ECE133A Discussion

Xin Jiang

Department of Electrical and Computer Engineering, UCLA

ECE133A Applied Numerical Computing

September 30, 2021

Course logistics

- weekly homework: due on Friday via Gradescope
- a project (tentative)
- midterm: open-book, Tuesday, May 4, 4pm–5:50pm (in class)
- final: open-book, Monday, June 7, 6:30pm-9:30pm
- course materials: on CCLE

Introduction to MATLAB

- you have free access to MATLAB via SEASNET student account
- the official site offers a nice start-up tutorial
- you are not expected to have a strong background in programming
- the programs you write will use only a tiny subset of MATLAB features

Introduction to Julia

- Julia is a new programming language for scientific computing
- friendly syntax for building math constructs like vectors, matrices
- official site: you can download the software and find a tutorial there
- Jupyter is a open-source web application on which you can create and share live codes, visualizations, and narrative text
- Julia companion for textbook

Outline

Matrices

Matrix inverse

orthogonal matrices

QR factorization

LU factorization

least squares

nonlinear least squares

Cholesky factorization

mathematical background

matrix norms

condition and stability

IEEE floating point numbers

Complexity

Flop count

- 1 flop = one basic arithmetic operation in $\mathbf{R}$ or $\mathbf{C}$
- flop count is the total number of operations in an algorithm
- keep dominant term (with coefficients)

$$(1/3)n^3 + 100n^2 + 10n + 5 \approx (1/3)n^3$$

Examples

- inner product between two n-vectors: $2n - 1 \approx 2n$ flops
- matrix–vector multiplication of $m \times n$ matrix A and n-vector x:

$$y = Ax \qquad (2n-1)m \approx 2mn \text{ flops}$$

- product of $m \times n$ matrix A and $n \times p$ matrix B:

$$C = AB \qquad mp(2n-1) \approx 2mnp \text{ flops}$$

Matrix representation: adjacency matrices

suppose A is the adjacency matrix of a directed graph with n vertices

$$A_{ij} = \begin{cases} 1 & \text{there is a edge from vertex } j \text{ to vertex } i \\ 0 & \text{otherwise} \end{cases}$$

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad A^2 = \begin{bmatrix} 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 2 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix}$$

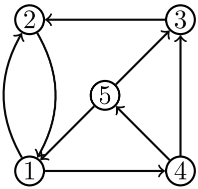

Matrix representation: adjacency matrices

examine the expression for the i, j element of the square of A:

$$(A^2)_{ij} = \sum_{k=1}^{n} A_{ik} A_{kj}$$

what's the graph associated with $B = I + A$?

now show the equivalence between

- all the elements of the matrix $(I + A)^{n-1}$ are positive
- for any vertex i and j, there is a directed path from i to j

Regression line

let a, b be two real n-vectors, and denote

$$m_a = \mathbf{avg}(a) = \frac{\mathbf{1}^T a}{n}, \qquad m_b = \mathbf{avg}(b) = \frac{\mathbf{1}^T b}{n},$$

$$s_a = \mathbf{std}(a) = \frac{1}{\sqrt{n}}\|a - m_a\mathbf{1}\|, \qquad s_b = \mathbf{std}(b) = \frac{1}{\sqrt{n}}\|b - m_b\mathbf{1}\|$$

$$\rho = \frac{1}{n}\frac{(a - m_a\mathbf{1})^T(b - m_b\mathbf{1})}{s_a s_b}$$

we fit a straight line to the points (a_k, b_k), by minimizing

$$J = \frac{1}{n}\sum_{k=1}^{n}(c_1 + c_2 a_k - b_k)^2 = \frac{1}{n}\|c_1\mathbf{1} + c_2 a - b\|^2$$

we found that the optimal coefficients are $c_2 = \rho s_a / s_b$ and $c_1 = m_b - m_a c_2$

show that for those values of c_1 and c_2, we have $J = (1 - \rho^2) s_b^2$

Outline

Matrix inverse

for a square matrix $A \in \mathbf{R}^{n \times n}$, **nonsingular = invertible**

$$B \text{ is the inverse of } A \quad \Longleftrightarrow \quad AB = I,\ BA = I$$

the following four properties are equivalent

1. A is left invertible
2. the columns of A are linearly independent
3. A is right invertible
4. the rows of A are linearly independent

Exercise: are the following matrices nonsingular?

