

Optimization (ND Methods)

What is the optimal solution? (ND)

$$f(\mathbf{x}^*) = \min_x f(\mathbf{x})$$

(First-order) Necessary condition

$$\begin{aligned} & f(\mathbf{x}) \\ & \approx f(\tilde{\mathbf{x}}) \\ & \nearrow \end{aligned}$$

1D: $f'(x) = 0$

ND : $\nabla f(\tilde{\mathbf{x}}^*) = \mathbf{0} \longrightarrow$ gives stationary solution \mathbf{x}^*

(Second-order) Sufficient condition

1D: $f''(x) > 0$

ND : $\underline{H}(\underline{\mathbf{x}}^*)$ is positive definite $\longrightarrow \mathbf{x}^*$ is minimizer

Taking derivatives...

$$f: \mathbb{R}^n \rightarrow \mathbb{R}$$

$$f(\underline{x}) = f(x_1, x_2, \dots, x_n)$$

$$\frac{\partial f}{\partial x_1} \quad \frac{\partial f}{\partial x_2} \quad \dots \quad \frac{\partial f}{\partial x_n} \implies \underline{\nabla f}(\underline{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \frac{\partial f}{\partial x_2} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix} \quad (n \times 1)$$

gradient of f

$\frac{d}{dx_i} \nabla f$

$$H(\underline{x}) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \dots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \dots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \dots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix} \quad (n \times n)$$

Symm!

$(\nabla f)_i = \frac{\partial f}{\partial x_i}$
 $(H)_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$

From linear algebra:

scalar $y^T H y$
vector $y \cdot H y$

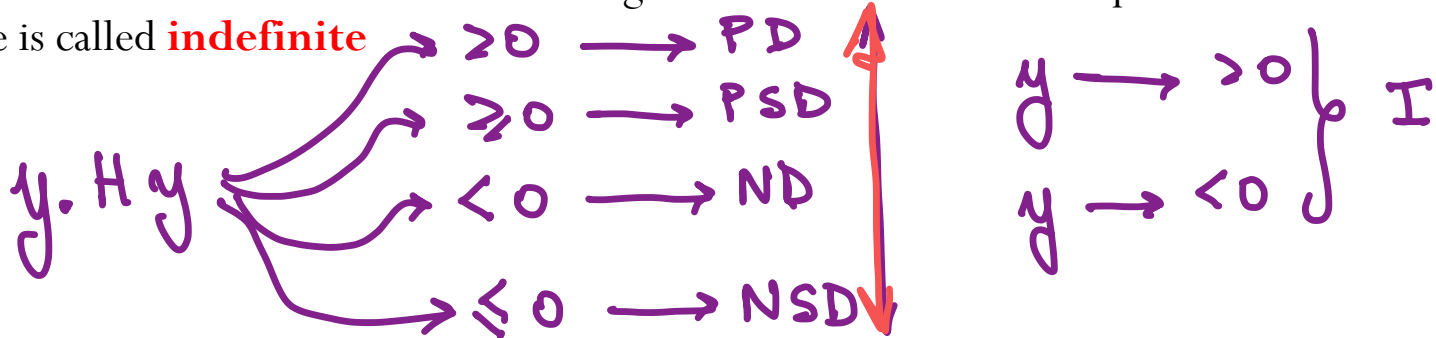
A symmetric $n \times n$ matrix H is positive definite if $y^T H y > 0$ for any $y \neq 0$

A symmetric $n \times n$ matrix H is positive semi-definite if $y^T H y \geq 0$ for any $y \neq 0$

A symmetric $n \times n$ matrix H is negative definite if $y^T H y < 0$ for any $y \neq 0$

A symmetric $n \times n$ matrix H is negative semi-definite if $y^T H y \leq 0$ for any $y \neq 0$

A symmetric $n \times n$ matrix H that is not negative semi-definite and not positive semi-definite is called indefinite



la. eig(H)

$$f(\mathbf{x}^*) = \min_{\mathbf{x}} f(\mathbf{x})$$

First order necessary condition: $\nabla f(\mathbf{x}) = \mathbf{0}$

Second order sufficient condition: **H(x) is positive definite**

How can we find out if the Hessian is positive definite?

