Terminology

$\exists$	means "there exists"
$\exists!$	means "there exists a unique"
$\forall$	means "for all"
$:=$	means that the left-hand side is a new notation for the expression on the right-hand side
iff	means "if and only if"

Conjugate directions

Motivation and definition

The conjugate gradient method is developed for a quadratic function of n variables

$$q(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\mathsf{T}\boldsymbol{H}\boldsymbol{x} + \boldsymbol{c}^\mathsf{T}\boldsymbol{x}, \quad \text{with } \boldsymbol{H} \text{ positive definite.}$$

Let $\boldsymbol{g}_k := \boldsymbol{\nabla} q(\boldsymbol{x}_k) = \boldsymbol{H}\boldsymbol{x}_k + \boldsymbol{c}$ and use exact line searches along directions $\boldsymbol{d}_k$ (to be determined):

$$\text{(EL)} \qquad \boldsymbol{g}_{k+1}^\mathsf{T}\boldsymbol{d}_k = 0.$$

Important property:

$$(\Delta\boldsymbol{g}) \qquad \begin{cases} \boldsymbol{g}_1 = \boldsymbol{H}\boldsymbol{x}_1 + \boldsymbol{c} \\ \boldsymbol{g}_2 = \boldsymbol{H}\boldsymbol{x}_2 + \boldsymbol{c} \end{cases} \quad \Rightarrow \quad \boldsymbol{g}_2 - \boldsymbol{g}_1 = \boldsymbol{H}(\boldsymbol{x}_2 - \boldsymbol{x}_1).$$

Idea of conjugate directions: (contours of a quadratic function)

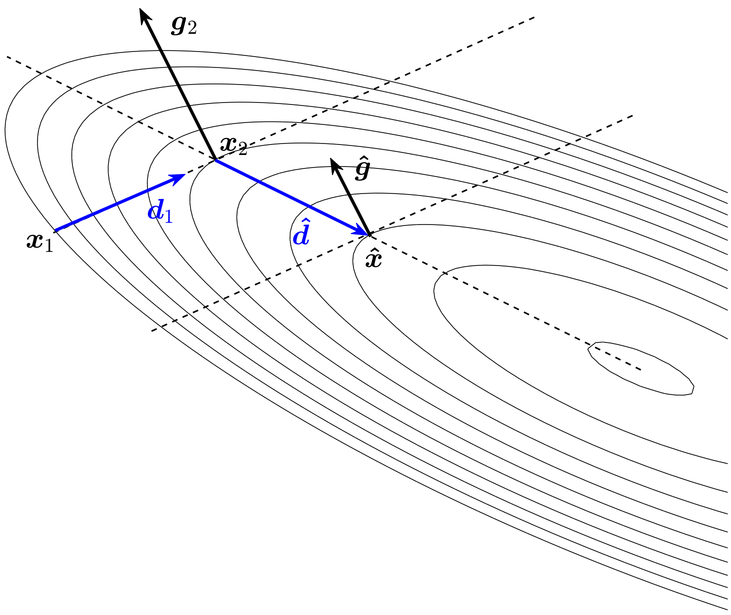

The property (EL) $\boldsymbol{g}_2^\mathsf{T}\boldsymbol{d}_1 = 0$ should be kept along the new search direction $\hat{\boldsymbol{d}}$. At any point along the new search direction, e.g., $\hat{\boldsymbol{x}} = \boldsymbol{x}_2 + \hat{\boldsymbol{d}}$, we want $\hat{\boldsymbol{g}}^\mathsf{T}\boldsymbol{d}_1 = 0$. Then $(\Delta\boldsymbol{g})$ implies

$$0 = \boldsymbol{d}_1^\mathsf{T}\hat{\boldsymbol{g}} - \boldsymbol{d}_1^\mathsf{T}\boldsymbol{g}_2 = \boldsymbol{d}_1^\mathsf{T}(\hat{\boldsymbol{g}} - \boldsymbol{g}_2) = \boldsymbol{d}_1^\mathsf{T}\boldsymbol{H}(\hat{\boldsymbol{x}} - \boldsymbol{x}_2) = \boldsymbol{d}_1^\mathsf{T}\boldsymbol{H}\hat{\boldsymbol{d}}.$$

DEFINITION. *The nonzero vectors $\boldsymbol{d}_1, \dots, \boldsymbol{d}_n$ are **$\boldsymbol{H}$-conjugate** iff*

$$\boldsymbol{d}_i^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_j = 0, \quad \forall i \neq j.$$

Note: $\boldsymbol{H}$ positive definite $\quad \Rightarrow \quad \boldsymbol{d}_j^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_j > 0.$

LEMMA: The $\boldsymbol{H}$-conjugate vectors $\boldsymbol{d}_1, \dots, \boldsymbol{d}_n$ are linearly independent.

