Numerical Methods for Computational Science and Engineering

Autumn Semester 2018

Prof. Rima Alaifari, SAM, ETH Zurich

Fixed Point Iteration in $\mathbb{R}^n$

Definition: A fixed point iteration

$$x^{(k+1)} = \Phi(x^{(k)})$$

is consistent with $F(x)=0$ if for $x \in U \cap D$

$$F(x)=0 \iff \Phi(x) = x.$$

Definition [Contractive mapping]

$\Phi : U \subset \mathbb{R}^n \to \mathbb{R}^n$ is contractive (w.r.t. norm $\|\cdot\|$ on $\mathbb{R}^n$) if

$$\exists L < 1 : \quad \|\Phi(x) - \Phi(y)\| \le L\|x-y\| \qquad \forall x,y \in U.$$

Implication of contractivity:

① If there exists x^* with $\Phi(x^*) = x^*$ then FPI will converge to x^*

$$\|x^{(k+1)} - x^*\| = \|\Phi(x^{(k)}) - \Phi(x^*)\|$$
$$\le L \cdot \|x^{(k)} - x^*\|$$

$$\|e^{(k+1)}\| \le \underbrace{L^k}_{\to 0} \cdot \|e^{(0)}\| \qquad L<1$$

② At least linear convergence.

③ Φ has at most one FP.

Why? Suppose we have 2 FPs: x_1^*, x_2^*

$$\|x_1^* - x_2^*\| = \|\Phi(x_1^*) - \Phi(x_2^*)\| \leq \underbrace{L}_{<1} \cdot \|x_1^* - x_2^*\|$$

$$\Rightarrow \|x_1^* - x_2^*\| = 0 \Rightarrow x_1^* = x_2^*$$

Existence of a FP:

Theorem 8.2.1 (Banach's Fixed Point Theorem). *If $D \subset \mathbb{K}^n$ ($\mathbb{K} = \mathbb{R}, \mathbb{C}$) closed and bounded and $\Phi : D \to D$ satisfies*

$$\exists L < 1 \text{ such that } \forall x, y \in D : \qquad \|\Phi(x) - \Phi(y)\| \leq L \cdot \|x - y\|,$$

then there is a unique fixed point $\mathbf{x}^* \in D$, $\Phi(\mathbf{x}^*) = \mathbf{x}^*$, which is the limit of the sequence of iterates $\mathbf{x}^{(k+1)} := \Phi(x^{(k)})$, for any $\mathbf{x}^{(0)} \in D$.

Local statement for differentiable Φ:

Lemma 8.2.1 (Sufficient Condition for Local Linear Convergence of FPI). *If $\Phi : U \subset \mathbb{R}^n \to \mathbb{R}^n$, $\Phi(\mathbf{x}^*) = \mathbf{x}^*$, Φ differentiable in $\mathbf{x}^*$, and $\|D\,\Phi(\mathbf{x}^*)\| < 1$, then the fixed point iteration*

$$\mathbf{x}^{(k+1)} := \Phi(\mathbf{x}^{(k)})$$

converges locally and at least linearly.

local contractivity

$$D\Phi(x) := \left[\frac{\partial \Phi_j}{\partial x_i}(x)\right]_{j,i=1}^{n} \in \mathbb{R}^{n,n}$$

Jacobian

Lemma 8.2.2 (Sufficient Condition for Linear Convergence of FPI). *Let U be convex and $\Phi : U \subset \mathbb{R}^n \to \mathbb{R}^n$ be continuously differentiable with*

$$L := \sup_{\mathbf{x} \in U} \|D\Phi(\mathbf{x})\| < 1 .$$

If $\Phi(\mathbf{x}^*) = \mathbf{x}^*$ for some interior point $\mathbf{x}^* \in U$, then the fixed point iteration $\mathbf{x}^{(k+1)} = \Phi(\mathbf{x}^{(k)})$ converges to $\mathbf{x}^*$ at least linearly with rate L.

- If Φ is locally contractive $\Rightarrow$ iteration converges locally around the FP
- At least linear convergence

When to terminate the iteration?

Find x^* with $F(x^*)=0$

When to stop iterating?

Ideally: $\|x^{(k)} - x^*\| \leq \tau$

Since x^* is not known $\rightsquigarrow$ not verifiable

Different criterion?

