Numerical Methods I: Orthogonalization and Newton's method

Georg Stadler
Courant Institute, NYU
stadler@cims.nyu.edu

October 5, 2017

Linear least squares and orthogonalization methods

Review: Least-squares problems

Given data points/measurements

$$(t_i, b_i), \quad i = 1, \ldots, m$$

and a model function ϕ that relates t and b:

$$b = \phi(t; x_1, \ldots, x_n),$$

where $x_1, \ldots, x_n$ are model function parameters. If the model is supposed to describe the data, the deviations/errors

$$\Delta_i = b_i - \phi(t_i, x_1, \ldots, x_n)$$

should be small. Thus, to fit the model to the measurements, one must choose $x_1, \ldots, x_n$ appropriately.

Review: Linear least-squares

We assume (for now) that the model depends linearly on $x_1, \dots, x_n$, e.g.:

$$\phi(t; x_1, \dots x_n) = a_1(t)x_1 + \ldots + a_n(t)x_n$$

Choosing the least square error, this results in

$$\min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{b}\|_2^2, \qquad \boxed{A}\,\boxed{x} - \boxed{b}$$

where $\boldsymbol{x} = (x_1, \dots, x_n)^T$, $\boldsymbol{b} = (b_1, \dots, b_m)^T$, and $a_{ij} = a_j(t_i)$.

In the following, we study the over-determined case, i.e., $m \geq n$.

Linear least-squares problems–QR factorization

Consider non-square matrices $A \in \mathbb{R}^{m\times n}$ with $m \geq n$ and rank(A) $= n$. Then the system

$$A\boldsymbol{x} = \boldsymbol{b}$$

does, in general, not have a solution (more equations than unknowns). We thus instead solve a minimization problem

$$\min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{b}\|^2.$$

The minimum $\bar{\boldsymbol{x}}$ of this optimization problem is characterized by the normal equations:

$$A^T A\bar{\boldsymbol{x}} = A^T\boldsymbol{b}.$$

Linear least-squares problems–QR factorization

To avoid the multiplication $A^T A$ and to use a suitable factorization of A that aids in solving the normal equation, we use the QR-factorization:

$$A = QR = [Q_1, Q_2] \begin{bmatrix} R_1 \\ 0 \end{bmatrix} = Q_1 R_1,$$

where $Q \in \mathbb{R}^{m \times m}$ is an orthonormal matrix ($QQ^T = I$), and $R \in \mathbb{R}^{m \times n}$ consists of an upper triangular matrix and a block of zeros.

Linear least-squares problems–QR factorization

How can the QR factorization be used to solve the normal equation?

$$\min_{\boldsymbol{x}} \|A\boldsymbol{x} - \boldsymbol{b}\|^2 = \min_{\boldsymbol{x}} \|Q^T(A\boldsymbol{x} - \boldsymbol{b})\|^2 = \min_{\boldsymbol{x}} \| \begin{bmatrix} \boldsymbol{b}_1 - R_1\boldsymbol{x} \\ \boldsymbol{b}_2 \end{bmatrix} \|^2,$$

where $Q^T\boldsymbol{b} = \begin{bmatrix} \boldsymbol{b}_1 \\ \boldsymbol{b}_2 \end{bmatrix}$.

Thus, the least squares solution is $\boldsymbol{x} = R^{-1}\boldsymbol{b}_1$ and the residual is $\|\boldsymbol{b}_2\|$.

Linear least-squares problems–QR factorization

How can we compute the QR factorization?

Givens rotations

Use sequence of rotations in 2D subspaces:
For $m \approx n$: $\sim n^2/2$ square roots, and $4/3n^3$ multiplications
For $m \gg n$: $\sim nm$ square roots, and $2mn^2$ multiplications

Householder reflections

Use sequence of reflections in 2D subspaces

For $m \approx n$: $2/3n^3$ multiplications
For $m \gg n$: $2mn^2$ multiplications

These methods compute an orthonormal basis of the columns of A. An alternative is the Gram Schmidt method—however, Gram Schmidt is unstable and thus sensitive to rounding errors (there are modified versions that are stable but require more computation).

QR factorization: Orthogonal transformations

Transform A by multiplication with orthogonal matrices

$$A \rightsquigarrow Q_1 A \rightsquigarrow Q_2 Q_1 A \rightsquigarrow \dots$$

$$\kappa_2(Q) = \|Q\|_2 \cdot \|Q^{-1}\|_2 = 1 \qquad \|Q\|_2 = \max_{\|x\|_2 = 1} \frac{\|Qx\|_2}{\|x\|_2} = 1$$

What basic orthogonal transformations exist in $\mathbb{R}^2$?