PROOF: Given $\sum_{i=1}^{n} \lambda_i \boldsymbol{d}_i = 0$, multiply by $\boldsymbol{d}_j^{\mathsf{T}} \boldsymbol{H}$ from the left with j arbitrary to get

$$\sum_{i=1}^{n} \lambda_i \boldsymbol{d}_j^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_i = 0 \quad \Leftrightarrow \quad \lambda_j \boldsymbol{d}_j^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_j = 0 \quad \Leftrightarrow \quad \lambda_j = 0.$$

□

Minimization along conjugate directions

With the minimizer at $\bar{\boldsymbol{x}}$, we can write (Exercise 1.5)

$$q(\boldsymbol{x}) = \frac{1}{2}(\boldsymbol{x} - \bar{\boldsymbol{x}})^{\mathsf{T}} \boldsymbol{H} (\boldsymbol{x} - \bar{\boldsymbol{x}}) + a \quad \text{with} \quad a = q(\bar{\boldsymbol{x}}).$$

Assume that we have $\boldsymbol{H}$-conjugate directions $\boldsymbol{d}_1, \dots, \boldsymbol{d}_n$ (how to find these comes later). Since these form a basis, we can make a variable change. With the invertible matrix

$$\boldsymbol{S} := \begin{pmatrix} \boldsymbol{d}_1 & \boldsymbol{d}_2 & \cdots & \boldsymbol{d}_n \end{pmatrix}$$

we can diagonalize $\boldsymbol{H}$:

$$\boldsymbol{S}^{\mathsf{T}} \boldsymbol{H} \boldsymbol{S} = \begin{pmatrix} \boldsymbol{d}_1^{\mathsf{T}} \\ \vdots \\ \boldsymbol{d}_n^{\mathsf{T}} \end{pmatrix} \boldsymbol{H} \begin{pmatrix} \boldsymbol{d}_1 & \boldsymbol{d}_2 & \cdots & \boldsymbol{d}_n \end{pmatrix} = \begin{pmatrix} \boldsymbol{d}_1^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_1 & \cdots & \boldsymbol{d}_1^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_n \\ \vdots & & \vdots \\ \boldsymbol{d}_n^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_1 & \cdots & \boldsymbol{d}_n^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_n \end{pmatrix} = \operatorname{diag}(\xi_1, \dots, \xi_n),$$

with $\xi_i := \boldsymbol{d}_i^{\mathsf{T}} \boldsymbol{H} \boldsymbol{d}_i > 0$. Therefore, we make the coordinate change

$$\boldsymbol{x} = \tilde{\boldsymbol{x}}(\boldsymbol{\alpha}) \quad \text{with} \quad \tilde{\boldsymbol{x}}(\boldsymbol{\alpha}) = \boldsymbol{x}_1 + \boldsymbol{S}\boldsymbol{\alpha} = \boldsymbol{x}_1 + \begin{pmatrix} \boldsymbol{d}_1 & \cdots & \boldsymbol{d}_n \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix} = \boldsymbol{x}_1 + \sum_{j=1}^{n} \alpha_j \boldsymbol{d}_j.$$

The minimizer satisfies $\bar{\boldsymbol{x}} = \boldsymbol{x}_1 + \boldsymbol{S}\bar{\boldsymbol{\alpha}}$; hence, $\boldsymbol{x} - \bar{\boldsymbol{x}} = \boldsymbol{S}(\boldsymbol{\alpha} - \bar{\boldsymbol{\alpha}})$ and

$$(*) \qquad q(\tilde{\boldsymbol{x}}(\boldsymbol{\alpha})) = \frac{1}{2}(\boldsymbol{\alpha} - \bar{\boldsymbol{\alpha}})^{\mathsf{T}} \boldsymbol{S}^{\mathsf{T}} \boldsymbol{H} \boldsymbol{S} (\boldsymbol{\alpha} - \bar{\boldsymbol{\alpha}}) + a = \frac{1}{2} \sum_{i=1}^{n} \xi_i (\alpha_i - \bar{\alpha}_i)^2 + a.$$