(1) residual based: stopping criterion $\|F(x^{(k)})\| \leq \tau$

(2) correction based: — " — $\|x^{(k+1)} - x^{(k)}\| \leq \tau$

or $\|x^{(k+1)} - x^{(k)}\| \leq \tau \|x^{(k+1)}\|$

Recall discussion about condition numbers:

$$\underbrace{\|F(x^{(k)}) - \underbrace{F(x^*)}_{=0}\|}_{\text{computable}} \text{ small } \not\Rightarrow \underbrace{\|x^{(k)} - x^*\|}_{\text{not computable}} \text{ small}$$

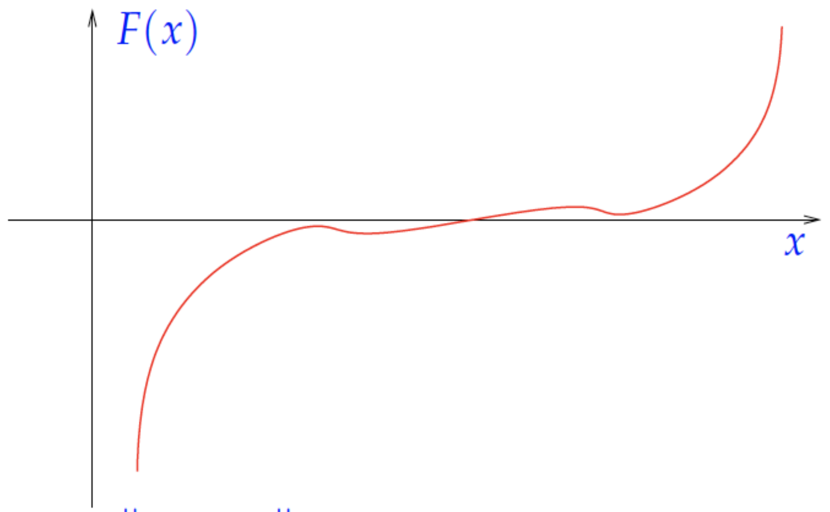

$\left\|F(\mathbf{x}^{(k)})\right\|$ small $\not\Rightarrow$ $|x - x^*|$ small

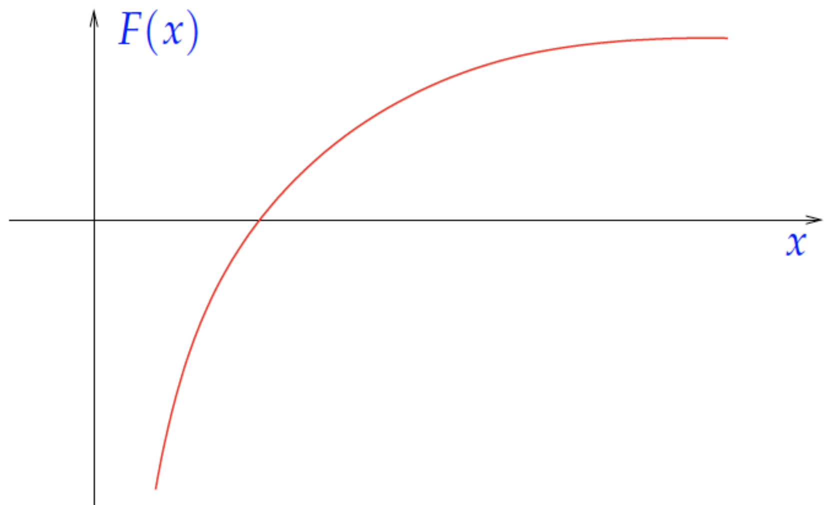

$\left\|F(\mathbf{x}^{(k)})\right\|$ small $\Rightarrow$ $|x - x^*|$ small

If iteration is linearly convergent:

$$\|x^{(k)} - x^*\| \leq \|x^{(k+1)} - x^{(k)}\| + \|x^{(k+1)} - x^*\|$$

$$\leq \|x^{(k+1)} - x^{(k)}\| + L \cdot \|x^{(k)} - x^*\|$$

$$\Rightarrow \quad (1-L) \cdot \|x^{(k)} - x^*\| \leq \|x^{(k+1)} - x^{(k)}\| \qquad (*)$$

$$\Rightarrow \quad \|x^{(k+1)} - x^*\| \leq L \cdot \|x^{(k)} - x^*\|$$

$$\overset{(*)}{\leq} \frac{L}{1-L} \|x^{(k+1)} - x^{(k)}\|$$

Idea: stopping criterion:

$$\frac{L}{1-L} \|x^{(k+1)} - x^{(k)}\| \leq \tau$$

This guarantees: $\|x^{(k+1)} - x^*\| \leq \tau$

Estimating L can be difficult

A pessimistic estimate $\tilde{L} > L$ will also do.

