

Unconstrained Optimization II

Asst. Prof. Dr.-Ing. Sudchai Boonto

September 10, 2024

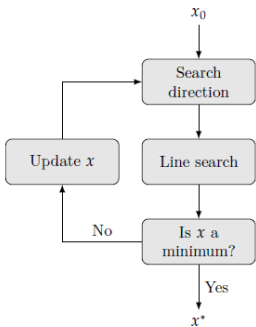
Department of Control System and Instrumentation Engineering
King Mongkut's University of Technology Thonburi
Thailand

Objective


At the end of this chapter you should be able to:

- Describe, implement, and use line-search-based methods.
- Explain the pros and cons of the various search direction methods.
- Understand steepest descent, conjugate gradient, etc.

Two Approaches to Finding and Optimum



Line search approach



Trust-region approach

Basic Concept

Consider a problem

$$\underset{\mathbf{x}}{\text{minimize}} \quad f(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^n$$

- Most numerical methods require a starting design or point which we call \mathbf{x}_0 .
- We then determine the *direction of travel* \mathbf{d}_0 .
- A *step size* α_0 is then determined based on minimizing f as much as possible and the design point is updated as $\mathbf{x}_1 = \mathbf{x}_0 + \alpha_0 \mathbf{d}_0$.
- The process of where to go and how far to go are repeated from \mathbf{x}_1 or $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{d}_k$.


Basic Concept: Example

Given $f(x_1, x_2) = x_1^2 + 5x_2^2$, a point $\mathbf{x}_0 = [3 \ 1]^T$, $f_0 = f(\mathbf{x}_0) = 14$.

1. construct $f(\alpha)$ along the direction $\mathbf{d} = [-3 \ -5]^T$ and provide a plot of $f(\alpha)$ versus α , for $\alpha \geq 0$.

We have $\mathbf{x}(\alpha) = \mathbf{x}_0 + \alpha\mathbf{d} = [3 - 3\alpha, 1 - 5\alpha]^T$ and

$$f(\alpha) = (3 - 3\alpha)^2 + 5(1 - 5\alpha)^2.$$



Basic Concept : Example

2. Find the slope $df(\alpha)/d\alpha$ at $\alpha = 0$. Verify that this equal $\nabla f(\mathbf{x}_0)^T \mathbf{d}$.

We have

$$\left. \frac{df(\alpha)}{d\alpha} \right|_{\alpha=0} = (-6(3 - 3\alpha) - 50(1 - 5\alpha))|_{\alpha=0} = -68$$

$$\nabla f(\mathbf{x}_0)^T \mathbf{d} = [2(3) \ 10(1)] \begin{bmatrix} -3 \\ -5 \end{bmatrix} = -68$$

3. Minimize $f(\alpha)$ with respect to α , to obtain step size α_0 . Given the corresponding new point \mathbf{x}_1 and value of $f_1 = f(\mathbf{x}_1)$


We have

$$\frac{df(\alpha)}{d\alpha} = -6(3 - 3\alpha) - 50(1 - 5\alpha) = 0 \implies 268\alpha = 68 \text{ or } \alpha = 0.2537$$


$$\mathbf{x}_1 = \begin{bmatrix} 3 \\ 1 \end{bmatrix} + \alpha_0 \begin{bmatrix} -3 \\ -5 \end{bmatrix} = \begin{bmatrix} 2.2388 \\ -0.2687 \end{bmatrix}, \quad f(\mathbf{x}_1) = 5.3732, \text{ less than } f_0 = 14.$$

Basic Concept : Example

2. Provide a plot showing contours of the function, steepest descent direction x_0 and x_1 .



Basic Concept : Example



We want to design the width and height of the rectangular cross-section to increase the bending stress defined by

$$\sigma_0 = \frac{6M}{wh^2}, \text{ where } M \text{ is a moment.}$$

With the initial design $(w, h) = (1, 3)$, we have

$$\sigma_0 = \frac{6(2000 \times 24)}{1(3^2)} = 32,000 \text{ psi}$$

Using $\mathbf{d} = [-1/\sqrt{5} \quad -2/\sqrt{5}]^T$ and $\alpha = 0.2$ we have

$$x_1 = x_0 + \alpha \mathbf{d} = \begin{bmatrix} 1 \\ 3 \end{bmatrix} + 0.2 \begin{bmatrix} -1/\sqrt{5} \\ -2/\sqrt{5} \end{bmatrix} = \begin{bmatrix} 0.9106 \\ 2.8211 \end{bmatrix}, \quad \sigma_1 = 71,342 \text{ psi}$$

Line Search : Exact Line Search

Assume we have chosen a descent direction \mathbf{d} . We need to choose the step factor α to obtain our next design point. One approach is to use *line search*, which selects the step factor that minimizes the one-dimensional function:

$$\underset{\alpha}{\text{minimize}} \quad f(\mathbf{x} + \alpha\mathbf{d})$$

To inform the search, we can use the derivative of the line search objective, which is simply the directional derivative along \mathbf{d} at $\mathbf{x} + \alpha\mathbf{d}$.

```
function LINE_SEARCH(f, d)  
    objective =  $\alpha \rightarrow f(\mathbf{x} + \alpha * \mathbf{d})$   
    a, b = bracket_minimum(objective)  
     $\alpha$  = minimize(objective, a, b)  
    return  $\mathbf{x} + \alpha * \mathbf{d}$   
end function
```

The exact line search is expensive, if we need to do it every step of the optimization. In Matlab environment, we can use commands `fminbnd` or `fminsearch`.