Rotations ($\det = 1$)

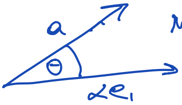

rotate: $\begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$

Reflections ($\det = -1$)

$$a \longmapsto a - 2\frac{\langle v, a\rangle}{\langle v, v\rangle} v$$

reflection at a hyperplane normal to v:

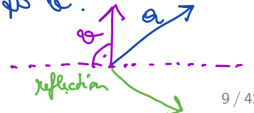

QR factorization: Givens rotations

Givens rotations:

$$\Omega_{k\ell} = \begin{array}{cc} \begin{array}{ccccccc} & & \downarrow k & & \downarrow \ell & & \end{array} & \\ \left[\begin{array}{ccccccc} 1 & & & & & & \\ & \ddots & & & & & \\ & & c & & s & & \\ & & & \ddots & & & \\ & & -s & & c & & \\ & & & & & \ddots & \\ & & & & & & 1 \end{array}\right] & \begin{array}{l} \\ \\ \leftarrow k \\ \\ \leftarrow \ell \\ \\ \\ \end{array} \end{array}$$

rotation in plane spanned by k^{th} & ℓ^{th} unit vector

In plane spanned by k^{th} & ℓ^{th} unit vector

$$\begin{pmatrix} c & s \\ -s & c \end{pmatrix} \begin{pmatrix} x_k \\ x_\ell \end{pmatrix} = \begin{pmatrix} r \\ 0 \end{pmatrix}$$

$\longrightarrow$ $r = \pm\sqrt{x_k^2 + x_\ell^2}$

$$s = \frac{x_\ell}{r}$$

$$c = \frac{x_k}{r}$$

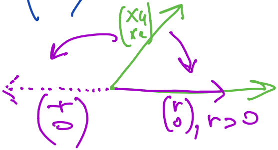

QR factorization: Givens rotations

$$A = \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix} \xrightarrow{\Omega_{34}} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \end{bmatrix} \xrightarrow{\Omega_{23}} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \end{bmatrix} \xrightarrow{\Omega_{12}} \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \end{bmatrix}$$

$$\xrightarrow{\Omega'_{34}} \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \end{bmatrix} \longrightarrow \quad \cdots \quad \longrightarrow \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} R_1 \\ 0 \end{bmatrix}$$

$$\underbrace{\Omega''_{34} \cdots \Omega_{23}\,\Omega_{34}}_{Q^T} A = R \implies A = QR$$

overall: $\frac{4}{3} n^3$ multiplications (and $\Theta(n^2)$ square roots)

For A in Hessenberg form, i.e.

$$A = \begin{bmatrix} \times & \times & \cdots & \times \\ \times & \times & & \\ 0 & \times & \times & \\ 0 & 0 & \times & \times \\ 0 & 0 & 0 & \times & \times \end{bmatrix} \longrightarrow \Theta(n) \text{ rotations}$$

QR factorization: Householder reflections

reflections $\qquad Q = I - 2\dfrac{vv^T}{v^Tv}$ $\qquad$ reflection on hyperplane $\perp$ to $v \in \mathbb{R}^n$

$$\begin{bmatrix} 1 & & \\ & \ddots & \\ & & 1 \end{bmatrix} - \frac{2}{v^Tv} \begin{bmatrix} v \end{bmatrix} \begin{bmatrix} \; v^T \; \end{bmatrix}$$

$$Qy = re_1 = \begin{pmatrix} r \\ 0 \\ \vdots \\ 0 \end{pmatrix} \qquad y \in \mathbb{R}^n$$

for $v = y - \alpha e_1$

$$\begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix} \xrightarrow{Q_1} \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \end{bmatrix} \xrightarrow{Q_2} \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \\ 0 & 0 & \times \end{bmatrix} \xrightarrow{Q_3} \begin{bmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \\ 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} R_1 \\ 0 \end{bmatrix}$$

QR factorization: Householder reflections

If $A \in \mathbb{R}^{n \times n}$ square matrix, $\det(A) \neq 0$

$A = LU$ factorization $\longrightarrow$ Solve $Ax = b$ by

(cost $\sim \frac{1}{3}n^3$)

| backward &

| forward subst.

$A = QR =$ 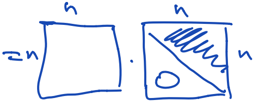 $\longrightarrow$ Solve $Ax = b$ by

$Rx = Q^T b$

| backward subst.

(cost $\sim \frac{4}{3}n^3$ Givens

$\sim \frac{2}{3}n^3$ Householder)

more stable as it only uses orthogonal transformation, but it's more expensive

Nonlinear systems

Fixed point ideas

We intend to solve the nonlinear equation $f : \mathbb{R} \to \mathbb{R}$

$$f(x) = 0, \quad x \in \mathbb{R}.$$

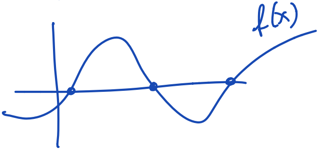

Reformulation as fixed point method:

$$x = \Phi(x)$$

Corresponding iteration: Choose x_0 (initialization) and compute $x_1, x_2, \ldots$ from

$$x_{k+1} = \Phi(x_k)$$

When does this iteration converge?