More generally:

A-priori termination criterion:

$$\|x^{(k)} - x^*\| \leq \frac{L^k}{1-L} \|x^{(1)} - x^{(0)}\|$$

A-posteriori termination criterion:

$$\|x^{(k)} - x^*\| \leq \frac{L}{1-L} \|x^{(k)} - x^{(k-1)}\|$$

8.4 Newton's method (in higher dimensions)

Extend idea from 1D:

First order approximation (linearization):

$$F(x) \approx F(x^{(k)}) + \underbrace{DF(x^{(k)})}_{\in \mathbb{R}^{n,n}} (x - x^{(k)})$$

Jacobian of F at $x^{(k)}$

Definition 8.3.1 (Newton's Method in $\mathbb{R}^n$). *Newton's iteration* in $\mathbb{R}^n$ may be defined by the recursive rule

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \mathrm{D}\,F(\mathbf{x}^{(k)})^{-1} F(\mathbf{x}^{(k)}), \tag{8.8}$$

if $\mathrm{D}\,F(\mathbf{x}^{(k)})$ is regular. We call $-\mathrm{D}\,F(\mathbf{x}^{(k)})^{-1} F(\mathbf{x}^{(k)})$ the *Newton correction* term.

[Recall 1D analogue: $x^{(k+1)} = x^{(k)} - \dfrac{f(x^{(k)})}{f'(x^{(k)})}$]

In 1D: $f'(x^{(k)}) \neq 0$

Now: Invertibility of $DF(x^{(k)})$

At each iteration: compute Newton correction

solve $DF(x^{(k)})\, y = -F(x^{(k)})$

y is then the Newton correction

Convergence of Newton's method:

If $F(x^*) = 0$ and $DF(x^*)$ is regular

then it is locally quadratically convergent.

Exact theorem: Theorem in lecture notes that states this rigorously $\rightsquigarrow$ hardly ever possible to verify in practice

Note: If $DF(x^*)$ is singular, no longer quadratic convergence

A few more considerations:

1.) Stopping criterion for Newton's method?

2.) Possibly larger region of convergence?

$\rightsquigarrow$ always at the cost of loosing quadratic convergence

3.) Reducing computational cost?

[Now: solving LSE in each iteration]

Ad 1.) Expect quadr. convergence

$$\|x^{(k+1)} - x^*\| \ll \|x^{(k)} - x^*\|$$

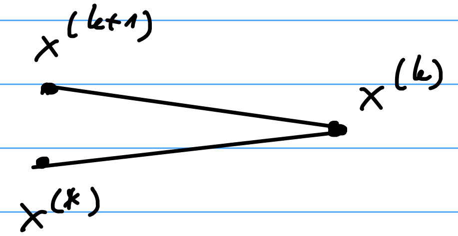

Roughly: $\|x^{(k)} - x^*\| \overset{(*)}{\approx} \|x^{(k+1)} - x^{(k)}\|$

$$\|x^{(k+1)} - x^{(k)}\| = \underbrace{\|DF(x^{(k)})^{-1} F(x^{(k)})\| \leq \tau \|x^{(k)}\|}_{\text{computable stopping criterion}}$$

roughly implies: $\|x^{(k)} - x^*\| \underset{(*)}{\lesssim} \tau \|x^{(k)}\|$

If $x^{(k)}$ was a good approximation, we would have solved one LSE too many for verifying the stopping criterion

In practice: "cheaper" stopping criterion

$$\underbrace{\| DF(x^{(k-1)})^{-1} F(x^{(k)}) \|}_{\text{simplified Newton correction}} \leq \tau \| x^{(k)} \|$$

Why: • We already have LU factorization of $DF(x^{(k-1)})$

• Due to quadratic conv.

$$DF(x^{(k)}) \approx DF(x^{(k-1)})$$

Failure of Newton's method:

① Local min./max.