Line Search : Exact Line Search

- One disadvantage of conducting a line search at each step is the computational cost of optimizing α to a high degree of precision.
- We could quickly find a reasonable value and then move on, selecting \mathbf{x}_{k+1} , and then picking a new direction \mathbf{d}_{k+1} .
- Some algorithms use a fixed step factor. Large steps will tend to result in fast convergence but risk overshooting the minimum.
- Smaller steps tend to be more stable but can result in slower convergence.
- A fixed step factor α is sometimes referred to as a **learning rate**.
- Another method is to use a **decaying step factor**:

$$\alpha_k = \alpha_1 \gamma_{k-1} \text{ for } \gamma \in (0, 1]$$

The decaying step factors are popular when minimizing noisy objective function, and always used in machine learning applications.

Line Search : Exact Line Search

Consider conducting a line search on $f(x_1, x_2, x_3) = \sin(x_1 x_2) + e^{(x_2 + x_3)} - x_3$ from $\mathbf{x} = [1, 2, 3]$ in the direction $\mathbf{d} = [0, -1, -1]$. The corresponding optimization problem is:

$$\begin{aligned} \underset{\alpha}{\text{minimize}} \quad & \sin((1 + 0\alpha)(2 - \alpha)) \\ & + e^{((2 - \alpha) + (3 - \alpha))} - (3 - \alpha) \end{aligned}$$


which simplifies to:

$$\underset{\alpha}{\text{minimize}} \quad \sin(2 - \alpha) + e^{(5 - 2\alpha)} + \alpha - 3$$

The minimum is at $\alpha \approx 3.127$ with $\mathbf{x} \approx [1, -1.126, -1.126]$. **Note I:**

$\nabla f(\alpha) = -\cos(2 - \alpha) - 2e^{(5 - 2\alpha)} + 1 = 0$. We can solve for α by using `vpsolve` in **Matlab**.

Note II: We can use Nonlinear search in Matlab or Julia like `fminbnd` from the original problem.



Line Search : Exact Line Search

Nonlinear optimization using `fminbnd`:

```
f = @(alpha) sin(2-alpha) + exp(5-2*alpha) + alpha - 3  
alpha_1 = fminbnd(f, 0, 4, opts)
```

Solve for zero gradient `vpasolve`:

```
syms alpha  
f1 = -cos(2-alpha) - 2*exp(5-2*alpha) + 1;  
alpha_1 = vpasolve(f1==0,alpha)
```

We should have the same result around 3.127.


Approximate Line Search

- It is often more computationally efficient to perform more iterations of a descent method than to do an exact line search at each iteration, especially if the function and derivative calculations are expensive.
- Many methods discussed so far can benefit from using **approximate line search** to find a suitable step size with a small number of evaluations.
- Since descent methods must descend, a step size α may be suitable if it causes a decrease in the objective function value. We need the *sufficient decrease* condition. (to protect that the reductions in f values is not too small.)
- The sufficient decrease in the objective function value:

$$f(\mathbf{x}_{k+1}) \leq f(\mathbf{x}_k) + \beta\alpha\nabla_{\mathbf{d}_k} f(\mathbf{x}_k)$$

with $\beta \in [0, 1]$ often set to $\beta = 1 \times 10^{-4}$.

Approximate Line Search



- If $\beta = 0$, then any decrease is acceptable. If $\beta = 1$, then the decrease has to be at least as much as what would be predicted by a first-order approximation.
- If \mathbf{d} is a valid descent direction, then there must exist a sufficiently small step size that satisfies the sufficient decrease condition.
- We can start with a large step size and decrease it by a constant reduction factor until the sufficient decrease condition is satisfied.
- The algorithm is known as *backtracking line search* because of how it backtracks along the descent direction.

Approximate Line Search

```
function BACKTRACKING_LINE_SEARCH( $f, \nabla f, x, d, \alpha; p = 0.5, \beta = 1e - 4$ )  
  while  $f(\mathbf{x} + \alpha * \mathbf{d}) > f(\mathbf{x}) + \beta\alpha(\nabla f^T(\mathbf{x})\mathbf{d})$  do  
     $\alpha = p\alpha$   
  end while  
  return  $\alpha$   
end function
```

- The first condition is insufficient to guarantee convergence to a local minimum. Very small step sizes will satisfy the first condition but can prematurely converge.
- Backtracking line search avoids premature convergence by accepting the largest satisfactory step size obtained by sequential downscaling and is guaranteed to converge to a local minimum.

Curvature Condition

The *curvature condition* requires the the directional derivative at the next iterate to be shallower (α is not too close to zero):



$$\nabla_{\mathbf{d}_k} f(\mathbf{x}_{k+1}) \geq \sigma \nabla_{\mathbf{d}_k} f(\mathbf{x}_k)$$

- Where σ controls how shallow the next directional derivative must be.
- It is common to set $\beta < \sigma < 1$ with $\sigma = 0.1$ when approximate linear search is used with the conjugate gradient method and to 0.9 when used with Newton's method.
- The *strong curvature condition*, which is more restrictive criterion in that is also required not to be too positive:

$$|\nabla_{\mathbf{d}_k} f(\mathbf{x}_{k+1})| \leq -\sigma \nabla_{\mathbf{d}_k} f(\mathbf{x}_k)$$

- Both sufficient decrease condition (for α_U) and strong curvature condition are called **strong Wolfe conditions**.(for α_L).

Curvature Condition



Wolfe Condition

Consider approximate line search on $f(x_1, x_2) = x_1^2 + x_1x_2 + x_2^2$ from $\mathbf{x} = [1, 2]$ in the direction $\mathbf{d} = [-1, -1]$, using a maximum step size of 10, a reduction factor of 0.5, a first Wolfe condition parameter $\beta = 1 \times 10^{-4}$ and a second Wolfe condition parameter $\sigma = 0.9$.