Fixed point ideas

Example: Solve the nonlinear equation

$$2x - \tan(x) = 0.$$

Iteration #1: $x_{k+1} = \Phi_1(x_k) = 0.5\tan(x_k)$

Iteration #2: $x_{k+1} = \Phi_2(x_k) = \arctan(2x_k)$

Iteration #3: $x_{k+1} = \Phi_3(x_k) = x_k - \frac{2x_k - \tan(x_k)}{1-\tan^2(x_k)}$

Convergence of fixed point methods

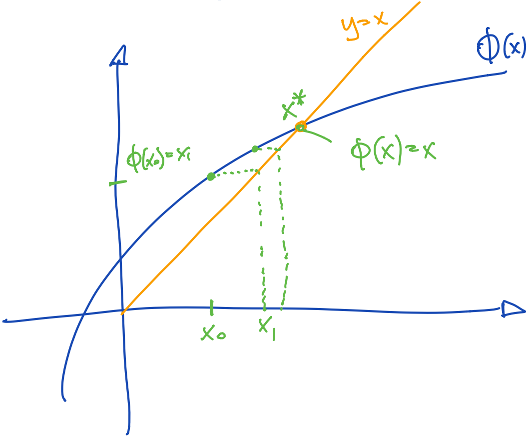

Convergence of fixed point methods

A mapping $\Phi : [a, b] \to \mathbb{R}$ is called contractive on $[a, b]$ if there is a $0 \leq \Theta < 1$ such that

$$|\Phi(x) - \Phi(y)| \leq \Theta |x - y| \text{ for all } x, y \in [a, b].$$

If Φ is continuously differentiable on $[a, b]$, then

$$\sup_{x,y \in [a,b]} \frac{|\Phi(x) - \Phi(y)|}{|x - y|} = \sup_{z \in [a,b]} |\Phi'(z)|$$

Convergence of fixed point methods

Let $\Phi : [a,b] \to [a,b]$ be contractive with constant $\Theta < 1$. Then:

- There exists a unique fixed point $\bar{x}$ with $\bar{x} = \Phi(\bar{x})$
- For any starting guess x_0 in $[a,b]$, the fixed point iteration converges to $\bar{x}$ and

$$|x_{k+1} - x_k| \leq \Theta |x_k - x_{k-1}| \quad \text{(linear convergence)}$$

$$|\bar{x} - x_k| \leq \frac{\Theta^k}{1-\Theta} |x_1 - x_0|.$$

The second expression allows to estimate the required number of iterations.

Convergence of fixed point methods

Proof: $|x_{k+1} - x_k| = |\phi(x_k) - \phi(x_{k-1})| \leq \Theta |x_k - x_{k-1}|$

$\Longrightarrow |x_{k+1} - x_k| \leq \Theta^k |x_1 - x_0|$

$$|x_{k+m} - x_k| \leq |x_{k+m} - x_{k+m-1}| + |x_{k+m-1} - x_{k+m-2}| + \dots + |x_{k+1} - x_k|$$

$$\leq \underbrace{\left(\Theta^{k+m-1} + \Theta^{k+m-2} + \dots + \Theta^k\right)}_{\Theta^k (1 + \Theta + \dots + \Theta^{m-1})} |x_1 - x_0|$$

$$\leq \frac{\Theta^k}{1-\Theta} |x_1 - x_0| \Longrightarrow x_k \longrightarrow x^* \text{ converges}$$

- x^* is a fixed point:

$$|x^* - \phi(x^*)| \leq |x^* - x_{k+1}| + |\underbrace{x_{k+1}}_{=\phi(x_k)} - \phi(x^*)| \leq \underbrace{|x^* - x_{k+1}| + \Theta |x_k - x^*|}_{\xrightarrow{k \to \infty} 0}$$

$\longrightarrow x^*$ is a fixed point.

uniquiness: x^*, y^* fixed points: $|x^* - y^*| = |\phi(x^*) - \phi(y^*)|$

$$\leq \Theta |x^* - y^*| \Longrightarrow x^* = y^* \quad \square$$

Newton's method

In one dimension, solve $f(x) = 0$:

Start with x_0, and compute $x_1, x_2, \dots$ from

$$x_{k+1} = x_k - \frac{f(x_k)}{f'(x_k)}, \quad k = 0, 1, \dots$$

Requires $f'(x_k) \neq 0$ to be well-defined (i.e., tangent has nonzero slope).