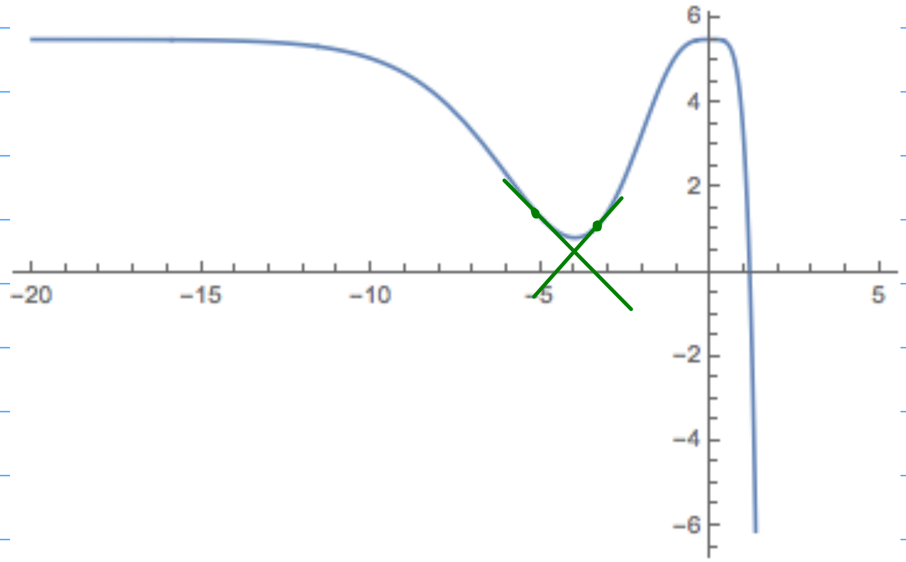

② Asymptotes

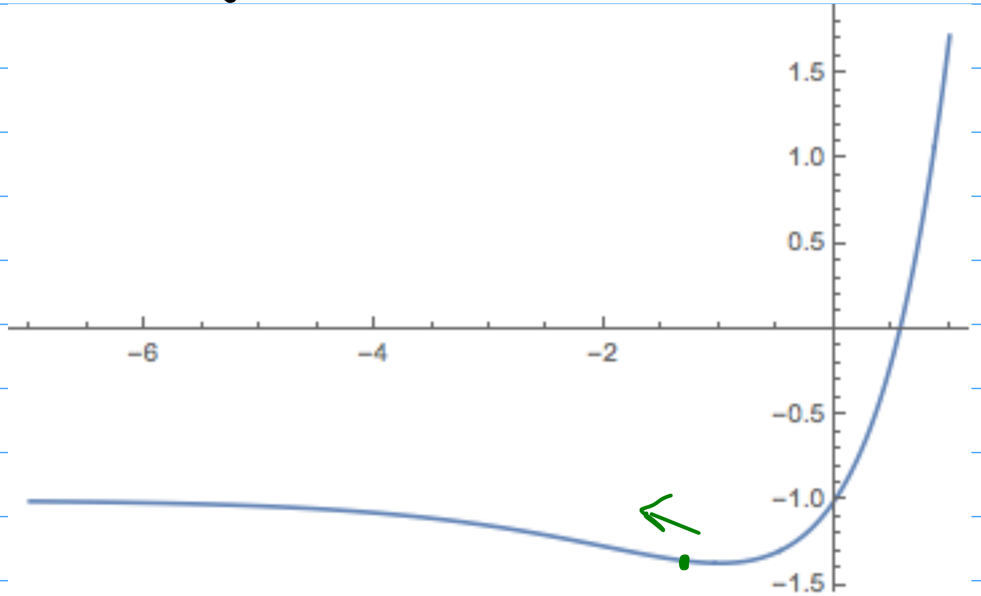

$f(x) = xe^x - 1$

$f'(x) = (1+x)e^x$

$f'(-1) = 0$

$x^{(0)} < -1 \Rightarrow x^{(k)} \to -\infty$ divergence

$x^{(0)} > -1 \Rightarrow x^{(k)} \to x^*$ convergence

③ Overshooting:

$f(x) = \arctan x \qquad f(0) = 0$

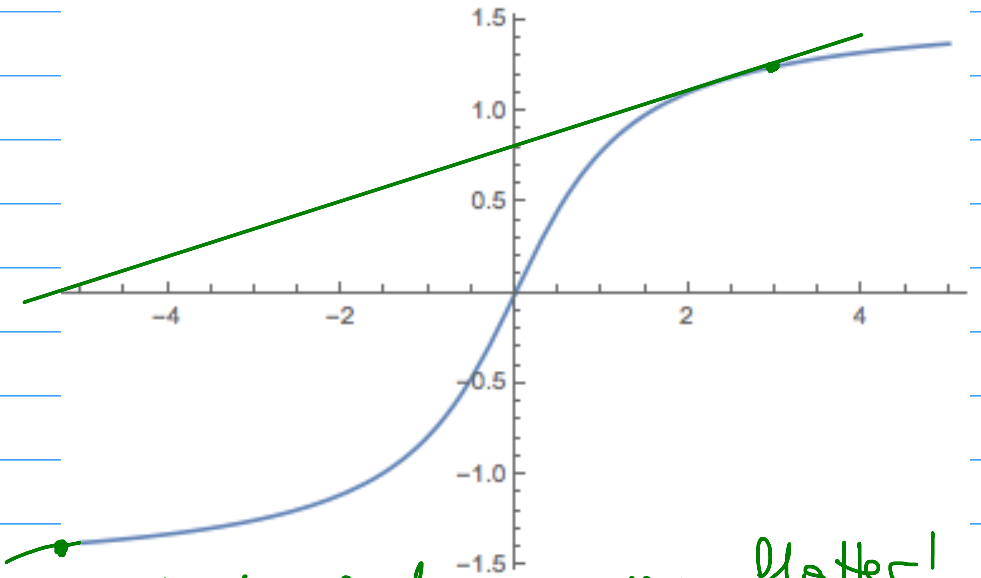

↑ this tangent is even flatter!