The first Wolfe condition is $f(\mathbf{x} + \alpha\mathbf{d}) \leq f(x) + \beta\alpha(\mathbf{g}^T \mathbf{d})$, where $\mathbf{g} = \nabla f(\mathbf{x}) = [4, 5]$. $\alpha = 10$ we have

$$f\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 10 \begin{bmatrix} -1 \\ -1 \end{bmatrix}\right) \leq 7 + 1 \times 10^{-4}(10) \left(\begin{bmatrix} 4 & 5 \end{bmatrix} \begin{bmatrix} -1 \\ -1 \end{bmatrix} \right)$$

$217 \leq 6.991$ (It is not satisfied.)

$\alpha = 0.5(10) = 5$ we have

$$f\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 5 \begin{bmatrix} -1 \\ -1 \end{bmatrix}\right) \leq 7 + 1 \times 10^{-4}(5) \begin{bmatrix} 4 & 5 \end{bmatrix} \begin{bmatrix} -1 \\ -1 \end{bmatrix}$$

$37 \leq 6.996$ (It is not satisfied.)

Wolfe Condition

$\alpha = 0.5(5) = 2.5$, we have

$$f\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix} + 2.5 \begin{bmatrix} -1 \\ -1 \end{bmatrix}\right) \leq 7 + 1 \times 10^{-4}(2.5) \begin{bmatrix} 4 & 5 \end{bmatrix}^T \begin{bmatrix} -1 \\ -1 \end{bmatrix}$$

$3.25 \leq 6.998$ (The first Wolfe condition is satisfied.)

The candidate design point $\mathbf{x}' = \mathbf{x} + \alpha\mathbf{d} = [-1.5, -0.5]^T$ is checked against the second Wolfe condition:

$$\nabla_{\mathbf{d}}f(\mathbf{x}') \geq \sigma \nabla_{\mathbf{d}}f(\mathbf{x})$$
$$\begin{bmatrix} -3.5 & -2.5 \end{bmatrix} \begin{bmatrix} -1 \\ -1 \end{bmatrix} \geq \sigma \begin{bmatrix} 4 & 5 \end{bmatrix} \begin{bmatrix} -1 \\ -1 \end{bmatrix}$$

$6 \geq -8.1$ (The second Wolfe condition is satisfied.)

Approximate line search terminates with $\mathbf{x} = [-1.5 \ -0.5]^T$.

The Steepest Descent Method

The steepest-descent method (also called *gradient descent*) is a simple and intuitive method for determining the search direction.

Direction Vector:

- Let \mathbf{x}_k be the current point at the k th iteration: $k = 0$ corresponds to the starting point.
- We need to choose a downhill direction \mathbf{d} and then a step size $\alpha > 0$ such that the new point $\mathbf{x}_k + \alpha\mathbf{d}$ is better. We desire $f(\mathbf{x}_k + \alpha\mathbf{d}) < f(\mathbf{x}_k)$.
- To see how \mathbf{d} should be chosen, we use the Taylor's expansion

$$\begin{aligned}f(\mathbf{x}_k + \alpha\mathbf{d}) &= f(\mathbf{x}_k) + \alpha\nabla f(\mathbf{x}_k)^T\mathbf{d} + O(\alpha^2) \\ \delta f &= \alpha\nabla f(\mathbf{x}_k)^T\mathbf{d} + O(\alpha^2)\end{aligned}$$

- For small enough α the term $O(\alpha^2)$ is dominated. Consequently, we have

$$\delta f \approx \alpha\nabla f(\mathbf{x}_k)^T\mathbf{d}$$

The Steepest Descent Method

- For a reduction in f or $\delta f < 0$, we require \mathbf{d} to be a *descent direction* or a direction that satisfies

$$\nabla f(\mathbf{x}_k)^T \mathbf{d} < 0$$

- The steepest descent method is based on choosing \mathbf{d} at the k th iteration, which we will denote as \mathbf{d}_k , as

$$\mathbf{d}_k = -\nabla f(\mathbf{x}_k) \implies \nabla f(\mathbf{x}_k)^T (-\nabla f(\mathbf{x}_k)) = -\|\nabla f(\mathbf{x}_k)\|^2 < 0$$

- This direction will be referred to as the **steepest descent direction**.

The Steepest Descent Method : Example

Given $f(\mathbf{x}) = x_1x_2^2$, $\mathbf{x}_0 = [1 \ 2]^T$. The gradient is $\nabla f(\mathbf{x}) = [x_2^2 \ 2x_1x_2]^T$.

1. Find the steepest descent direction at \mathbf{x}_0

$$\mathbf{d} = -\nabla f(\mathbf{x})|_{\mathbf{x}=\mathbf{x}_0} = [-4 \ -4]^T$$

$$\mathbf{d}(\text{normalized}) = \left[-\frac{1}{\sqrt{2}} \ -\frac{1}{\sqrt{2}}\right]^T$$

2. Is $\mathbf{d} = [-1 \ 2]^T$ a direction of descent?

$$\nabla f(\mathbf{x})^T|_{\mathbf{x}=\mathbf{x}_0} \mathbf{d} = [4 \ 4] \begin{bmatrix} -1 \\ 2 \end{bmatrix} = 4 > 0$$

It is not a descent direction.

The Steepest Descent Method

After we have the direction vector \mathbf{d}_k at the point \mathbf{x}_k , how far to go along this direction?