The starting point $\boldsymbol{x}_1$ corresponds to $\boldsymbol{\alpha} = \boldsymbol{0}$. Then we minimize along the direction $\boldsymbol{d}_1$ and find that the minimum of

$$q(\boldsymbol{x}_1 + \lambda \boldsymbol{d}_1) = q(\boldsymbol{x}_1 + \boldsymbol{S}\boldsymbol{\alpha}_1) = \frac{1}{2}\xi_1(\lambda - \bar{\alpha}_1)^2 + \frac{1}{2}\sum_{i=2}^{n} \xi_i (0 - \bar{\alpha}_i)^2 + a \quad \text{with} \quad \boldsymbol{\alpha}_1 := \begin{pmatrix} \lambda \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

occurs at $\lambda = \bar{\alpha}_1$, which defines $\boldsymbol{x}_2 := \boldsymbol{x}_1 + \bar{\alpha}_1 \boldsymbol{d}_1$. From this point, we search along $\boldsymbol{d}_2$, which corresponds to setting $\boldsymbol{\alpha} = \boldsymbol{\alpha}_2 := (\bar{\alpha}_1, \lambda, 0, \ldots, 0)^\mathsf{T}$, and the minimizer of

$$\begin{aligned} q(\boldsymbol{x}_2 + \lambda \boldsymbol{d}_2) &= q(\boldsymbol{x}_1 + \bar{\alpha}_1 \boldsymbol{d}_1 + \lambda \boldsymbol{d}_2) = q(\boldsymbol{x}_1 + \boldsymbol{S}\boldsymbol{\alpha}_2) \\ &= \frac{1}{2}\xi_2(\lambda - \bar{\alpha}_2)^2 + \frac{1}{2}\sum_{i=3}^{n} \xi_i (0 - \bar{\alpha}_i)^2 + a, \end{aligned}$$

is clearly $\lambda = \bar{\alpha}_2$. Thus, we obtain

$$\begin{aligned} \boldsymbol{x}_2 &:= \boldsymbol{x}_1 + \bar{\alpha}_1 \boldsymbol{d}_1 = \boldsymbol{x}_1 + \boldsymbol{S}\bar{\boldsymbol{\alpha}}_1 = \tilde{\boldsymbol{x}}(\bar{\boldsymbol{\alpha}}_1) \quad \text{with} \quad \bar{\boldsymbol{\alpha}}_1 := \begin{pmatrix} \bar{\alpha}_1 & 0 & \ldots & 0 \end{pmatrix}^\mathsf{T}, \\ \boldsymbol{x}_3 &:= \boldsymbol{x}_2 + \bar{\alpha}_2 \boldsymbol{d}_2 = \boldsymbol{x}_1 + \boldsymbol{S}\bar{\boldsymbol{\alpha}}_2 = \tilde{\boldsymbol{x}}(\bar{\boldsymbol{\alpha}}_2) \quad \text{with} \quad \bar{\boldsymbol{\alpha}}_2 := \begin{pmatrix} \bar{\alpha}_1 & \bar{\alpha}_2 & 0 & \ldots & 0 \end{pmatrix}^\mathsf{T}, \\ &\ \vdots \\ \boldsymbol{x}_{k+1} &:= \boldsymbol{x}_k + \bar{\alpha}_k \boldsymbol{d}_k = \boldsymbol{x}_1 + \boldsymbol{S}\bar{\boldsymbol{\alpha}}_k = \tilde{\boldsymbol{x}}(\bar{\boldsymbol{\alpha}}_k) \quad \text{with} \quad \bar{\boldsymbol{\alpha}}_k := \begin{pmatrix} \bar{\alpha}_1 & \ldots & \bar{\alpha}_k & 0 & \ldots & 0 \end{pmatrix}^\mathsf{T}. \end{aligned}$$

(There is an explicit formula for $\bar{\alpha}_k$ for a quadratic function.) The minimizer is found after n line searches, since $q(\boldsymbol{x}_{n+1}) = q(\tilde{\boldsymbol{x}}(\bar{\boldsymbol{\alpha}}_n)) = a = q(\bar{\boldsymbol{x}})$. We have proved the following theorem.