Ad 2.): A remedy for overshooting

Damped Newton method

Idea: In each iteration step, check whether the distance $\|x^{(k+1)} - x^{(k)}\|$ is decreasing.

E.g. $\|x^{(k+2)} - x^{(k+1)}\| \le \frac{1}{2} \|x^{(k+1)} - x^{(k)}\|$

If this is violated (see Example ③):

⇝ don't take a full Newton step

⇝ damping

$$x^{(k+1)} = x^{(k)} - \lambda^{(k)} DF(x^{(k)})^{-1} F(x^{(k)})$$

$\lambda^{(k)} \in (0, 1]$ damping factor

How to choose $\lambda^{(k)}$?

Strategy: choose the largest possible $\lambda^{(k)}$ so that distance between iterates decreases

Choose the maximal $0 < \lambda^{(k)} \leq 1$ such that

$$\left\| \Delta\overline{\mathbf{x}}(\lambda^{(k)}) \right\| \leq \left(1 - \frac{\lambda^{(k)}}{2}\right) \cdot \left\| \Delta\mathbf{x}^{(k)} \right\|, \qquad (*)$$

where $\Delta\mathbf{x}^{(k)} := \mathrm{D}\,F(\mathbf{x}^{(k)})^{-1}F(\mathbf{x}^{(k)})$ denotes the current Newton correction and

$$\Delta\overline{\mathbf{x}}(\lambda^{(k)}) := \mathrm{D}\,F(\mathbf{x}^{(k)})^{-1}F(\underbrace{\mathbf{x}^{(k)} - \lambda^{(k)}\Delta\mathbf{x}^{(k)}}_{\tilde{x}^{(k+1)}})$$

is a tentative simplified Newton correction.

$$\Delta x^{(k)} = DF(x^{(k)})^{-1} F(x^{(k)})$$

$$x^{(k+1)} = x^{(k)} - DF(x^{(k)})^{-1} F(x^{(k)})$$

$$\Delta x^{(k)} = x^{(k)} - x^{(k+1)}$$

$$\Delta \overline{x}(\lambda^{(k)}) \approx x^{(k+2)} - \tilde{x}^{(k+1)}$$

In practice: $\lambda^{(k)} = 1$ [no damping]

and check $(*)$,

if $(*)$ doesn't hold $\lambda^{(k)} \leftarrow \frac{\lambda^{(k)}}{2}$

until $(*)$ fulfilled for the first time

Example 1: $F(x) = \arctan x$, $x^{(0)} = 20$

k	$\lambda^{(k)}$	$x^{(k)}$	$F(x^{(k)})$
1	0.03125	0.94199967624205	0.75554074974604
2	0.06250	0.85287592931991	0.70616132170387
3	0.12500	0.70039827977515	0.61099321623952
4	0.25000	0.47271811131169	0.44158487422833
5	0.50000	0.20258686348037	0.19988168667351
6	1.00000	-0.00549825489514	-0.00549819949059
7	1.00000	0.00000011081045	0.00000011081045
8	1.00000	-0.00000000000001	-0.00000000000001