- We need to develop a numerical procedure to determine the step size α_k along \mathbf{d}_k .
- If we move along \mathbf{d}_k the design variables and the objective function depend only on α as

$$\begin{aligned}\mathbf{x}(\alpha) &= \mathbf{x}_k + \alpha \mathbf{d}_k, & f(\alpha) &= f(\mathbf{x}_k + \alpha \mathbf{d}_k) \\ \alpha_k &= \arg \underset{\alpha}{\text{minimize}} f(\mathbf{x}_k + \alpha \mathbf{d}_k)\end{aligned}$$

- The optimization above implies that the directional derivative equals zero. We have

$$\begin{aligned}\nabla f(\mathbf{x}_k + \alpha \mathbf{d}_k)^T \mathbf{d}_k &= 0, & \mathbf{d}_{k+1} &= -\frac{\nabla f(\mathbf{x}_k + \alpha \mathbf{d}_k)}{\|\nabla f(\mathbf{x}_k + \alpha \mathbf{d}_k)\|} \\ \mathbf{d}_{k+1}^T \mathbf{d}_k &= 0, & \mathbf{d}_{k+1} &\text{ and } \mathbf{d}_k \text{ are orthogonal.}\end{aligned}$$

The Steepest Descent Method

- In the steepest descent method, the direction vector is $-\nabla f(\mathbf{x}_k)$ resulting in the slope at the current point $\alpha = 0$ being

$$\left. \frac{df(\alpha)}{d\alpha} \right|_{\alpha=0} = \nabla f(\mathbf{x}_k)^T (-\nabla f(\mathbf{x}_k)) = -\|\nabla f(\mathbf{x}_k)\|^2 < 0$$

Implying a move in a downhill direction.

The Steepest Descent Method

- Starting from an initial point, we determine a direction vector and a step size, and obtain a new point as $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{d}_k$.
- The question is to know when to stop the iterative process. We have two stop criteria to discuss here.
- **First** Befor performing line search, the necessary condition for optimality is checked:

$$\|\nabla f(\mathbf{x}_k)\| \leq \epsilon_G,$$

where ϵ_G is a tolerance on the gradient and is supplied by the user. If the condition is satisfied the the process is terminated.

The Steepest Descent Method : Stopping Criteria

- **Second:** We check the successive reductions in f as a criterion for stopping.

$$|f(\mathbf{x}_{k+1}) - f(\mathbf{x}_k)| \leq \varepsilon_A + \varepsilon_R |f(\mathbf{x}_k)|$$

where ε_A = absolute tolerance on the change in function value and ε_R = relative tolerance. Only if the condition is satisfied for two consecutive iterations is the descen process stopped.

Steepest Descent Algorithm : Algorithm

Require: $\mathbf{x}_0, \varepsilon_G, \varepsilon_A, \varepsilon_R$

$k = 0$

while true **do**

 Compute $\nabla f(\mathbf{x}_k)$

if $\|\nabla f(\mathbf{x}_k)\| \leq \varepsilon_G$ **then**

 Stop

else if **then**

$\mathbf{d}_k = -\nabla f(\mathbf{x}_k) / \|\nabla f(\mathbf{x}_k)\|$

end if

$\alpha_k = \text{line_search}(f, \mathbf{d}_k)$,

$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{d}_k$,

if $|f(\mathbf{x}_{k+1}) - f(\mathbf{x}_k)| \leq \varepsilon_A + \varepsilon_R |f(\mathbf{x}_k)|$ **then**

 Stop

else

$k = k + 1, \mathbf{x}_k = \mathbf{x}_{k+1}$

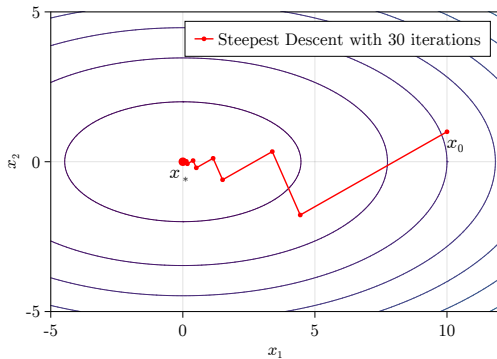
end if

end while

Steepest Descent Algorithm: Zig-Zags Property

- The steepest descent method zig-zags its way towards the optimum point. Consider

$$f(x_1, x_2) = x_1^2 + 5x_2^2.$$



Steepest Descent Algorithm: Zig-Zags Property

From the fact that α_k is obtained by minimizing $f(\mathbf{x}_k + \alpha \mathbf{d}_k)$. Thus

$$\begin{aligned}\frac{\partial f(\mathbf{x}_k + \alpha \mathbf{d}_k)}{\partial \alpha} &= 0 \\ \frac{\partial f(\mathbf{x}_{k+1})}{\partial \alpha} &= \frac{\partial f(\mathbf{x}_{k+1})}{\partial \mathbf{x}_{k+1}} \frac{\partial \mathbf{x}_{k+1}}{\partial \alpha} = \frac{\partial f(\mathbf{x}_{k+1})}{\partial \mathbf{x}_{k+1}} \frac{\partial (\mathbf{x}_k + \alpha \mathbf{d}_k)}{\partial \alpha} = 0 \\ \nabla f(\mathbf{x}_{k+1})^T \mathbf{d}_k &= 0\end{aligned}$$

Set $\mathbf{d} = -\nabla f_k$, we have

$$-\nabla f(\mathbf{x}_{k+1})^T \nabla f(\mathbf{x}_k) = 0$$


From the last line, it means the $k + 1$ direction is perpendicular to the k direction. If you use the approximation line search, this perpendicular property is lost, but the zig-zags are still there.

Steepest Descent Algorithm : The Bean function

Find the minimum of the bean function

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

using the steepest-descent algorithm with an exact line search, and a convergence tolerance of $\|\nabla f\| \leq 10^{-6}$.



$$\mathbf{x}^* = \begin{bmatrix} 1.2134 \\ 0.8241 \end{bmatrix}, f(\mathbf{x}^*) = 0.0919$$

Steepest Descent Algorithm : Convergence Characteristics

- The speed of convergence of the method is related to the spectral condition number of the Hessian matrix. The spectral condition number κ of a symmetric positive definite matrix A is defined as the ratio of the largest to the smallest eigenvalue, or




$$\kappa = \frac{\lambda_{\max}}{\lambda_{\min}}$$

- For well-conditioned Hessian matrices, the condition number is close to unity, contours are more circular, and the method is at its best.
- The higher the condition number, the more ill-conditioned is the Hessian, the contours are more elliptical, more is the amount of zig-zagging near as the optimum is approached, the smaller are the step sizes and thus the poorer is the rate of convergence.