- $A = ab^T$ where a and b are n-vectors and $n > 1$
- $A = I - ab^T$ where a and b are n-vectors with $\|a\|\|b\| < 1$

Examples on matrix inverse

suppose A is a nonsingular $n \times n$ matrix, u, v are n-vectors, $v^T A^{-1} u \neq -1$

show that $A + uv^T$ is nonsingular with inverse

$$(A + uv^T)^{-1} = A^{-1} - \frac{1}{1 + v^T A^{-1} u} A^{-1} u v^T A^{-1}$$

consider the $(n+1) \times (n+1)$ matrix $A = \begin{bmatrix} I & a \\ a^T & 0 \end{bmatrix}$, where a is an n-vector

1. when is A invertible?
2. assuming A is invertible, give an expression for the inverse matrix A^{-1}

Example: Vandermonde matrix

$$A = \begin{bmatrix} 1 & t_1 & t_1^2 & \cdots & t_1^{n-1} \\ 1 & t_2 & t_2^2 & \cdots & t_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & t_n & t_n^2 & \cdots & t_n^{n-1} \end{bmatrix} \quad \text{with } t_i \neq t_j \text{ for } i \neq j$$

we show that A is nonsingular by showing that $Ax = 0$ only if $x = 0$

- $Ax = 0$ means $p(t_1) = p(t_2) = \cdots = p(t_n) = 0$ where

  $$p(t) = x_1 + x_2 t + x_3 t^2 + \cdots + x_n t^{n-1}$$

  $p(t)$ is a polynomial of degree $n-1$ or less
- if $x \neq 0$, then $p(t)$ cannot have more than $n-1$ distinct real roots
- therefore $p(t_1) = \cdots = p(t_n) = 0$ is only possible if $x = 0$

Polynomial interpolation

in this problem we construct polynomials

$$p(t) = x_1 + x_2 t + x_3 t^2 + \cdots + x_n t^{n-1}$$

to interpolate points on the graph of the function $f(t) = 1/(1 + 25t^2)$

we first generate n pairs (t_i, y_i). We then solve a set of linear equations

$$\begin{bmatrix} 1 & t_1 & t_1^2 & \cdots & t_1^{n-1} \\ 1 & t_2 & t_2^2 & \cdots & t_2^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & t_n & t_n^2 & \cdots & t_n^{n-1} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_{n-1} \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{n-1} \\ y_n \end{bmatrix}$$

to find the coefficients x_i

we then plot the polynomials and the function f in the interval $[-1, 1]$

the figures below show the interpolation for $n = 5, 10, 15, 16$, respectively

Example on interpolation

express the following problem as a set of linear equations $Ax = b$

find a rational function

$$f(t) = \frac{x_1 + x_2 t + x_3 t^2}{1 + x_4 t + x_5 t^2}$$

that satisfies the five conditions

$$f(0) = b_1, \quad f^{(1)}(0) = b_2, \quad \frac{f^{(2)}(0)}{2} = b_3, \quad \frac{f^{(3)}(0)}{6} = b_4, \quad \frac{f^{(5)}(0)}{24} = b_5,$$

where $b_1, \ldots, b_5$ are given

Left inverse and right inverse

for tall matrices $A \in \mathbf{R}^{m \times n}$ $(m > n)$, the following properties are equivalent

1. A is left invertible
2. the columns of A are linearly independent
3. $A^T A$ is nonsingular

the pseudo-inverse of such matrices is given by $A^\dagger = (A^T A)^{-1} A^T$

for wide matrices $A \in \mathbf{R}^{m \times n}$ $(m < n)$, the following properties are equivalent

1. A is right invertible
2. the rows of A are linearly independent
3. AA^T is nonsingular

the pseudo-inverse of such matrices is given by $A^\dagger = A^T (AA^T)^{-1}$

Pseudo-inverse

tall matrix ($m > n$) with independent cols	wide matrix ($m < n$) with independent rows	nonsingular matrix ($m = n$)
$A^\dagger = (A^TA)^{-1}A^T$	$A^\dagger = A^T(AA^T)^{-1}$	$A^\dagger = A^{-1}$
A^TA is nonsingular	AA^T is nonsingular	
$A^\dagger A = I$	$AA^\dagger = I$	

	existence	unique
inverse	square nonsingular	Y
left inverse	matrix with linearly independent cols	N
right inverse	matrix with linearly independent rows	N
pseudo-inverse	all matrices	Y

Example on pseudo-inverse

$$(AB)^\dagger = B^\dagger A^\dagger? \qquad (\blacktriangle)$$

consider the following example

$$A = \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}, \quad AB = \begin{bmatrix} 2 & 1 \end{bmatrix}$$

the pseudo-inverses are

$$A^\dagger = \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix}, \quad B^\dagger = \begin{bmatrix} 1/2 & 0 \\ 0 & 1 \end{bmatrix}, \quad (AB)^\dagger = \begin{bmatrix} 2/5 \\ 1/5 \end{bmatrix}$$

we have $(AB)(B^\dagger A^\dagger) = I$ but $B^\dagger A^\dagger \neq (AB)^\dagger$

- is ($\blacktriangle$) true when A has linearly independent columns and B is nonsingular?
- is ($\blacktriangle$) true when A is nonsingular and B has linearly independent columns?