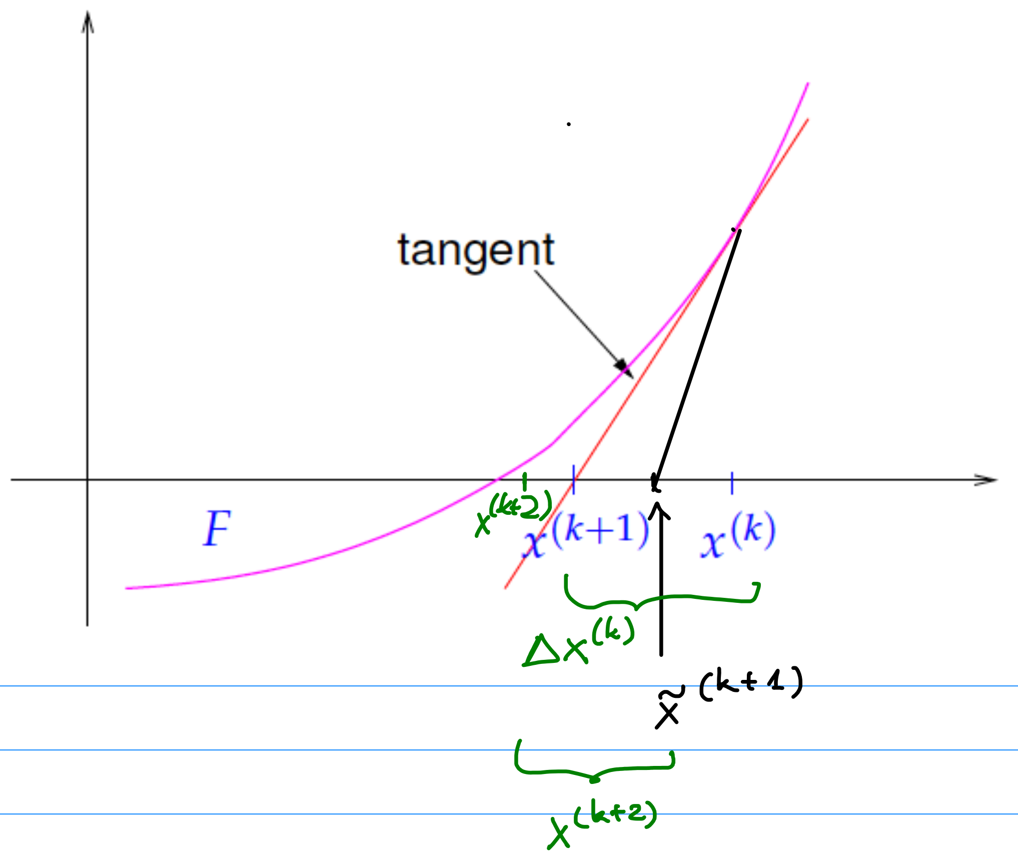

$\lambda^{(k)} = \frac{1}{2}$

Example 2: $F(x) = x e^{x} - 1$, $x^{(0)} = -1.5$

k	$\lambda^{(k)}$	$x^{(k)}$	$F(x^{(k)})$
1	0.25000	-4.4908445351690	-1.0503476286303
2	0.06250	-6.1682249558799	-1.0129221310944
3	0.01562	-7.6300006580712	-1.0037055902301
4	0.00390	-8.8476436930246	-1.0012715832278
5	0.00195	-10.5815494437311	-1.0002685596314

$\lambda_{min} = 0.001$

Ad 3.) Cheaper (approximate) Newton corrections?

Secant method in 1D:

$$f'(x^{(k)}) \approx \frac{f(x^{(k)}) - f(x^{(k-1)})}{x^{(k)} - x^{(k-1)}}$$

dimension $n > 1$?

Approximation of $DF(x^{(k)})$: $J_k \in \mathbb{R}^{n,n}$

$$J_k \left(x^{(k)} - x^{(k-1)} \right) = F(x^{(k)}) - F(x^{(k-1)}) \qquad (*)$$

Modified Newton's method:

$$x^{(k)} = x^{(k-1)} - J_{k-1}^{-1} F(x^{(k-1)})$$

$$\Leftrightarrow \quad J_{k-1}\left(x^{(k)}-x^{(k-1)}\right) = -F\left(x^{(k-1)}\right) \quad (**)$$

$(*)-(**)$:

$$\underbrace{\left(J_k - J_{k-1}\right)\left(x^{(k)}-x^{(k-1)}\right) = F\left(x^{(k)}\right)}_{\text{underdetermined}}$$

$\rightsquigarrow$ opt for a cheap choice:

$$J_k - J_{k-1} = \underbrace{\frac{F(x^{(k)})\left(x^{(k)}-x^{(k-1)}\right)^T}{\|x^{(k)}-x^{(k-1)}\|_2^2}}_{\text{!!! rank-1 matrix !!!}}$$

Start with $J_0 = DF(x^{(0)})$

get J_1 by a rank-1 update

$\rightsquigarrow$ keep going iteratively:

$$J_k = J_{k-1} + \frac{F(x^{(k)})\left(x^{(k)}-x^{(k-1)}\right)^T}{\|x^{(k)}-x^{(k-1)}\|_2^2}$$

Broyden's quasi-Newton method:

$$\mathbf{x}^{(k+1)} := \mathbf{x}^{(k)} + \Delta\mathbf{x}^{(k)}, \quad \Delta\mathbf{x}^{(k)} := -\mathbf{J}_k^{-1}F(\mathbf{x}^{(k)}),$$

$$\mathbf{J}_{k+1} := \mathbf{J}_k + \frac{F(\mathbf{x}^{(k+1)})(\Delta\mathbf{x}^{(k)})^\top}{\left\|\Delta\mathbf{x}^{(k)}\right\|_2^2}.$$

We can calculate J_k^{-1} from J_{k-1}^{-1} through Sherman-Morrison-Woodbury formula

Remark: In general, iterative methods for nonlinear systems should have a convergence monitor: check at each iteration whether convergence is to be expected or not

Example of NMT [damped Newton]:

if repeated failure: stop & report error

8.5 Unconstrained Optimization

Optimization problems we have already seen:

- Least-squares solution:

  Find $x \in \mathbb{K}^n$ s.t. $\|Ax-b\|_2 \to \min$

- Generalized solution:

  Find least-sq. solution x to $Ax=b$ s.t.