Steepest Descent Algorithm : Convergence Characteristics

Consider a function $f(x_1, x_2) = x_1^2 + \beta x_2^2$, we have

$$\mathbf{H}(f) = \nabla^2 f = \begin{bmatrix} 2 & 0 \\ 0 & 2\beta \end{bmatrix}, \text{ with } \beta = 1, 5, 15$$



The Conjugate Gradient Method

- The Conjugate Gradient Method [Fletcher and Powell 1963] is a dramatic improve over the steepest descent method. The steepest descent perform poorly in narrow valleys.
- It can find the minimum of a quadratic function of n variables in n iterations.
- The conjugate Gradient method is also powerful on general functions.
- Consider the problem of minimizing a quadratic function

$$\underset{\mathbf{x}}{\text{minimize}} \quad q(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{A}\mathbf{x} + \mathbf{b}^T \mathbf{x} + c$$

where \mathbf{A} is a symmetric and positive definite.

- The conjugate directions, or directions that are mutually conjugate with respect to A , as vectors which satisfy

$$\mathbf{d}_i^T \mathbf{A} \mathbf{d}_j = 0, \quad i \neq j, 0 \leq i, j \leq n$$

The Conjugate Gradient Method : The method

- The mutually conjugate vectors are the basis vectors of \mathbf{A} . They are generally not orthogonal to one another.
- The algorithm is started with the direction of steepest descent:

$$\mathbf{d}_1 = -\mathbf{g}_1$$

- Use line search to find the next design point For quadratic functions, the step factor α can be computed exactly. The update is then:

$$\mathbf{x}_2 = \mathbf{x}_1 + \alpha_1 \mathbf{d}_1$$

The Conjugate Gradient Method : The method

- Suppose we want to derive the optimal step factor for a line search on a quadratic function:

$$\underset{\alpha}{\text{minimize}} \quad f(\mathbf{x} + \alpha \mathbf{d})$$

We have

$$\begin{aligned} \frac{\partial f(\mathbf{x} + \alpha \mathbf{d})}{\partial \alpha} &= \frac{\partial}{\partial \alpha} \left[\frac{1}{2} (\mathbf{x} + \alpha \mathbf{d})^T \mathbf{A} (\mathbf{x} + \alpha \mathbf{d}) + \mathbf{b}^T (\mathbf{x} + \alpha \mathbf{d}) + c \right] \\ &= \mathbf{d}^T \mathbf{A} (\mathbf{x} + \alpha \mathbf{d}) + \mathbf{b}^T \mathbf{d} = \mathbf{d}^T \mathbf{A} (\mathbf{x} + \alpha \mathbf{d}) + \mathbf{d}^T \mathbf{b} \\ &= \mathbf{d}^T (\mathbf{A} \mathbf{x} + \mathbf{b}) + \alpha \mathbf{d}^T \mathbf{A} \mathbf{d} \end{aligned}$$

Setting $\frac{\partial f(\mathbf{x} + \alpha \mathbf{d})}{\partial \alpha} = 0$ results in:

$$\alpha = - \frac{\mathbf{d}^T (\mathbf{A} \mathbf{x} + \mathbf{b})}{\mathbf{d}^T \mathbf{A} \mathbf{d}}$$

The Conjugate Gradient Method : The method

- Subsequent iterations choose \mathbf{d}_{k+1} based on the current gradient and a contribution from the previous descent direction:

$$\mathbf{d}_k = -\mathbf{g}_k + \beta_k \mathbf{d}_{k-1}$$

for scalar parameter β . Larger values of β indicate that the previous descent direction contributes more strongly.

- To find the best value for β for a known \mathbf{A} , using the fact that \mathbf{d}_k is conjugate to \mathbf{d}_{k-1} :

$$\begin{aligned}\mathbf{d}_k^T \mathbf{A} \mathbf{d}_{k-1} &= 0 \quad \implies (-\mathbf{g}_k + \beta_k \mathbf{d}_{k-1})^T \mathbf{A} \mathbf{d}_{k-1} = 0 \\ -\mathbf{g}_k^T \mathbf{A} \mathbf{d}_{k-1} + \beta_k \mathbf{d}_{k-1}^T \mathbf{A} \mathbf{d}_{k-1} &= 0 \quad \implies \beta_k = \frac{\mathbf{g}_k^T \mathbf{A} \mathbf{d}_{k-1}}{\mathbf{d}_{k-1}^T \mathbf{A} \mathbf{d}_{k-1}}\end{aligned}$$

- The conjugate gradient method can be applied to nonquadratic functions as well.

The Conjugate Gradient Method : The method

We do not know the value of A that best approximates f around \mathbf{x}_k . Several choices for β_k tend to work well:

- Fletcher-Reeves:

$$\beta_k = -\frac{\mathbf{g}_k^T \mathbf{g}_k}{\mathbf{g}_{k-1}^T \mathbf{g}_{k-1}}$$

- Polak-Ribière:

$$\beta_k = \frac{\mathbf{g}_k^T (\mathbf{g}_k - \mathbf{g}_{k-1})}{\mathbf{g}_{k-1}^T \mathbf{g}_{k-1}}$$

- Convergence for the Polak-Ribière method can be guaranteed if we modify it to allow for automatic resets:

$$\beta \leftarrow \max(\beta, 0)$$

The Conjugate Gradient Method : Example

Consider $f = x_1^2 + 4x_2^2$, $\mathbf{x}_0 = \begin{bmatrix} 1 & 1 \end{bmatrix}^T$. We will perform two iterations of the conjugate gradient algorithm. The first step is the steepest descent iteration. Thus

$$\mathbf{d}_0 = -\nabla f(\mathbf{x}_0) = -\begin{bmatrix} 2 & 8 \end{bmatrix}^T$$