Outline

Orthogonal matrices

Tall matrix with orthonormal columns

$$A^T A = I, \quad AA^T \neq I$$

- properties: preservation of inner products, norms, distance, and angles
- left-invertibility
- projection of x on the range of A: $AA^T b$

Orthogonal matrices: a square real matrix with orthonormal columns

$$Q^T Q = I, \quad QQ^T = I, \quad Q^{-1} = Q^T$$

- examples: permutation matrix, plane rotation, reflector
- linear equation with orthogonal matrix

Exercise: when is a matrix lower-triangular and orthogonal?

Examples on orthogonal matrices

let Q be an $n \times n$ orthogonal matrix, partitioned as

$$Q = \begin{bmatrix} Q_1 & Q_2 \end{bmatrix}$$

where $Q_1 \in \mathbf{R}^{n \times m}$ and $Q_2 \in \mathbf{R}^{n \times (n-m)}$ (assume $0 < m < n$)
consider the matrix $A = Q_1 Q_1^T - Q_2 Q_2^T$

1. show that $A = 2Q_1 Q_1^T - I = I - 2Q_2 Q_2^T$
2. show that A is orthogonal

for what property of the matrix B is a matrix of the form

$$A = \frac{1}{\sqrt{2}} \begin{bmatrix} I & B^T \\ -B & I \end{bmatrix}$$

orthogonal? nonsingular?

Example on orthogonal matrices

let a be an n-vector with $\|a\| = 1$; define the $2n \times 2n$ matrix

$$A = \begin{bmatrix} aa^T & I - aa^T \\ I - aa^T & aa^T \end{bmatrix}$$

1. show that A is orthogonal
2. now suppose $n = 2$; given the plots of b and c, indicate on the figure the 2-vectors x, y that solve the 4×4 equation

$$\begin{bmatrix} aa^T & I - aa^T \\ I - aa^T & aa^T \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ c \end{bmatrix}$$

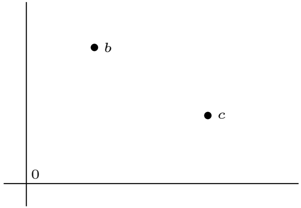

line through a and the origin

Outline

Triangular matrices

- definition
- forward/back substitution
- inverse of a nonsingular triangular matrix A is also triangular, with

$$(A^{-1})_{ii} = 1/A_{ii}$$

- A^{-1} is computed by solving $AX = I$ column by column ($(1/3)n^3$ flops)

Exercise: the trace of a matrix is the sum of its diagonal elements; *i.e.*,

$$\operatorname{tr} A = \sum_{i=1}^{n} A_{ii}$$

what is the complexity of computing $\operatorname{tr}(A^{-1})$ if A is triangular and nonsingular

QR factorization

suppose $A \in \mathbf{R}^{m \times n}$ has linearly independent columns; A can be factored as

$$A = QR$$

where

- Q is $m \times n$ with orthonormal columns
- R is $n \times n$ and upper-triangular with nonzero diagonal elements
- by convention, we require $R_{ii} > 0$

Properties

- pseudo-inverse: $A^\dagger = R^{-1}Q^T$
- $\mathrm{range}(A) = \mathrm{range}(Q)$
- projection of x on the range of A: $AA^\dagger x = QQ^T x$
- algorithms: Gram–Schimdt, Householder ($2mn^2$ flops)
- application: linear equations, least squares

Outline

LU factorization

LU factorization (with row pivoting)

$$A = PLU$$

- P permutation matrix, L unit lower triangular, U upper triangular
- exists if and only if A is nonsingular, but not unique
- complexity: $(2/3)n^3$ if A is $n \times n$

Solving linear equations $Ax = b$ by LU factorization

1. factor A as $A = PLU$ ($(2/3)n^3$ flops)
2. solve $(PLU)x = b$ in three steps
   (a) permutation: $z_1 = P^T b$ (0 flop)
   (b) forward substitution: solve $Lz_2 = z_1$ (n^2 flops)
   (c) back substitution: solve $Ux = z_2$ (n^2 flops)

total complexity: $(2/3)n^3 + 2n^2 \approx (2/3)n^3$ flops

Examples on solving linear equations

suppose A is an $n \times n$ matrix, and u and v are n-vectors
in each of the following cases, what is the complexity of computing the matrix

$$B = A^{-1}uv^TA^{-1}$$

1. A is diagonal with nonzero diagonal elements
2. A is lower-triangular with nonzero diagonal elements
3. A is orthogonal
4. A is a general nonsingular matrix

assume we already have the LU factorization $A = PLU$
describe an algorithm for each of the following problems

1. compute the jth column of A^{-1}
2. compute the sum of columns of A^{-1}
3. compute the sum of rows of A^{-1}

Examples on solving linear equations

consider a square $(n+m) \times (n+m)$ matrix

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}$$

with $A \in \mathbf{R}^{n\times n}$ and $D \in \mathbf{R}^{m\times m}$
describe efficient algorithms for computing the Schur complement

$$S = D - CA^{-1}B$$

of each of the following types of matrices A

1. A is diagonal with nonzero diagonal elements
2. A is lower triangular with nonzero diagonal elements
3. A is a general nonsingular matrix