$$\boxed{Hy = \lambda y} \rightarrow (\lambda, y) \rightarrow \text{are eigenpairs of } H$$

$$y^T H y = \lambda y^T y = \lambda \|y\|_2^2$$

$$\lambda = \frac{y^T H y}{\|y\|_2^2}$$

always positive

* $\lambda_i > 0 \quad \forall i \Rightarrow y^T H y > 0 \quad \forall y \Rightarrow H$ is pos. def $\Rightarrow x^*$ is minimizer

* $\lambda_i < 0 \quad \forall i \Rightarrow y^T H y < 0 \quad \forall y \Rightarrow H$ is neg def $\Rightarrow x^*$ is maximizer

* $\left. \begin{array}{l} \lambda_i > 0 \\ \lambda_i < 0 \end{array} \right\} \rightarrow H$ is indefinite $\rightarrow x^*$ is saddle point

Types of optimization problems

$$f(\mathbf{x}^*) = \min_{\mathbf{x}} f(\mathbf{x})$$

f : nonlinear, continuous
and smooth

Gradient-free methods

Evaluate $f(\mathbf{x})$

Gradient (first-derivative) methods

Evaluate $f(\mathbf{x}), \nabla f(\mathbf{x})$

Second-derivative methods

Evaluate $f(\mathbf{x}), \nabla f(\mathbf{x}), \nabla^2 f(\mathbf{x})$

H(x)

Example (ND)

Consider the function $f(x_1, x_2) = 2x_1^3 + 4x_2^2 + 2x_2 - 24x_1$

Find the stationary point and check the sufficient condition

$$\nabla f = \begin{bmatrix} 6x_1^2 - 24 \\ 8x_2 + 2 \end{bmatrix}; \quad H = \begin{bmatrix} 12x_1 & 0 \\ 0 & 8 \end{bmatrix}$$

$$1) \nabla f = \underline{0} \Rightarrow \begin{bmatrix} 6x_1^2 - 24 \\ 8x_2 + 2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \Rightarrow \begin{aligned} 6x_1^2 = 24 &\rightarrow x_1^2 = 4 \rightarrow x_1 = \pm 2 \\ 8x_2 = -2 &\Rightarrow x_2 = -0.25 \end{aligned}$$

stationary points: $x^* = \begin{bmatrix} +2 \\ -0.25 \end{bmatrix}$ $x^* = \begin{bmatrix} -2 \\ -0.25 \end{bmatrix}$

$$2) H \begin{pmatrix} -2 \\ -0.25 \end{pmatrix} = \begin{bmatrix} -24 & 0 \\ 0 & 8 \end{bmatrix} \Rightarrow \text{indefinite} \downarrow \text{saddle point} \quad \left\{ \begin{aligned} H \begin{pmatrix} 2 \\ -0.25 \end{pmatrix} = \begin{bmatrix} 24 & 0 \\ 0 & 8 \end{bmatrix} \text{ pos. def.} \\ \downarrow \text{Minimizer!} \end{aligned} \right.$$


Optimization in ND: Steepest Descent Method

$$\min_x f(x)$$
$$\boxed{-\nabla f}$$

Given a function

$f(\mathbf{x}): \mathcal{R}^n \rightarrow \mathcal{R}$ at a point \mathbf{x} , the function will decrease its value in the direction of steepest descent: $-\nabla f(\mathbf{x})$

What is the steepest descent direction?