Assuming the direction vectors are not normalized to be unit vectors,

$$f(\alpha) = f(\mathbf{x}_0 + \alpha\mathbf{d}_0) = (1 - 2\alpha)^2 + 4(1 - 8\alpha)^2$$

which yields $\alpha_0 = 0.1308$, $\mathbf{x}_1 = \mathbf{x}_0 + \alpha_0\mathbf{d}_0 = \begin{bmatrix} 0.7385 & -0.0462 \end{bmatrix}^T$. The next iteration (using Fletcher-Reeves method):

$$\beta_0 = \frac{\|\nabla f(\mathbf{x}_1)\|^2}{\|\nabla f(\mathbf{x}_0)\|^2} = 2.3176/68 = 0.0341$$

The Conjugate Gradient Method : Example

$$\mathbf{d}_1 = -\nabla f^T(\mathbf{x}_1) + \beta_0 \mathbf{d}_0 = \begin{bmatrix} -1.4770 \\ 0.3692 \end{bmatrix} + 0.0341 \begin{bmatrix} -2 \\ -8 \end{bmatrix} = \begin{bmatrix} -1.5451 \\ 0.0966 \end{bmatrix}$$

We have

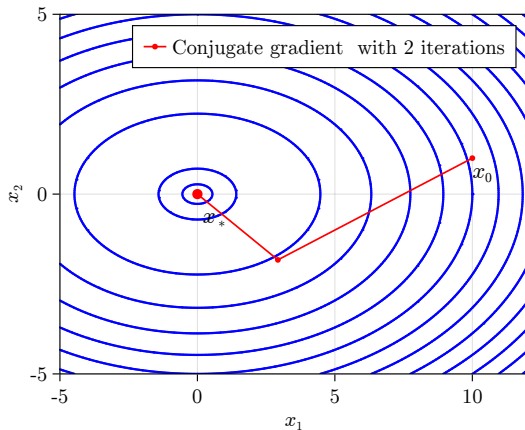
$$f(\alpha) = f(\mathbf{x}_1 + \alpha \mathbf{d}_1) = (0.7385 - 1.5451\alpha)^2 + 4(-0.0462 + 0.0966\alpha)^2$$

which yields

$$\alpha_1 = 0.4780$$

$$\mathbf{x}_2 = \mathbf{x}_1 + \alpha_1 \mathbf{d}_1 = \begin{bmatrix} 0.7385 \\ -0.0462 \end{bmatrix} + 0.4780 \begin{bmatrix} -1.5451 \\ 0.0966 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

The Conjugate Gradient Method : Example



The Conjugate Gradient Method : Algorithm

Require: $\mathbf{x}_0, \varepsilon_G$

$k = 0$

while $\|\nabla f_k\| > \varepsilon_G$ do

 if $k = 0$ then

$$\mathbf{d}_k = -\frac{\nabla f_x}{\|\nabla f_k\|}$$

 else

$$\beta_k = \frac{\nabla f_x^T \nabla f_k}{\nabla f_{k-1}^T \nabla f_{k-1}}$$

$$\mathbf{d}_k = -\frac{\nabla f_k}{\|\nabla f_k\|} + \beta_k \mathbf{d}_{k-1}$$

 end if

$\alpha_k = \text{line_search}(f, \mathbf{d}_k)$

$\mathbf{x}_k = \mathbf{x}_{k-1} + \alpha_k \mathbf{d}_k$


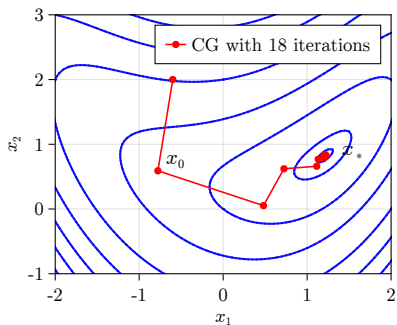
$k = k + 1$

end while

The Conjugate Gradient Method : Example

The minimum of the bean function,

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



Newton's Method

- The function value and gradient can help to determine the direction to travel, but it does not directly help to determine how far to step to reach a local minimum.
- Second-order information allows us to make a quadratic approximation of the objective function and approximate the right step size to reach a local minimum.
- As we have seen with a quadratic fit search, we can analytically obtain the location where a quadratic approximation has a zero gradient. We can use that location as the next iteration to approach a local minimum.
- The quadratic approximation about a point \mathbf{x}_k comes from the second-order Taylor expansion:



$$q(\mathbf{x}) = f(\mathbf{x}_k) + (\mathbf{x} - \mathbf{x}_k)f'(\mathbf{x}_k) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_k)^2 f''(\mathbf{x}_k)$$
$$\frac{\partial}{\partial \mathbf{x}} q(\mathbf{x}) = f'(\mathbf{x}_k) + (\mathbf{x} - \mathbf{x}_k)f''(\mathbf{x}_k) = 0$$
$$\mathbf{x}_{k+1} = \mathbf{x}_k - \frac{f'(\mathbf{x}_k)}{f''(\mathbf{x}_k)}$$

Newton's Method : Example

Suppose we want to minimize the following single-variable function:

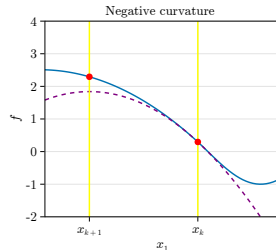
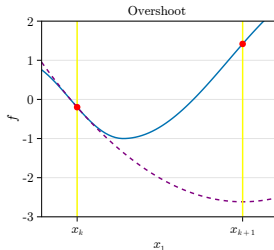

$$f(x) = (x - 2)^4 + 2x^2 - 4x + 4, \quad f'(x) = 4(x - 2)^3 + 4x - 4, \\ f''(x) = 12(x - 2)^2 + 4$$

with $x_0 = 3$, we can form the quadratic using the function value and the first and second derivatives evaluated at the point.