Outline

Least squares

the least squares problem is an unconstrained optimization problem

$$\text{minimize} \quad \|Ax - b\|^2$$

with variable $x \in \mathbf{R}^n$ and coefficients $A \in \mathbf{R}^{m \times n}$, $b \in \mathbf{R}^m$

- assume A has linearly independent columns
- normal equation: $A^T A\hat{x} = A^T b$
- suppose QR factorization of A is given by $A = QR$

$$\hat{x} = A^\dagger b = (A^T A)^{-1} A^T b = R^{-1} Q^T b$$

1. compute QR factorization $A = QR$ ($2mn^2$ flops)
2. matrix-vector product $d = Q^T b$ ($2mn$ flops)
3. solve $Rx = d$ by back substitution (n^2 flops)

Typical least squares problems

suppose $\hat{x}$ is the solution for the least squares problem

$$\text{minimize} \quad \|Ax - b\|^2;$$

and $\hat{y}$ is the solution for the least squares problem

$$\text{minimize} \quad \|\tilde{A}y - \tilde{b}\|^2$$

show that $\hat{y} = g(\hat{x})$ by verifying

$$\tilde{A}^T \tilde{A} g(\hat{x}) = \tilde{A}^T \tilde{b}, \quad \text{where} \quad A^T A \hat{x} = A^T b$$

Exercise: suppose QR factorization $\begin{bmatrix} A & b \end{bmatrix} = QR$ can be partitioned as

$$Q = \begin{bmatrix} Q_1 & Q_2 \end{bmatrix}, \quad R = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}$$

show that the LS solution $\tilde{x}_{\text{ls}} = R_{11}^{-1} R_{12}$ and $R_{22} = \|A\tilde{x}_{\text{ls}} - b\|$

Example: K-fold cross-validation

given $m \times n$ matrices $A_1, \ldots, A_K$, and m-vectors $b_1, \ldots, b_K$
matrices C_k is constructed by stacking $A_1, \ldots, A_K$, but skipping A_k

$$C_k = \begin{bmatrix} A_1 \\ \vdots \\ A_{k-1} \\ A_{k+1} \\ \vdots \\ A_K \end{bmatrix}, \quad d_k = \begin{bmatrix} b_1 \\ \vdots \\ b_{k-1} \\ b_{k+1} \\ \vdots \\ b_K \end{bmatrix}$$

C_k has size $((K-1)m) \times n$; assume C_k has linearly independent columns

define $\hat{x}^{(k)}$ as the solution of the least squares problem

$$\text{minimize} \quad \|C_k x - d_k\|^2$$

what is the complexity for computing K least squares solutions $\hat{x}^{(1)}, \ldots, \hat{x}^{(K)}$?

Least squares data fitting

1. identify the unknown variable x
2. transfer nonlinear functions into a linear function of x
3. write the problem into least-squares form

Exercise: A8.3, A8.6
the m data points (t_i, y_i) are well approximated by a function of the form

$$f(t) = \frac{e^{\alpha t+\beta}}{1+e^{\alpha t+\beta}}$$

formulate the following problem as a least squares problem:

find values of the parameters α, β such that

$$\frac{e^{\alpha t_i+\beta}}{1+e^{\alpha t_i+\beta}} \approx y_i, \quad i = 1, \dots, m$$

Multi-objective least squares

many other problems can be transformed into a least squares problem

- multi-objective least squares

$$\text{minimize} \quad \lambda_1 \|A_1 x - b_1\|^2 + \cdots + \lambda_k \|A_k x - b_k\|^2$$

  with all positive λ_i's

- Tokhonov regularization ($\lambda > 0$)

$$\text{minimize} \quad \|Ax - y\|^2 + \lambda \|x\|^2$$

  where the solution is

$$\hat{x} = (A^T A + \lambda I)^{-1} A^T y = A^T (A A^T + \lambda I)^{-1} y$$

  this avoids the QR factorization when A is very wide ($m \ll n$)

Example: regularized least squares image deblurring

the $\mathbf{vec}$ operation creates an n^2-vector x by converting an $n \times n$ matrix X in the column-major order:

$$x = \mathbf{vec}(X) = \begin{bmatrix} X_{1:n,1} \\ X_{1:n,2} \\ \vdots \\ X_{1:n,n} \end{bmatrix}$$

conversely, $\mathbf{mat}$ is the inverse operation of $\mathbf{vec}$, *i.e.*,

$$X = \mathbf{mat}(x) = \begin{bmatrix} x_{1:n} & x_{(n+1):2n} & \cdots & x_{(n(n-1)+1):n^2} \end{bmatrix}$$