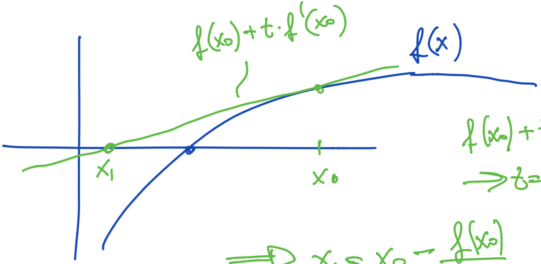

$f(x_0) + t f'(x_0) = 0$

$\Rightarrow t = \dfrac{f(x_0)}{f'(x_0)}$

$\Longrightarrow x_1 = x_0 - \dfrac{f(x_0)}{f'(x_0)}$

Newton's method

$$F = \begin{bmatrix} F_1 \\ F_2 \\ \vdots \end{bmatrix}$$

$$F'(x_0) = \begin{bmatrix} \frac{\partial F_1}{\partial x_1} & \cdots \\ \vdots & \\ \frac{\partial F_1}{\partial x_n} & \cdots \end{bmatrix} \in \mathbb{R}^{n \times n}$$

Let $F : \mathbb{R}^n \to \mathbb{R}^n$, $n \geq 1$ and solve

$$F(\boldsymbol{x}) = 0.$$

Taylor expansion about starting point $\boldsymbol{x}^0$:

$$F(\boldsymbol{x}) = F(\boldsymbol{x}^0) + F'(\boldsymbol{x}^0)(\boldsymbol{x} - \boldsymbol{x}^0) + o(|\boldsymbol{x} - \boldsymbol{x}^0|) \quad \text{for } \boldsymbol{x} \to \boldsymbol{x}^0.$$

Hence:

$$\boldsymbol{x}^1 = \boldsymbol{x}^0 - F'(\boldsymbol{x}^0)^{-1} F(\boldsymbol{x}^0)$$

Newton iteration: Start with $\boldsymbol{x}^0 \in \mathbb{R}^n$, and for $k = 0, 1, \ldots$ compute

$$F'(\boldsymbol{x}^k)\Delta\boldsymbol{x}^k = -F(\boldsymbol{x}^k), \quad \boldsymbol{x}^{k+1} = \boldsymbol{x}^k + \Delta\boldsymbol{x}^k$$

Requires that $F'(\boldsymbol{x}^k) \in \mathbb{R}^{n \times n}$ is invertible.

Newton's method

Newton iteration: Start with $\boldsymbol{x}^0 \in \mathbb{R}^n$, and for $k = 0, 1, \ldots$ compute

$$F'(\boldsymbol{x}^k)\Delta\boldsymbol{x}^k = -F(\boldsymbol{x}^k), \quad \boldsymbol{x}^{k+1} = \boldsymbol{x}^k + \Delta\boldsymbol{x}^k$$

Equivalently:

$$\boldsymbol{x}^{k+1} = \boldsymbol{x}^k - F'(\boldsymbol{x}^k)^{-1}F(\boldsymbol{x}^k)$$

Newton's method is affine invariant, that is, the sequence is invariant to affine transformations:

Convergence of Newton's method

Assumptions on F: $D \subset \mathbb{R}^n$ open and convex, $F : D \to \mathbb{R}^n$ continuously differentiable with $F'(\boldsymbol{x})$ is invertible for all $\boldsymbol{x}$, and there exists $\omega \geq 0$ such that

$$\|F'(\boldsymbol{x})^{-1}(F'(\boldsymbol{x}+s\boldsymbol{v}) - F'(\boldsymbol{x}))\boldsymbol{v}\| \leq s\omega\|\boldsymbol{v}\|^2$$

for all $s \in [0,1]$, $\boldsymbol{x} \in D, \boldsymbol{v} \in \mathbb{R}^n$ with $\boldsymbol{x} + \boldsymbol{v} \in D$.
Assumptions on $\boldsymbol{x}^*$ and $\boldsymbol{x}^0$: There exists a solution $\boldsymbol{x}^* \in D$ and a starting point $\boldsymbol{x}^0 \in D$ such that

$$\rho := \|\boldsymbol{x}^* - \boldsymbol{x}^0\| \leq \frac{2}{\omega} \text{ and } B_\rho(\boldsymbol{x}^*) \subset D$$

Theorem: Then, the Newton sequence $\boldsymbol{x}^k$ stays in $B_\rho(\boldsymbol{x}^*)$ and $\lim_{k\to\infty} \boldsymbol{x}^k = \boldsymbol{x}^*$, and

$$\|\boldsymbol{x}^{k+1} - \boldsymbol{x}^*\| \leq \frac{\omega}{2}\|\boldsymbol{x}^k - \boldsymbol{x}^*\|^2$$