THEOREM. *Minimization with exact line searches of a quadratic function along a sequence of $\boldsymbol{H}$-conjugate directions terminates after (at most) n steps.*

LEMMA CD. *Minimization with exact line searches of a quadratic function from the starting point $\boldsymbol{x}_1$ along a sequence of $\boldsymbol{H}$-conjugate directions $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n$ results in a sequence of points $\boldsymbol{x}_2, \ldots, \boldsymbol{x}_{n+1}$ such that, for $k = 1, \ldots, n$,*

$$\boldsymbol{g}_{k+1}^\mathsf{T} \boldsymbol{d}_i = \boldsymbol{\nabla} q(\boldsymbol{x}_{k+1})^\mathsf{T} \boldsymbol{d}_i = 0 \quad \textit{for} \quad i = 1, \ldots, k.$$

(The new gradient is perpendicular to all previous search directions.)

Remark: Since $\boldsymbol{g}_{n+1}$ is perpendicular to all the basis vectors $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n$, it has to be the zero vector.

PROOF: We differentiate $(*)$, which is

$$q(\tilde{\boldsymbol{x}}(\boldsymbol{\alpha})) = \frac{1}{2}\sum_{i=1}^{n} \xi_i (\alpha_i - \bar{\alpha}_i)^2 + a, \quad \text{where} \quad \tilde{\boldsymbol{x}}(\boldsymbol{\alpha}) = \boldsymbol{x}_1 + \sum_{j=1}^{n} \alpha_j \boldsymbol{d}_j.$$

with respect to α_i and use the the chain rule for the left-hand side to get ($\boldsymbol{\nabla}$ refers to derivatives with respect to $\boldsymbol{x}$)

$$\boldsymbol{\nabla} q(\tilde{\boldsymbol{x}}(\boldsymbol{\alpha}))^\mathsf{T} \boldsymbol{d}_i = \xi_i (\alpha_i - \bar{\alpha}_i).$$

In this formula, we let $\boldsymbol{\alpha} = \bar{\boldsymbol{\alpha}}_k = (\bar{\alpha}_1, \ldots, \bar{\alpha}_k, 0, \ldots, 0)^\mathsf{T}$ to get the result. $\square$

Derivation of the Conjugate Gradient method

DEFINITION. *A real inner product (scalar product) of two vectors* $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$ *is a real-valued function denoted by* $\langle \boldsymbol{u}, \boldsymbol{v} \rangle$ *with the properties*

$$\begin{aligned} \langle \boldsymbol{u}, \boldsymbol{v} \rangle &= \langle \boldsymbol{v}, \boldsymbol{u} \rangle, \\ \langle a\boldsymbol{u} + b\boldsymbol{v}, \boldsymbol{w} \rangle &= a\langle \boldsymbol{u}, \boldsymbol{w} \rangle + b\langle \boldsymbol{v}, \boldsymbol{w} \rangle, \quad a, b \in \mathbb{R}, \\ \langle \boldsymbol{u}, \boldsymbol{u} \rangle &\geq 0 \quad \textit{with equality iff } \boldsymbol{u} = \boldsymbol{0}. \end{aligned}$$

One example is $\langle \boldsymbol{u}, \boldsymbol{v} \rangle = \boldsymbol{u}^\mathsf{T} \boldsymbol{v} = \boldsymbol{u}^\mathsf{T} \boldsymbol{I} \boldsymbol{v}$.

Another is $\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\boldsymbol{H}} = \boldsymbol{u}^\mathsf{T} \boldsymbol{H} \boldsymbol{v}$ with $\boldsymbol{H}$ positive definite.

We use the terminology

$$\boldsymbol{u} \text{ and } \boldsymbol{v} \text{ are } \boldsymbol{H}\text{-orthogonal} \quad \Leftrightarrow \quad \langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\boldsymbol{H}} = 0 \quad \Leftrightarrow \quad \boldsymbol{u} \text{ and } \boldsymbol{v} \text{ are } \boldsymbol{H}\text{-conjugate}$$

We now use the *Gram-Schmidt orthogonalization process* with $\langle \boldsymbol{u}, \boldsymbol{v} \rangle_{\boldsymbol{H}}$ to derive the Conjugate Gradient method:

Step 0: Given the starting point $\boldsymbol{x}_1$, compute $\boldsymbol{g}_1 = \boldsymbol{\nabla} q(\boldsymbol{x}_1)$.

Step 1: Set $\boldsymbol{d}_1 := -\boldsymbol{g}_1$ (steepest descent)
Do exact line search, set $\boldsymbol{x}_2 := \boldsymbol{x}_1 + \bar{\alpha}_1 \boldsymbol{d}_1$ and compute $\boldsymbol{g}_2 = \boldsymbol{\nabla} q(\boldsymbol{x}_2)$.