Example: regularized least squares image deblurring

we write the discrete Fourier transform in terms of the $n \times n$ DFT matrix W:

$$V = WUW \qquad \texttt{V=fft2(U)}$$

$$U = (1/n^2)W^H V W^H \qquad \texttt{U=ifft2(V)}$$

then we can rewrite the discrete Fourier transform in vector form with $u = \mathbf{vec}(U)$ and $v = \mathbf{vec}(V)$, *i.e.*,

$$v = \widetilde{W}u \qquad \texttt{v=reshape(fft2(reshape(u,n,n)), n}\wedge\texttt{2, 1)}$$

$$u = \widetilde{W}^{-1}v \qquad \texttt{u=reshape(ifft2(reshape(v,n,n)), n}\wedge\texttt{2, 1)}$$

where $\widetilde{W} = W \otimes W \in \mathbf{R}^{n^2 \times n^2}$

since $(1/n)W^H W = I$, we have

$$\widetilde{W}^H \widetilde{W} = n^2 I, \quad \widetilde{W}\widetilde{W}^H = n^2 I, \quad \widetilde{W}^{-1} = \frac{1}{n^2}\widetilde{W}^H$$

Example: regularized least squares image deblurring

now we are ready to discuss the image deblurring problem

it is a regularized least squares problem:

$$\text{minimize} \quad \|Ax - y\|^2 + \lambda(\|D_{\rm v}x\|^2 + \|D_{\rm h}x\|^2),$$

where $A = T(B)$, $D_{\rm v} = T(E)$, and $D_{\rm h} = T(E^T)$;
the coefficient matrices $B \in \mathbf{R}^{n\times n}$ and $E \in \mathbf{R}^{n\times n}$ are given

define function $T : \mathbf{R}^{n\times n} \to \mathbf{R}^{n\times n}$:

$$T(X) = \frac{1}{n^2}\widetilde{W}^H \operatorname{diag}(\widetilde{W}x)\widetilde{W},$$

where $x = \mathbf{vec}(X)$

this structure is called *block-circulant with circulant blocks* (BCCB)

the normal equation is given by

$$(A^H A + \lambda D_{\rm v}^H D_{\rm v} + \lambda D_{\rm h}^H D_{\rm h})x = A^H y$$

Example: regularized least squares image deblurring

Least norm problem

$$\begin{array}{ll} \text{minimize} & \|x\|^2 \\ \text{subject to} & Cx = d \end{array}$$

the variable is $x \in \mathbf{R}^n$, and $C \in \mathbf{R}^{p \times n}$ with $p < n$

Assumption: the coefficient matrix has linearly independent rows

Solution: the solution of the above least norm problem is

$$\hat{x} = C^{\dagger} d = C^T (CC^T)^{-1} d.$$

Constrained least squares

$$\begin{array}{ll} \text{minimize} & \|Ax - b\|^2 \\ \text{subject to} & Cx = d \end{array}$$

the variable is $x \in \mathbf{R}^n$; $A \in \mathbf{R}^{m\times n}$, $b \in \mathbf{R}^m$, $C \in \mathbf{R}^{p\times n}$, and $d \in \mathbf{R}^p$

we make following assumptions in our discussion:

1. the stacked $(m+p)\times n$ matrix $\begin{bmatrix} A \\ C \end{bmatrix}$ has linearly independent columns
2. C has linearly independent rows

hence, $\hat{x}$ solves the constrained LS problem iff there exists a z such that

$$\begin{bmatrix} A^TA & C^T \\ C & 0 \end{bmatrix} \begin{bmatrix} \hat{x} \\ z \end{bmatrix} = \begin{bmatrix} A^Tb \\ d \end{bmatrix}$$

(the assumptions ensure that the matrix on the lefthand side is nonsingular)

Example on constrained least squares

solve the following constrained least squares problems

1. $A \in \mathbf{R}^{m \times n}$ has linearly independent columns, $b \in \mathbf{R}^n$, $c \in \mathbf{R}^n$, and $d \in \mathbf{R}$

$$\begin{array}{ll} \text{minimize} & \|Ax - b\|^2 \\ \text{subject to} & c^T x = d \end{array}$$

   where the optimization variable is $x \in \mathbf{R}^n$

2. $A \in \mathbf{R}^{m \times n}$ has linearly independent columns, $b \in \mathbf{R}^m$, $c \in \mathbf{R}^m$

$$\begin{array}{ll} \text{minimize} & \|x - b\|^2 + \|y - c\|^2 \\ \text{subject to} & A^T x = A^T y \end{array}$$

   where the optimization variable $x, y \in \mathbf{R}^m$

Example on constrained least squares

let A be an $m \times n$ matrix with linearly independent columns

1. show that $\tilde{x}^{(i)}$ is the solution for the constrained least squares problem

$$\begin{array}{ll} \text{minimize} & \|Ax\|^2 \\ \text{subject to} & e_i^T x = -1 \end{array} \qquad \Longrightarrow \qquad \tilde{x}^{(i)} = -\frac{1}{e_i^T (A^T A)^{-1} e_i} (A^T A)^{-1} e_i$$