Newton's Method : Disadvantage

- The update rule in Newton's method involves dividing by the second derivative. The update is undefined if the second derivative is zero, which occurs when the quadratic approximation is a horizontal line.
- Instability also occurs when the second derivative is very close to zero, in which case the next iterate will lie very far from the current design point, far from where the local quadratic approximation is valid.
- Poor local approximations can lead to poor performance with Newton's method.



Newton's Method : Multivariate Optimization

- The multivariate second-order Taylor expansion at \mathbf{x}_k is

$$f(\mathbf{x}) \approx q(\mathbf{x}) = f(\mathbf{x}_k) + \nabla f(\mathbf{x}_k)^T (\mathbf{x} - \mathbf{x}_k) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_k)^T \mathbf{H}_k (\mathbf{x} - \mathbf{x}_k)$$
$$\frac{dq(\mathbf{x}_k)}{d(\mathbf{x} - \mathbf{x}_k)} = \nabla f(\mathbf{x}_k) + \mathbf{H}_k (\mathbf{x} - \mathbf{x}_k) = 0$$

We then solve for the next iterate, thereby obtaining Newton's method in multivariate form:

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

- If $f(\mathbf{x})$ is quadratic and its Hessian is positive definite, then the update converges to the global minimum in one step. For general functions, Newton's method is often terminated once \mathbf{x} ceases to change by more than a given tolerance.

Newton's Method : Example

With $\mathbf{x}_1 = \begin{bmatrix} 9 & 8 \end{bmatrix}$, we will use Newton's method to minimize Booth's function:

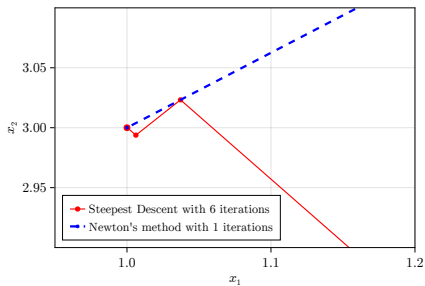

$$f(\mathbf{x}) = (x_1 + 2x_2 - 7)^2 + (2x_1 + x_2 - 5)^2,$$
$$\nabla f(\mathbf{x}) = \begin{bmatrix} 10x_1 + 8x_2 - 34, & 8x_1 + 10x_2 - 38 \end{bmatrix}^T, \quad \mathbf{H}(\mathbf{x}) = \begin{bmatrix} 10 & 8 \\ 8 & 10 \end{bmatrix}$$

The first iteration of Newton's method yields:

$$\begin{aligned} \mathbf{x}_2 &= \mathbf{x}_1 - \mathbf{H}_1^{-1} \mathbf{g}_1 = \begin{bmatrix} 9 \\ 8 \end{bmatrix} - \begin{bmatrix} 10 & 8 \\ 8 & 10 \end{bmatrix}^{-1} \begin{bmatrix} 10(9) + 8(8) - 34 \\ 8(9) + 10(8) - 38 \end{bmatrix} \\ &= \begin{bmatrix} 9 \\ 8 \end{bmatrix} - \begin{bmatrix} 10 & 8 \\ 8 & 10 \end{bmatrix}^{-1} \begin{bmatrix} 120 \\ 114 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix} \end{aligned}$$

The gradient at \mathbf{x}_2 is zero, so we have converged after a single iteration. The Hessian is positive definite everywhere, so \mathbf{x}_2 is the global minimum.

Newton's Method : Multivariate Optimization



Newton's Method : Algorithm

Require: $\mathbf{x}_0, \varepsilon_G, \nabla f_k, \mathbf{H}_k$

$k = 0$

while $\|\nabla f_k\| > \varepsilon_G$ and $k \leq k_{\max}$ do

$\Delta = \mathbf{H}(\mathbf{x})^{-1} \nabla f(\mathbf{x})$

$\mathbf{x} = \mathbf{x} + \Delta$

$k = k + 1$

end while

return \mathbf{x}

Secant Method

Newton's method is efficient because the second-order information results in better search directions, but it has the significant shortcoming of requiring the Hessian. Quasi-Newton methods are designed to address this issue. The basic idea is that we can use first-order information (gradients) along each step in the iteration path to building an approximation of the Hessian.

- The *secant method* uses the last two iterates to approximate the second derivative:

$$f''(\mathbf{x}_k) \approx \frac{f'(\mathbf{x}_k) - f'(\mathbf{x}_{k-1})}{\mathbf{x}_k - \mathbf{x}_{k-1}}$$

- This estimate is substituted into Newton's method:

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

- The secant method requires an additional initial design point. It suffers from the same problems as Newton's method.

Quasi-Newton Methods

- As the secant method approximates f'' in the univariate case, *quasi-Newton* methods approximate the inverse Hessian. Quasi-Newton method updates have the form:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \mathbf{Q}_k \nabla f_k,$$

where α_k is a scalar step factor and \mathbf{Q}_k approximates the inverse of the Hessian at \mathbf{x}_k

- These methods typically set \mathbf{Q}_1 to the identity matrix, and they then apply updates to reflect information learned with each iteration. To simplify the equations for the various quasi-Newton methods, we define the following:

$$\boldsymbol{\gamma}_{k+1} = \nabla f_{k+1} - \nabla f_k$$

$$\boldsymbol{\delta}_{k+1} = \mathbf{x}_{k+1} - \mathbf{x}_k$$

Quasi-Newton Methods : Davidon-Fletcher-Powell (DFP)

The *Davidon-Fletcher-Powell* (DFP) method uses:

$$\mathbf{Q}_{k+1} = \mathbf{Q}_k - \frac{\mathbf{Q}_k \boldsymbol{\gamma}_k \boldsymbol{\gamma}_k^T \mathbf{Q}_k}{\boldsymbol{\gamma}_k^T \mathbf{Q}_k \boldsymbol{\gamma}_k} + \frac{\boldsymbol{\delta}_k \boldsymbol{\delta}_k^T}{\boldsymbol{\delta}_k^T \boldsymbol{\gamma}_k}$$

The update for \mathbf{Q} in the DFP method has three properties:

- \mathbf{Q} remains symmetric and positive definite.
- If $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c$, then $\mathbf{Q} = \mathbf{A}^{-1}$. Thus the DFP has the same convergence properties as the conjugate gradient method.
- For high-dimensional problems, storing and updating \mathbf{Q} can be significant compared to other methods like the conjugate gradient method.