Step 2: Define $\boldsymbol{d}_2$ so that $\langle \boldsymbol{d}_2, \boldsymbol{d}_1 \rangle_{\boldsymbol{H}} = 0$ using projection with respect to the $\boldsymbol{H}$-inner product:

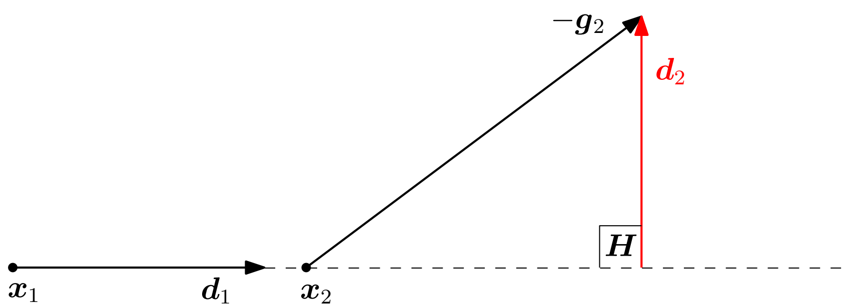

Define (use the projection formula in linear algebra):

$$\boldsymbol{d}_2 := -\boldsymbol{g}_2 - \frac{\langle -\boldsymbol{g}_2, \boldsymbol{d}_1 \rangle_{\boldsymbol{H}}}{\langle \boldsymbol{d}_1, \boldsymbol{d}_1 \rangle_{\boldsymbol{H}}} \boldsymbol{d}_1$$

Do exact line search, set $\boldsymbol{x}_3 := \boldsymbol{x}_2 + \bar{\alpha}_2 \boldsymbol{d}_2$ and compute $\boldsymbol{g}_3 = \boldsymbol{\nabla} q(\boldsymbol{x}_3)$.

$\vdots$

Step k**:** Project $-\boldsymbol{g}_k$ onto the subspace spanned by $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_{k-1}$:

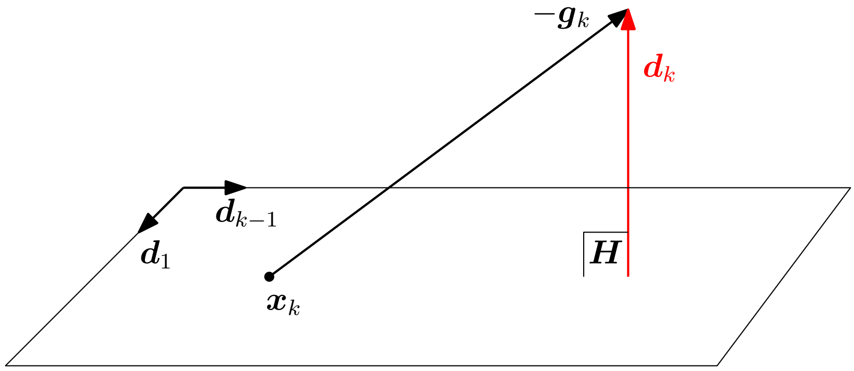

Define (sum of projections on each basis vector):

$$\text{(D)} \qquad \begin{aligned} \boldsymbol{d}_k &:= -\boldsymbol{g}_k - \sum_{i=1}^{k-1} \frac{\langle -\boldsymbol{g}_k, \boldsymbol{d}_i \rangle_{\boldsymbol{H}}}{\langle \boldsymbol{d}_i, \boldsymbol{d}_i \rangle_{\boldsymbol{H}}} \boldsymbol{d}_i \\ &= -\boldsymbol{g}_k + \sum_{i=1}^{k-1} \beta_i \boldsymbol{d}_i, \quad \text{with} \quad \beta_i := \frac{\langle \boldsymbol{g}_k, \boldsymbol{d}_i \rangle_{\boldsymbol{H}}}{\langle \boldsymbol{d}_i, \boldsymbol{d}_i \rangle_{\boldsymbol{H}}}. \end{aligned}$$

The task is now to simplify this formula. Recall:

$$(\Delta \boldsymbol{g}) \qquad \boldsymbol{g}_2 - \boldsymbol{g}_1 = \boldsymbol{H}(\boldsymbol{x}_2 - \boldsymbol{x}_1)$$

$$\text{(Lemma CD)} \qquad \boldsymbol{g}_{k+1}^\top \boldsymbol{d}_i = 0, \quad i = 1, \ldots, k$$