2. show that $\hat{x}^{(i)}$ is the solution for the constrained least squares problem

$$\begin{array}{ll} \text{minimize} & \|Ax - b\|^2 \\ \text{subject to} & e_i^T x = 0 \end{array} \qquad \Longrightarrow \qquad \hat{x}^{(i)} = \hat{x} - \frac{\hat{x}_i}{e_i^T (A^T A)^{-1} e_i} (A^T A)^{-1} e_i$$

   where $\hat{x}$ is the minimizer of $\|Ax - b\|^2$

Least squares summary

- (linear) least squares

$$\text{minimize} \quad \|Ax - b\|^2 \qquad \Longrightarrow \qquad \hat{x} = (A^T A)^{-1} A^T b$$

- least norm

$$\begin{array}{ll} \text{minimize} & \|x\|^2 \\ \text{subject to} & Cx = d \end{array} \qquad \Longrightarrow \qquad \hat{x} = C^T (CC^T)^{-1} d$$

- constrained least squares

$$\begin{array}{ll} \text{minimize} & \|Ax - b\|^2 \\ \text{subject to} & Cx = d \end{array} \qquad \Longrightarrow \qquad \begin{bmatrix} A^T A & C^T \\ C & 0 \end{bmatrix} \begin{bmatrix} \hat{x} \\ z \end{bmatrix} = \begin{bmatrix} A^T b \\ d \end{bmatrix}$$

Outline

Nonlinear least squares

$$\text{minimize} \quad g(x) = \|f(x)\|^2 = \sum_{i=1}^{m} f_i^2(x)$$

- **Gauss–Newton method:** at iteration k, we solve a least squares problem

$$\text{minimize} \quad \|f(x^{(k)}) + Df(x^{(k)})(x - x^{(k)})\|^2$$

$$\implies \quad x^{(k+1)} = x^{(k)} - (A^T A)^{-1} A^T f(x^{(k)}), \quad \text{where } A = Df(x^{(k)})$$

- **Levenberg–Marquardt:** at iteration k, we solve a regularized version

$$\text{minimize} \quad \|f(x^{(k)}) + Df(x^{(k)})(x - x^{(k)})\|^2 + \lambda^{(k)}\|x - x^{(k)}\|^2$$

$$\implies x^{(k+1/2)} = x^{(k)} - (A^T A + \lambda^{(k)} I)^{-1} A^T f(x^{(k)}), \quad \text{where } A = Df(x^{(k)})$$

$$\implies \begin{cases} x^{(k+1)} = x^{(k+1/2)},\ \lambda^{(k+1)} = \beta_1 \lambda^{(k)} & \text{if } \|f(x^{(k+1/2)})\|^2 < \|f(x^{(k)})\|^2 \\ x^{(k+1)} = x^{(k)},\ \lambda^{(k+1)} = \beta_2 \lambda^{(k)} & \text{otherwise} \end{cases}$$

Example: fitting an ellipse to points in a plane

an ellipse in a plane can be described as the set of points

$$\hat{f}(t;\theta) = \begin{bmatrix} c_1 + r\cos(\alpha + t) + \delta\cos(\alpha - t) \\ c_2 + r\sin(\alpha + t) + \delta\sin(\alpha - t) \end{bmatrix},$$

where $t \in [0, 2\pi]$, and $\theta = (c_1, c_2, r, \delta, \alpha)$
we consider the problem of fitting an ellipse to N points $x^{(1)}, \ldots, x^{(N)}$ in a plane:

$$\text{minimize} \quad \sum_{i=1}^{N} \|\hat{f}(t^{(i)};\theta) - x^{(i)}\|^2$$

where the optimization variables are $t^{(1)}, \ldots, t^{(N)}$ and θ
formulate this as a nonlinear least squares problem, and then give expression for the derivatives of the residuals

Example: fitting an ellipse to points in a plane

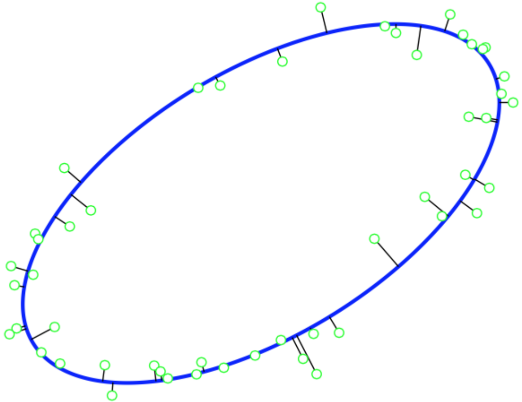

Outline

Matrices

Matrix inverse

orthogonal matrices

QR factorization

LU factorization

least squares

nonlinear least squares

Cholesky factorization

mathematical background

matrix norms

condition and stability

IEEE floating point numbers

Positive definite matrices

- a symmetric $n \times n$ matrix A is positive definite if

$$x^T Ax > 0 \quad \text{for all } x \neq 0$$

- every positive definite matrix is nonsingular
- every positive definite matrix has positive diagonal elements
- if the $n \times n$ matrix A is positive definite, then

$$B^T AB$$

  is positive definite for any $B \in \mathbf{R}^{n \times m}$ with linearly independent columns
- $A = B^T B$ is positive definite if B has linearly independent columns

Positive semidefinite matrices

- a symmetric $n \times n$ matrix A is positive semidefinite if

$$x^T Ax \geq 0 \quad \text{for all } x$$

- if A is positive semidefinite, but not positive definite, then it is singular
- every positive semidefinite matrix has nonnegative diagonal elements
- if the $n \times n$ matrix A is positive semidefinite, then

$$B^T AB$$

  is positive semidefinite for any $n \times m$ matrix B
- every Gram matrix $A = B^T B$ is positive semidefinite

Examples on positive definiteness

are the following matrices positive definite?