Quasi-Newton Methods : Davidon-Fletcher-Powell (DFP)

Require: $\mathbf{x}_0, \varepsilon_G, \nabla f_k$

$k = 0, \mathbf{Q} = I$

while $\|\nabla f_k\| > \varepsilon_G$ && $k \leq k_{\max}$ do

$\mathbf{g} = \nabla f(\mathbf{x})$

$\mathbf{x}' = \text{line_search}(f, \mathbf{x}, -\mathbf{Q} * \mathbf{g})$

$\mathbf{g}' = \nabla f(\mathbf{x}')$

$\delta = \mathbf{x}' - \mathbf{x}$

$\gamma = \mathbf{g}' - \mathbf{g}$

$\mathbf{Q} = \mathbf{Q} - \mathbf{Q}\gamma\gamma^T\mathbf{Q}/\gamma^T\mathbf{Q}\gamma + \delta\delta^T/\delta^T\gamma$

 return \mathbf{x}'

end while

Quasi-Newton Methods : Broyden-Fletcher-Goldfarb-Shanno (BFGS)

An alternative to DFP, the *Broyden-Fletcher-Goldfarb-Shanno* (BFGS) method uses:

$$\mathbf{Q}_{k+1} = \mathbf{Q}_k - \left(\frac{\delta_k \gamma_k^T \mathbf{Q}_k + \mathbf{Q}_k \gamma_k \delta_k^T}{\delta_k^T \gamma_k} \right) + \left(1 + \frac{\gamma_k^T \mathbf{Q}_k \gamma_k}{\delta_k^T \gamma_k} \right) \frac{\delta_k \delta_k^T}{\delta_k^T \gamma_k}$$

Require: $\mathbf{x}_0, \varepsilon_G, \nabla f_k$

$k = 0, \mathbf{Q} = I$

while $\|\nabla f_k\| > \varepsilon_G$ && $k \leq k_{\max}$ do

$\mathbf{g} = \nabla f(\mathbf{x})$

$\mathbf{x}' = \text{line_search}(f, \mathbf{x}, -\mathbf{Q} * \mathbf{g})$

$\mathbf{g}' = \nabla f(\mathbf{x}')$

$\delta = \mathbf{x}' - \mathbf{x}$

$\gamma = \mathbf{g}' - \mathbf{g}$


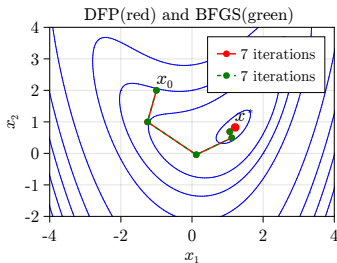

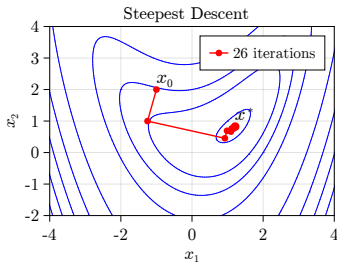
$\mathbf{Q} = \mathbf{Q} - (\delta \gamma^T \mathbf{Q} + \mathbf{Q} \gamma \delta^T / \delta^T \gamma) + (1 + \gamma^T \mathbf{Q} \gamma / \delta^T \gamma) \delta \delta^T / \delta^T \gamma$

 return \mathbf{x}'

end while


Compare Four Methods

Minimizing of bean function



Compare Four Methods

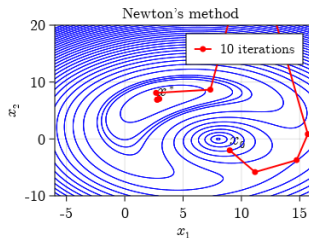
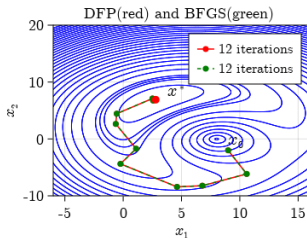


Minimizing the total potential energy for a spring system:



$$\text{minimize}_{x_1, x_2} \frac{1}{2} k_1 \left(\sqrt{(l_1 + x_1)^2 + x_2^2} - l_1 \right)^2 + \frac{1}{2} k_2 \left(\sqrt{(l_2 - x_1)^2 + x_2^2} - l_2 \right)^2 - mgx_2$$

By letting $l_1 = 12$, $l_2 = 8$, $k_1 = 1$, $k_2 = 10$, $mg = 7$ (with appropriate units).

Compare Four Methods



1. Joaquim R. R. A. Martins, Andrew Ning, "*Engineering Design Optimization*," Cambridge University Press, 2021.
2. Mykel J. Kochenderfer, and Tim A. Wheeler, "*Algorithms for Optimization*," The MIT Press, 2019.
3. Ashok D. Belegundu, Tirupathi R. Chandrupatla, "*Optimization Concepts and Applications in Engineering*," Cambridge University Press, 2019.
4. Kalyanmoy D., "*Optimization for Engineering Design: Algorithms and Examples*," 2nd, PHI Learning Private Limited, 2012.