r_{jk} = \mathbf{q}_j^T \mathbf{a}_k$ 
         $\mathbf{q}_k = \mathbf{q}_k - r_{jk} \mathbf{q}_j$ 
    end
     $r_{kk} = \|\mathbf{q}_k\|_2$ 
     $\mathbf{q}_k = \mathbf{q}_k / r_{kk}$ 
end
  
```

- ▶ Resulting \mathbf{q}_k and r_{jk} form reduced QR factorization of \mathbf{A}

Modified Gram-Schmidt

- ▶ Classical Gram-Schmidt procedure often suffers loss of orthogonality in finite-precision arithmetic
- ▶ Also, separate storage is required for \mathbf{A} , \mathbf{Q} , and \mathbf{R} , since original \mathbf{a}_k are needed in inner loop, so \mathbf{q}_k cannot overwrite columns of \mathbf{A}
- ▶ Both deficiencies are 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

Modified Gram-Schmidt QR Factorization

- ▶ Modified Gram-Schmidt algorithm

```
for  $k = 1$  to  $n$   
     $r_{kk} = \|\mathbf{a}_k\|_2$   
     $\mathbf{q}_k = \mathbf{a}_k / r_{kk}$   
    for  $j = k + 1$  to  $n$   
         $r_{kj} = \mathbf{q}_k^T \mathbf{a}_j$   
         $\mathbf{a}_j = \mathbf{a}_j - r_{kj} \mathbf{q}_k$   
    end  
end
```

⟨ interactive example ⟩

Rank Deficiency

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 is often not clear cut in practice, so relative tolerance is used to determine rank

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} is extremely close to singular (exactly singular to 3-digit accuracy of problem statement)
- ▶ If \mathbf{R} is 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 are nearly linearly dependent

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

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

where \mathbf{R} is $k \times k$, upper triangular, and nonsingular, and permutation matrix \mathbf{P} performs column interchanges

QR with Column Pivoting, continued

- ▶ *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

$$\mathbf{x} = \mathbf{P} \begin{bmatrix} \mathbf{z} \\ \mathbf{0} \end{bmatrix}$$

- ▶ *Minimum-norm solution* can be computed, if desired, at expense of additional processing to annihilate \mathbf{S}
- ▶ $\text{rank}(\mathbf{A})$ is usually unknown, so rank is determined by monitoring norms of remaining unreduced columns and terminating factorization when maximum value falls below chosen tolerance

⟨ interactive example ⟩

Singular Value Decomposition

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} , are 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} are called left and right *singular vectors*

Example: SVD

► SVD of $\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \\ 10 & 11 & 12 \end{bmatrix}$ is given by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T =$

$$\begin{bmatrix} .141 & .825 & -.420 & -.351 \\ .344 & .426 & .298 & .782 \\ .547 & .0278 & .664 & -.509 \\ .750 & -.371 & -.542 & .0790 \end{bmatrix} \begin{bmatrix} 25.5 & 0 & 0 \\ 0 & 1.29 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} .504 & .574 & .644 \\ -.761 & -.057 & .646 \\ .408 & -.816 & .408 \end{bmatrix}$$

< interactive example >

Applications of SVD

- ▶ *Minimum norm solution* to $\mathbf{Ax} \cong \mathbf{b}$ is given by

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

For ill-conditioned or rank deficient problems, “small” singular values can be omitted from summation to stabilize solution

- ▶ *Euclidean matrix norm*: $\|\mathbf{A}\|_2 = \sigma_{\max}$

- ▶ *Euclidean condition number of matrix*: $\text{cond}_2(\mathbf{A}) = \frac{\sigma_{\max}}{\sigma_{\min}}$

- ▶ *Rank of matrix*: $\text{rank}(\mathbf{A}) =$ number of nonzero singular values

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} is 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 $\mathbf{x} = \mathbf{A}^+ \mathbf{b}$

Orthogonal Bases

- ▶ SVD of matrix, $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, provides orthogonal bases for subspaces relevant to \mathbf{A}
- ▶ 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 of \mathbf{A}

Lower-Rank Matrix Approximation

- ▶ Another way to write SVD is

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

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 computed using only $m + n$ multiplications
- ▶ Condensed approximation to \mathbf{A} is 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.

[⟨ interactive example ⟩](#)

Total Least Squares

- ▶ Ordinary least squares is applicable when right-hand side \mathbf{b} is subject to random error but matrix \mathbf{A} is known accurately
- ▶ When all data, including \mathbf{A} , are subject to error, then *total* least squares is 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 $[\mathbf{A}, \mathbf{b}]$

Comparison of Methods for Least Squares

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 $m \approx n$, both 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 is proportional to $mn^2 + n^3$, with proportionality constant ranging from 4 to 10, depending on algorithm used

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

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 sufficient accuracy
- ▶ For rank-deficient or nearly rank-deficient problems, 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