- $A = \begin{bmatrix} -1 & 2 & 3 \\ 2 & 5 & -3 \\ 3 & -3 & 2 \end{bmatrix}$
- $A = I - uu^T$ where u is an n-vector with $\|u\| < 1$
- $A = \begin{bmatrix} I & B \\ B^T & I + B^T B \end{bmatrix}$ where B is an $m \times n$ matrix

Cholesky factorization

every positive definite $n \times n$ matrix A can be factored as

$$A = R^T R$$

where $R \in \mathbf{R}^{n \times n}$ is upper triangular with positive diagonal elements

- complexity of computing R is $(1/3)n^3$ flops
- practical method for testing positive definiteness
- used in solving $Ax = b$ when A is positive definite

Cholesky factorization algorithm

$$\begin{bmatrix} A_{11} & A_{1,2:n} \\ A_{2:n,1} & A_{2:n,2:n} \end{bmatrix} = \begin{bmatrix} R_{11} & 0 \\ R_{1,2:n}^T & R_{2:n,2:n}^T \end{bmatrix} \begin{bmatrix} R_{11} & R_{1,2:n} \\ 0 & R_{2:n,2:n} \end{bmatrix}$$

$$= \begin{bmatrix} R_{11}^2 & R_{11}R_{1,2:n} \\ R_{11}R_{1,2:n}^T & R_{1,2:n}^T R_{1,2:n} + R_{2:n,2:n}^T R_{2:n,2:n} \end{bmatrix}$$

1. compute first row of R:

$$R_{11} = \sqrt{A_{11}}, \qquad R_{1,2:n} = \frac{1}{R_{11}} A_{1,2:n}$$

2. compute $2,2$ block $R_{2:n,2:n}$ from

$$A_{2:n,2:n} - R_{1,2:n}^T R_{1,2:n} = R_{2:n,2:n}^T R_{2:n,2:n}$$

   which is a Cholesky factorization of order $n-1$

Examples on Cholesky factorization

- simple exercises: A11.8
- block matrix example: A11.13

$$B = \begin{bmatrix} A & u \\ u^T & 1 \end{bmatrix}$$

- a more complicated example: A11.21

$$A = \begin{bmatrix} 1 & \mathbf{avg}(a) & \mathbf{avg}(b) \\ \mathbf{avg}(a) & \mathbf{rms}(a)^2 & (a^T n)/n \\ \mathbf{avg}(b) & (b^T a)/n & \mathbf{rms}(b)^2 \end{bmatrix} = \frac{1}{n} \begin{bmatrix} n & \mathbf{1}^T a & \mathbf{1}^T b \\ a^T \mathbf{1} & a^T a & a^T b \\ b^T \mathbf{1} & b^T a & b^T b \end{bmatrix}$$

- exploit structure: A is positive definite with negative off-diagonal entries
  1. show that its Cholesky factor R has negative above diagonal entries
  2. show that R^{-1} has positive above diagonal entries
  3. show that all entries of A^{-1} is positive

Outline

Mathematical background

- gradient of differentiable function $g\colon \mathbf{R}^n \to \mathbf{R}$

$$\nabla g(z) = \left(\frac{\partial g}{\partial x_1}(z), \ldots, \frac{\partial g}{\partial x_n}(z) \right) \in \mathbf{R}^n$$

- Hessian of g at z is a symmetric $n \times n$ matrix $\nabla^2 g(z)$ with entries

$$(\nabla^2 g(z))_{ij} = \frac{\partial^2 g}{\partial x_i \partial x_j}(z)$$

- composition with affine mapping: if $g(x) = h(Cx + d)$, then

$$\nabla g(x) = C^T \nabla h(Cx + d) \qquad \nabla^2 g(x) = C^T \nabla^2 h(Cx + d) C$$

Mathematical background

- affine approximation of g at z

$$\hat{g}(x) = g(z) + g(z)^T(x - z)$$

- quadratic approximation of g at z

$$\tilde{g}(x) = g(z) + \nabla g(z)^T(x - z) + \frac{1}{2}(x - z)^T \nabla^2 g(z)(x - z)$$