Steepest Descent Method

$$\tilde{x}_2 = \tilde{x}_1 - \nabla f(\tilde{x}_1)$$

Start with initial guess:

$$\mathbf{x}_0 = \begin{bmatrix} 3 \\ 3 \end{bmatrix}$$

Check the update:

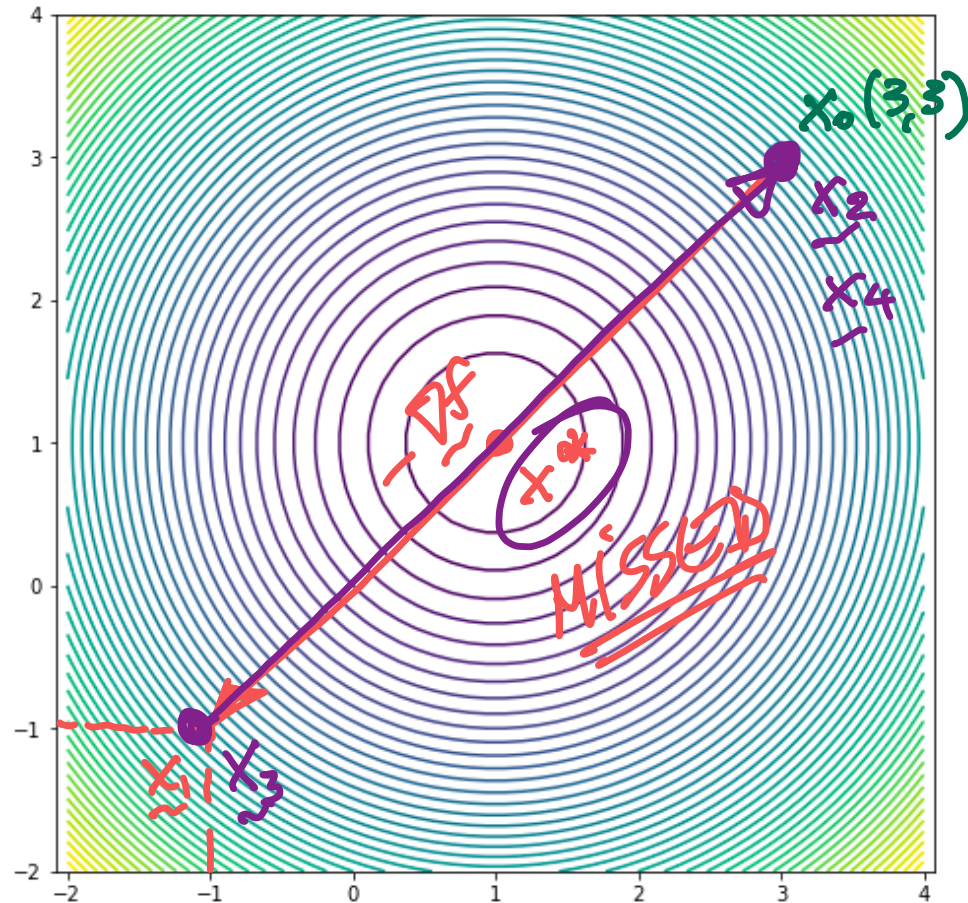
$$\tilde{x}_1 = \tilde{x}_0 - \nabla f(\tilde{x}_0)$$

$$\nabla f = \begin{bmatrix} 2(x_1 - 1) \\ 2(x_2 - 1) \end{bmatrix}$$

$$\nabla f(\tilde{x}_0) = \begin{bmatrix} 4 \\ 4 \end{bmatrix}$$

$$\tilde{x}_1 = \begin{bmatrix} 3 \\ 3 \end{bmatrix} - \begin{bmatrix} 4 \\ 4 \end{bmatrix} = \begin{bmatrix} -1 \\ -1 \end{bmatrix}$$

$$f(x_1, x_2) = (x_1 - 1)^2 + (x_2 - 1)^2$$



Steepest Descent Method

Update the variable with:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \nabla f(\mathbf{x}_k)$$


How far along the gradient should we go? What is the “best size” for α_k ?

$$\underline{x}_1 = \underline{x}_0 - \underline{\underline{0.5}} \nabla f(x_0)$$