  $\|x\|_2 \to \min$

- Best low-rank approximation:

  Given $A \in \mathbb{K}^{m,n}$, find $\tilde{A} \in \mathbb{K}^{m,n}$ with $\operatorname{rank}(\tilde{A}) \leq k$

  s.t. $\|A - \tilde{A}\|_{2/F} \to \min$ over rank-k matrices

General problem formulation:

$$F: \mathbb{R}^n \to \mathbb{R}$$

Find a max/min of F!

Example from machine learning:

Maximum likelihood estimation

Suppose some quantity can be modeled by a probability distribution

For example: Height of fir trees of a certain age

$\rightsquigarrow$ normal distribution

Can we estimate mean μ & variance σ through randomized samples?

Samples $\{h_1, \dots, h_n\}$

$$f(h; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\,\sigma} e^{-(h-\mu)^2/2\sigma^2}$$

$f(h_i; \mu, \sigma)$ likelihood to observe height h_i

height of tree i is independent of height of tree j

$$\underbrace{\mathbb{P}(\underbrace{\{h_1, \dots, h_n\}}_{\text{fixed}}; \mu, \sigma)}_{\text{function in } \mu \,\&\, \sigma} = \prod_{i=1}^{n} f(h_i; \mu, \sigma)$$

$\rightsquigarrow$ Maximize $\mathbb{P}$ to estimate μ & σ

In practice: maximize $\log \mathbb{P}$ instead

[location of max is the same, but preferred due to numerical aspects]

Remark: maximizing $F \Leftrightarrow$ minimizing $-F$

$\rightsquigarrow$ we will only consider min. problems

Global vs. local minimum:

- x^* is a global minimum of $F: \mathbb{R}^n \to \mathbb{R}$ if
$$F(x^*) \leq F(x) \quad \forall\, x \in \mathbb{R}^n$$

- x^* is a local minimum of $F: \mathbb{R}^n \to \mathbb{R}$ if $\exists\, \varepsilon > 0$ s.t. $\forall\, x$ with $\underbrace{\|x - x^*\| \leq \varepsilon}_{\varepsilon\text{-ball around } x^*}$
$$F(x^*) \leq F(x)$$

Optimization with differentiable objective function

$F: \mathbb{R}^n \to \mathbb{R}$ differentiable

direction ∇F : direction of greatest increase

$-\nabla F$: direction of steepest descent

Why? Locally around $\bar{x}$: $F(x) \approx F(\bar{x}) + \nabla F(\bar{x})^T (x - \bar{x})$

If we choose $x = \bar{x} + \tau \nabla F(\bar{x})$

$$F(x) = F(\bar{x} + \tau \nabla F(\bar{x})) \approx F(\bar{x}) + \tau \|\nabla F(\bar{x})\|^2$$

If $\tau > 0$: function value increases

$\tau < 0$: —"— decreases

Stationary point : $\nabla F(x) = 0$

$\rightsquigarrow$ could be local/global max/min, saddle point

If F is twice diff. $\rightsquigarrow$ consider the Hessian matrix at the stationary point:

$$H_F(x) = \left(\frac{\partial^2 F}{\partial x_i \partial x_j}(x) \right)_{i,j=1}^{n}$$

Taylor expansion:

around stationary point x^*

$$F(x) \approx F(x^*) + \underbrace{\nabla F(x^*)^T (x-x^*)}_{=0} + \frac{1}{2}(x-x^*)^T H_F(x^*)(x-x^*)$$

$$F(x) \approx F(x^*) + \underbrace{\frac{1}{2}(x-x^*)^T H_F(x^*)(x-x^*)}_{(*)}$$

increase / decrease / not clear

$H_F(x^*)$ pos. def.: $\quad (*) > 0$

$\rightsquigarrow$ locally: increase of function values around x^*

$$F(x) \geq F(x^*)$$

$\Rightarrow$ local minimum

$H_F(x^*)$ neg. def. $\Rightarrow$ local maximum

indefinite $\Rightarrow$ saddle point

$H_F(x^*)$ not invertible $\rightsquigarrow$ whole region of saddle points

Check positive definiteness: by checking whether Cholesky factorization exists

Optimization with convex objective function

Definition [convex function]:

A function $F: \mathbb{R}^n \to \mathbb{R}$ is called convex if for all $x, y \in \mathbb{R}^n$ and all $\alpha \in (0,1)$

$$F((1-\alpha)x + \alpha y) \leq (1-\alpha)F(x) + \alpha F(y)$$

$$(<) \qquad \text{(strictly convex)}$$

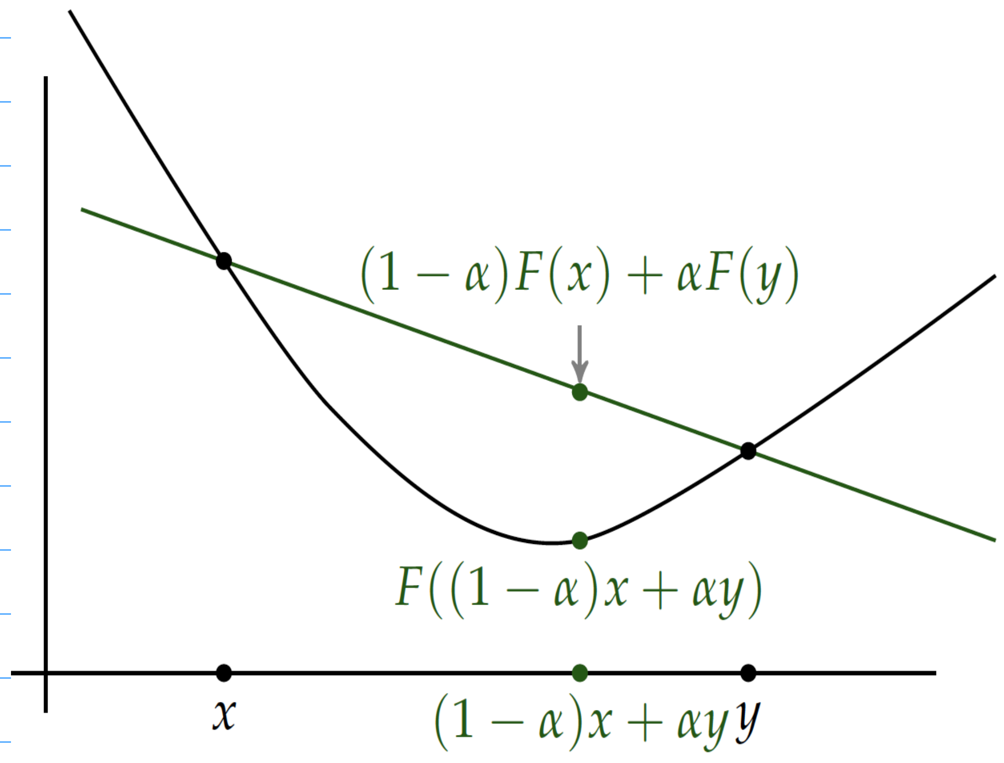

Lemma [minimum of convex function]:

If x^* is a local minimum of F and F is convex, then x^* is a global minimum.

Derive:

Suppose F is convex, x^* is a local minimum but not a global minimum

$\Rightarrow \exists\ x_0 \in \mathbb{R}^n$ s.t. $F(x_0) < F(x^*)$

Convexity: for $\alpha \in (0,1)$:

$$F(\underbrace{x^* + \alpha(x_0 - x^*)}_{=\alpha x_0 + (1-\alpha)x^*}) \leq \alpha \underbrace{F(x_0)}_{<F(x^*)} + (1-\alpha)F(x^*)$$

$$< F(x^*)$$

Take a sequence $\alpha_k \to 0 \quad k \to \infty$

$$x_k := x^* + \alpha_k(x_0 - x^*) \longrightarrow x^*$$

BUT $\underline{\underline{F(x_k) < F(x^*)}}$

$\Rightarrow$ there is a neighborhood I around x^* s.t.

$$F(x^*) \leq F(x) \qquad x \in I$$

$\Rightarrow x^*$ cannot be a local minimum ⚡

Methods for Minimization of $F: \mathbb{R}^n \to \mathbb{R}$

Gradient descent

level sets:

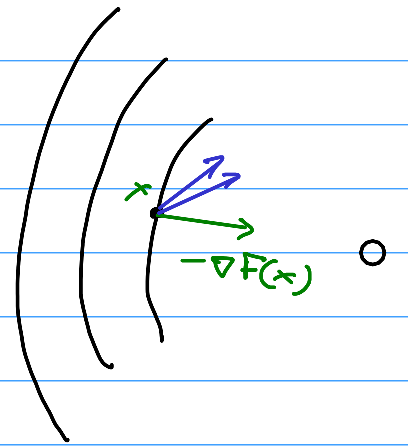

Any direction Δx s.t. $\nabla F(x)^T \Delta x < 0$ is a descent direction

choice $\Delta x = -\nabla F(x)$ gradient descent direction

$\rightsquigarrow$ greedy approach

Guarantee for gradient descent direction:

if $\nabla F(x) \neq 0$ and $\alpha > 0$ suff. small

$$F(x - \alpha \nabla F(x)) \leq F(x)$$

Gradient descent iteration:

$$x^{(k+1)} = x^{(k)} - t^{(k)} \nabla F(x^{(k)})$$

How to choose step size $t^{(k)}$?

This is a 1D problem