- Jacobian of differentiable function $f \colon \mathbf{R}^n \to \mathbf{R}^m$

$$Df(z) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1}(z) & \frac{\partial f_1}{\partial x_2}(z) & \cdots & \frac{\partial f_1}{\partial x_n}(z) \\ \frac{\partial f_2}{\partial x_1}(z) & \frac{\partial f_2}{\partial x_2}(z) & \cdots & \frac{\partial f_2}{\partial x_n}(z) \\ \vdots & \vdots & & \vdots \\ \frac{\partial f_m}{\partial x_1}(z) & \frac{\partial f_m}{\partial x_2}(z) & \cdots & \frac{\partial f_m}{\partial x_n}(z) \end{bmatrix} = \begin{bmatrix} \nabla f_1(z)^T \\ \vdots \\ \nabla f_m(z)^T \end{bmatrix}$$

Basic optimization theory

- local optimum and global optimum
- optimality conditions for twice differentiable function g
  - necessary: if $x^\star$ is locally optimal, then

$$\nabla g(x^\star) = 0 \quad \text{and} \quad \nabla^2 g(x^\star) \text{ is positive semidefinite}$$

  - sufficient: $x^\star$ is locally optimal only if

$$\nabla g(x^\star) = 0 \quad \text{and} \quad \nabla^2 g(x^\star) \text{ is positive definite}$$

  - if g is a convex function, then

$$x^\star \text{ is optimal} \quad \Longleftrightarrow \quad \nabla g(x^\star) = 0$$

Outline

Matrices

Matrix inverse

orthogonal matrices

QR factorization

LU factorization

least squares

nonlinear least squares

Cholesky factorization

mathematical background

matrix norms

condition and stability

IEEE floating point numbers

Properties of matrix norms

Properties satisfied by all matrix norms

- *nonnegative:* $\|A\|_2 \geq 0$ for all A
- *positive definiteness:* $\|A\|_2 = 0$ only if $A = 0$
- *homogeneity:* $\|\beta A\|_2 = |\beta|\|A\|_2$
- *triangle inequality:* $\|A + B\|_2 \leq \|A\|_2 + \|B\|_2$

Additional properties satisfied by the 2-norm $\|A\|_2 = \max_{x \neq 0}(\|Ax\|/\|x\|)$

- $\|Ax\| \leq \|A\|_2\|x\|$
- $\|AB\|_2 \leq \|A\|_2$
- if A is nonsingular, then $\|A\|_2\|A^{-1}\|_2 \geq 1$
- if A is nonsingular, then $1/\|A^{-1}\|_2 = \min_{x \neq 0}(\|Ax\|/\|x\|)$
- $\|A^T\|_2 = \|A\|$

Example on matrix norms

$A \in \mathbf{R}^{m \times n}$ has linearly independent columns and QR factorization $A = QR$

1. show that the norm of A satisfies

$$\|A\|_2 \geq \max\{R_{11}, R_{22}, \ldots, R_{nn}\}, \quad \|A^\dagger\|_2 \geq \frac{1}{\min\{R_{11}, R_{22}, \ldots, R_{nn}\}}$$

   (we follow the convention that $R_{ii} > 0$)

2. show that $\|AA^\dagger\|_2 = 1$ (even when $AA^\dagger \neq I$)

Example on matrix norms

1. if A is a square matrix with $\|I-A\|_2 < 1$. then A is nonsingular
2. if A is a nonsingular matrix, then

$$\|A^{-1}\|_2 \leq \|A^{-1}-I\|_2 + 1, \quad \|A^{-1}-I\|_2 \leq \|A^{-1}\|_2\|I-A\|_2$$

3. if A is a square matrix with $\|I-A\|_2 < 1$, then

$$\|A^{-1}\|_2 \leq \frac{1}{1-\|I-A\|_2}, \quad \kappa(A) \leq \frac{1+\|I-A\|_2}{1-\|I-A\|_2}$$

Outline

Condition and stability

Problem condition a mathematical problem is

- *well conditioned* if small changes in problem parameters (or problem data) lead to small changes in the solution;
- *ill-conditioned* if small changes in problem parameters (or problem data) can cause large changes in the solution

Cancellation occurs when

- we subtract two numbers that are almost equal;
- one or both numbers are subject to error

Numerical stability
refers to the accuracy of an *algorithm* in the presence of rounding errors

Outline

IEEE floating point numbers

Binary floating point numbers

$$x = \pm(.d_1 d_2 \ldots d_n)_2 \cdot 2^e$$

Machine precision $\epsilon_{\mathrm{M}} = 2^{-53} \approx 1.1102 \cdot 10^{-16}$

Rounding

- a floating point number system is a finite set of numbers
- all other numbers must be rounded

Rounding rules

- numbers are rounded to the nearest floating point number
- ties are resolved by rounding to the number with least significant bit 0 ("round to nearest even")

Example on IEEE floating point numbers

the figure shows the function

$$f(x) = \frac{(1+x)-1}{1+(x-1)}$$

evaluated in IEEE double precision arithmetic in the interval $[10^{-16}, 10^{-15}]$