We also need that (D) and Lemma CD imply

$$\boldsymbol{g}_{k+1}^\top \boldsymbol{g}_j = \boldsymbol{g}_{k+1}^\top \left(-\boldsymbol{d}_j + \sum_{i=1}^{j-1} \beta_i \boldsymbol{d}_i \right) = 0, \quad j = 1, \ldots, k.$$

(*The new gradient is also perpendicular to all previous gradients.*) Numerator of β_i:

$$\begin{aligned} \langle \boldsymbol{g}_k, \boldsymbol{d}_i \rangle_{\boldsymbol{H}} &= \boldsymbol{g}_k^\top \boldsymbol{H} \boldsymbol{d}_i = \frac{1}{\bar{\alpha}_i} \boldsymbol{g}_k^\top \boldsymbol{H}(\bar{\alpha}_i \boldsymbol{d}_i) = \frac{1}{\bar{\alpha}_i} \boldsymbol{g}_k^\top \boldsymbol{H}(\boldsymbol{x}_{i+1} - \boldsymbol{x}_i) \\ &\overset{(\Delta \boldsymbol{g})}{=} \frac{1}{\bar{\alpha}_i} \boldsymbol{g}_k^\top (\boldsymbol{g}_{i+1} - \boldsymbol{g}_i) = \begin{cases} 0 & \text{if } i < k-1, \\ \frac{1}{\bar{\alpha}_{k-1}} \boldsymbol{g}_k^\top \boldsymbol{g}_k & \text{if } i = k-1. \end{cases} \end{aligned}$$

Hence, (D) is simplified to $\boldsymbol{d}_k = -\boldsymbol{g}_k + \beta_{k-1} \boldsymbol{d}_{k-1}$. Denominator of β_i $(= \beta_{k-1})$:

$$\begin{aligned} \langle \boldsymbol{d}_i, \boldsymbol{d}_i \rangle_{\boldsymbol{H}} &= \boldsymbol{d}_i^\top \boldsymbol{H} \boldsymbol{d}_i = \frac{1}{\bar{\alpha}_i} \boldsymbol{d}_i^\top \boldsymbol{H}(\bar{\alpha}_i \boldsymbol{d}_i) = \frac{1}{\bar{\alpha}_i} \boldsymbol{d}_i^\top \boldsymbol{H}(\boldsymbol{x}_{i+1} - \boldsymbol{x}_i) \\ &= \frac{1}{\bar{\alpha}_i} \boldsymbol{d}_i^\top (\boldsymbol{g}_{i+1} - \boldsymbol{g}_i) = -\frac{1}{\bar{\alpha}_i} \boldsymbol{d}_i^\top \boldsymbol{g}_i \\ &\overset{\text{(D)}}{=} -\frac{1}{\bar{\alpha}_i} (-\boldsymbol{g}_i + \beta_{i-1} \boldsymbol{d}_{i-1})^\top \boldsymbol{g}_i = \frac{1}{\bar{\alpha}_i} \boldsymbol{g}_i^\top \boldsymbol{g}_i = \frac{1}{\bar{\alpha}_{k-1}} \boldsymbol{g}_{k-1}^\top \boldsymbol{g}_{k-1}. \end{aligned}$$

Thus, (D) becomes

$$\boldsymbol{d}_k = -\boldsymbol{g}_k + \frac{\boldsymbol{g}_k^\top \boldsymbol{g}_k}{\boldsymbol{g}_{k-1}^\top \boldsymbol{g}_{k-1}} \boldsymbol{d}_{k-1} = -\boldsymbol{g}_k + \frac{\|\boldsymbol{g}_k\|^2}{\|\boldsymbol{g}_{k-1}\|^2} \boldsymbol{d}_{k-1}$$

Conjugate gradient method for a general function $f(\boldsymbol{x})$:
Cyclic-coordinate search along the basis vectors $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_n$ defined by

$$\boldsymbol{d}_k = -\boldsymbol{\nabla} f(\boldsymbol{x}_k) + \frac{\|\boldsymbol{\nabla} f(\boldsymbol{x}_k)\|^2}{\|\boldsymbol{\nabla} f(\boldsymbol{x}_{k-1})\|^2} \boldsymbol{d}_{k-1}$$

$\oplus$ Only vectors involved – suitable for large problems

$\ominus$ Requires accurate line searches