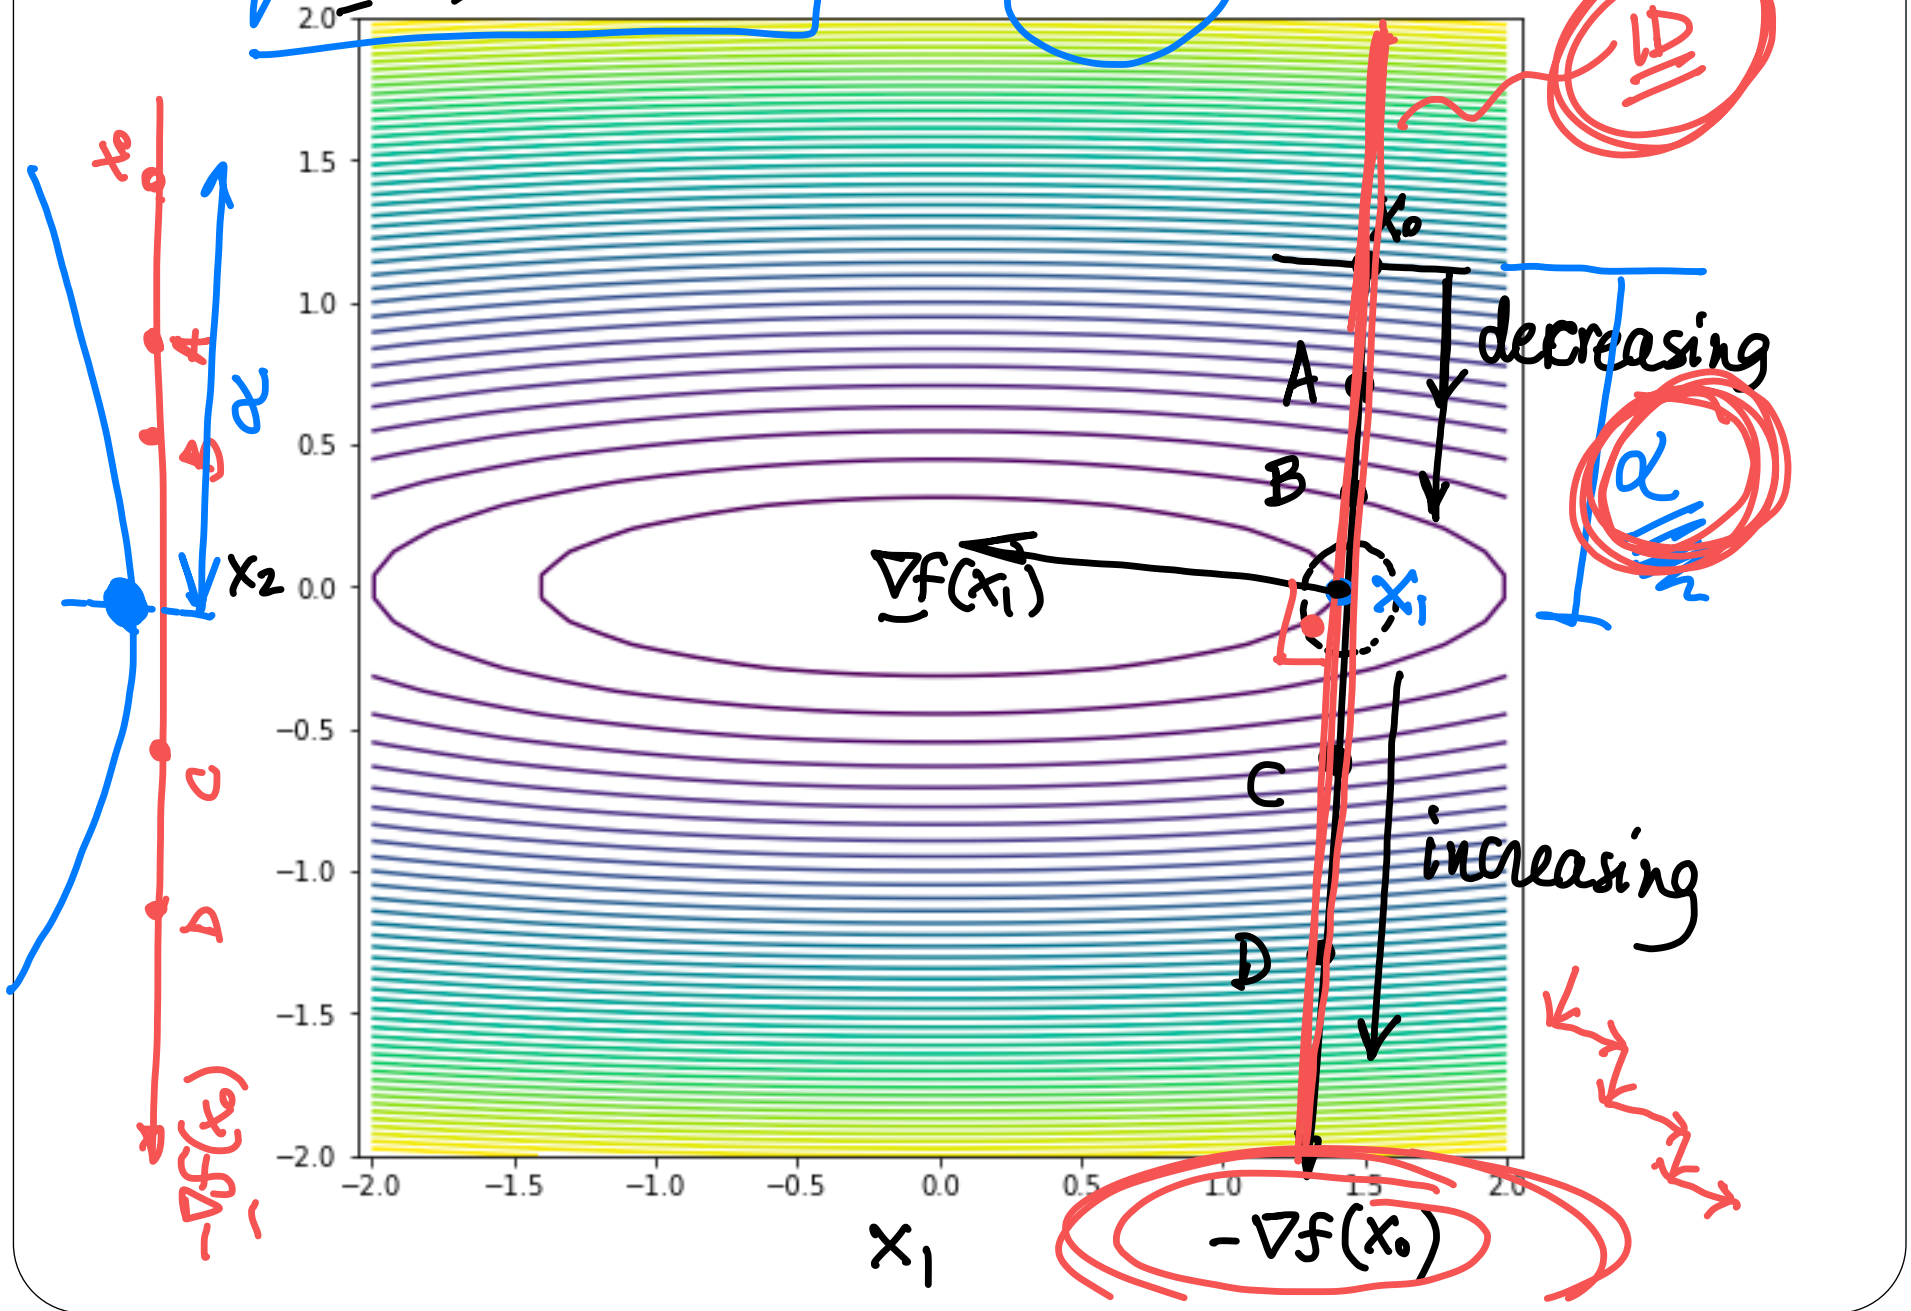
$$\boxed{\alpha = 0.5}$$

How can we get α ?

$$f(x_1, x_2) = (x_1 - 1)^2 + (x_2 - 1)^2$$



Find $\underline{x}_1 = \underline{x}_0 - \alpha \underline{\nabla} f(\underline{x}_0)$ s.t. $f(x_1)$ is minimized



Steepest Descent Method

Algorithm:

Initial guess: \underline{x}_0

Evaluate: $\underline{s}_k = -\underline{\nabla} f(\underline{x}_k)$

Perform a line search to obtain α_k (for example, Golden Section Search)

$$\alpha_k = \operatorname{argmin}_{\alpha} f(\underline{x}_k + \alpha \underline{s}_k)$$

Update: $\underline{x}_{k+1} = \underline{x}_k + \alpha_k \underline{s}_k$

1D optimization problem

several
fc
eval.

$$\underline{x}_{k+1} = \underline{x}_k + \alpha \underline{s}_k$$

Line Search

$$f(x_{k+1})$$

$$x_{k+1} = x_k - \alpha_k \nabla f(x_k)$$

we want to find α_k s.t.

$$\min_{\alpha} f(x_k - \alpha \nabla f(x_k))$$

x_{k+1}

1st order condition $\frac{df}{d\alpha} = 0 \rightarrow$ gives α

$$\frac{df}{d\alpha} = \frac{\partial f}{\partial x_{k+1}} \frac{\partial x_{k+1}}{\partial \alpha} = \nabla f(x_{k+1}) \cdot \nabla f(x_k) = 0$$
$$\nabla f(x_{k+1}) \cdot \nabla f(x_k) = 0$$

Zig-zag
pattern
convergence.

$\nabla f(x_{k+1})$ is orthogonal to
 $\nabla f(x_k)$

Example

$$\min_{x_1, x_2} f(x_1, x_2)$$

Consider minimizing the function

$$\underline{\underline{f(x_1, x_2) = 10(x_1)^3 - (x_2)^2 + x_1 - 1}}$$

Given the initial guess

$$x_1 = 2, x_2 = 2$$

$$\tilde{x}_0 = \begin{bmatrix} 2 \\ 2 \end{bmatrix}$$

what is the direction of the first step of gradient descent?

$$\nabla f = \begin{bmatrix} 30x_1^2 + 1 \\ -2x_2 \end{bmatrix}$$

$$\underline{\underline{\nabla f(\tilde{x}_0) = \begin{bmatrix} 121 \\ -4 \end{bmatrix}}}$$

steepest descent
direction

$$\Rightarrow \begin{bmatrix} -121 \\ +4 \end{bmatrix}$$

Newton's Method

Using Taylor Expansion, we build the approximation:

$$\underbrace{f(\underline{x} + \underline{s})}_{\text{non linear}} = f(\underline{x}) + \nabla f(\underline{x})^T \underline{s} + \frac{1}{2} \underline{s}^T \underline{H} \underline{s} = \hat{f}(\underline{s})$$

quadratic approx of f

1st order condition: $\nabla \hat{f} = 0$

$$\nabla f(\underline{x}) + \underline{H} \underline{s} = 0$$

\underline{H} is symmetric
 $\underline{H} = \underline{H}^T$

$$\underline{H}(\underline{x}) \underline{s} = -\nabla f(\underline{x})$$

→ solve lin sys to find Newton step \underline{s}

Newton's Method

Algorithm:

Initial guess: \mathbf{x}_0

Solve: $\mathbf{H}_f(\mathbf{x}_k) \mathbf{s}_k = -\nabla f(\mathbf{x}_k)$ → solve $\underline{O(n^3)}$ \mathbf{s}_k

Update: $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{s}_k$



$$H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$$

$O(n^2)$

Note that the Hessian is related to the curvature and therefore contains the information about how large the step should be.

Try this out!

$$f(x, y) = 0.5x^2 + 2.5y^2$$



When using the Newton's Method to find the minimizer of this function, estimate the number of iterations it would take for convergence?

- A) 1 B) 2-5 C) 5-10 D) More than 10 E) Depends on the initial guess

Newton's Method Summary

Algorithm:

Initial guess: \mathbf{x}_0

Solve: $\mathbf{H}_f(\mathbf{x}_k) \mathbf{s}_k = -\nabla f(\mathbf{x}_k)$

Update: $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{s}_k$

About the method...

- Typical quadratic convergence 😊
- Need second derivatives ☹️
- Local convergence (start guess close to solution)
- Works poorly when Hessian is nearly indefinite
- Cost per iteration: $O(n^3)$