4.4.4 Quasi-Newton Methods

The basic idea of quasi-Newton methods is to approximate the Hessian matrix (or its inverse) which we need to compute in the Newton method. There are of course infinitely many ways to do so, but we choose the ones which satisfy the *secant equation*:

$$\boldsymbol{H}_{k+1}\boldsymbol{y}_k = \boldsymbol{s}_k$$

where $\boldsymbol{y}_k = \boldsymbol{\nabla f}(\boldsymbol{x}_{k+1}) - \boldsymbol{\nabla f}(\boldsymbol{x}_k)$, $\boldsymbol{s}_k = \boldsymbol{x}_{k+1} - \boldsymbol{x}_k$.

The general scheme of the quasi-Newton method is as follows.

Quasi-Newton Method	
Step 0:	Let $\boldsymbol{x}_0 \in \mathbb{R}^n$, $\boldsymbol{H}_0 := \boldsymbol{I}$, $k := 0$. Compute $f(\boldsymbol{x}_0), \boldsymbol{\nabla f}(\boldsymbol{x}_0)$
Step 1:	Set $\boldsymbol{p}_k := \boldsymbol{H}_k \boldsymbol{\nabla f}(\boldsymbol{x}_k)$
Step 2:	Find $\boldsymbol{x}_{k+1} := \boldsymbol{x}_k - h_k \boldsymbol{p}_k$ by "approximate line search" on the scalar h_k
Step 3:	Compute $f(\boldsymbol{x}_{k+1})$ and $\boldsymbol{\nabla f}(\boldsymbol{x}_{k+1})$
Step 4:	Compute $\boldsymbol{H}_{k+1}$ from $\boldsymbol{H}_k$, $k := k+1$ and go to Step 1

The most popular updates for $\boldsymbol{H}_{k+1}$ are:

1. *BFGS (Broyden-Fletcher-Goldfarb-Shanno)*

$$\boldsymbol{H}_{k+1} := \left(\boldsymbol{I} - \frac{\boldsymbol{s}_k(\boldsymbol{y}_k)^T}{\langle \boldsymbol{s}_k, \boldsymbol{y}_k \rangle}\right) \boldsymbol{H}_k \left(\boldsymbol{I} - \frac{\boldsymbol{y}_k(\boldsymbol{s}_k)^T}{\langle \boldsymbol{s}_k, \boldsymbol{y}_k \rangle}\right) + \frac{\boldsymbol{s}_k(\boldsymbol{s}_k)^T}{\langle \boldsymbol{s}_k, \boldsymbol{y}_k \rangle}$$

2. *DFP (Davidon-Fletcher-Powell)*

$$\boldsymbol{H}_{k+1} := \boldsymbol{H}_k + \frac{\boldsymbol{s}_k(\boldsymbol{s}_k)^T}{\langle \boldsymbol{y}_k, \boldsymbol{s}_k \rangle} - \frac{\boldsymbol{H}_k \boldsymbol{y}_k (\boldsymbol{y}_k)^T \boldsymbol{H}_k}{\langle \boldsymbol{y}_k, \boldsymbol{H}_k \boldsymbol{y}_k \rangle}$$

3. *Symmetric-Rank-One*

$$\boldsymbol{H}_{k+1} := \boldsymbol{H}_k + \frac{(\boldsymbol{s}_k - \boldsymbol{H}_k \boldsymbol{y}_k)(\boldsymbol{s}_k - \boldsymbol{H}_k \boldsymbol{y}_k)^T}{\langle \boldsymbol{s}_k - \boldsymbol{H}_k \boldsymbol{y}_k, \boldsymbol{y}_k \rangle}$$

In the same way for the conjugate gradient method, we can show that the quasi-Newton method converges in finite number of iterations for a strictly convex quadratic function. Moreover, under some strict convexity conditions at the neighborhood of the local minimum, it is possible to show that its iterates converge super-linearly [Nocedal].

4.5 Exercises

1. Let $f : \mathbb{R}^n \to \mathbb{R}$, $g : \mathbb{R}^n \to \mathbb{R}^m$ continuously differentiable functions and $\boldsymbol{h} \in \mathbb{R}^m$. Define the following optimization problem.

$$\begin{cases} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & g(\boldsymbol{x}) = \boldsymbol{h} \\ & \boldsymbol{x} \in \mathbb{R}^n \end{cases}$$

   Write the Karush-Kuhn-Tucker (KKT) conditions corresponding to the above problem.

2. In view of Theorem 4.13, find a twice continuously differentiable function on $\mathbb{R}^n$ which satisfies $\boldsymbol{\nabla f}(\boldsymbol{x}^*) = 0, \quad \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*) \succeq \boldsymbol{O}$, but $\boldsymbol{x}^*$ is not a local minimum of $f(\boldsymbol{x})$.

3. Let $f : \mathbb{R}^n \to \mathbb{R}$ be a continuous differentiable and convex function. If $\boldsymbol{x}^* \in \mathbb{R}^n$ is such that $\boldsymbol{\nabla f}(\boldsymbol{x}^*) = \boldsymbol{0}$, then show that $\boldsymbol{x}^*$ is a global minimum for $f(\boldsymbol{x})$.

4. Determine the Cauchy step-size $h_k \in \mathbb{R}$ for the following strictly convex quadratic function $f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^T\boldsymbol{Q}\boldsymbol{x} + \boldsymbol{q}^T\boldsymbol{x} + \gamma$, where $\boldsymbol{Q}$ is a $n \times n$ real positive definite matrix.

5. Give a geometric interpretation of the following step-size strategies:

   Let $0 < c_1 < c_2 < 1$,

   - Wolfe condition

   $$f(\boldsymbol{x}_k - h\boldsymbol{\nabla f}(\boldsymbol{x}_k)) \leq f(\boldsymbol{x}_k) - c_1 h \|\boldsymbol{\nabla f}(\boldsymbol{x}_k)\|_2^2,$$
   $$\langle \boldsymbol{\nabla f}(\boldsymbol{x}_k - h\boldsymbol{\nabla f}(\boldsymbol{x}_k)), \boldsymbol{\nabla f}(\boldsymbol{x}_k)\rangle \leq c_2 \|\boldsymbol{\nabla f}(\boldsymbol{x}_k)\|_2^2.$$

   - Strong Wolfe condition

   $$f(\boldsymbol{x}_k - h\boldsymbol{\nabla f}(\boldsymbol{x}_k)) \leq f(\boldsymbol{x}_k) - c_1 h \|\boldsymbol{\nabla f}(\boldsymbol{x}_k)\|_2^2,$$
   $$|\langle \boldsymbol{\nabla f}(\boldsymbol{x}_k - h\boldsymbol{\nabla f}(\boldsymbol{x}_k)), \boldsymbol{\nabla f}(\boldsymbol{x}_k)\rangle| \leq c_2 \|\boldsymbol{\nabla f}(\boldsymbol{x}_k)\|_2^2.$$

6. Consider a sequence $\{\beta_k\}_{k=0}^{\infty}$ which converges to zero.

   The sequence is said to converge *Q-linearly* if there exists a scalar $\rho \in (0, 1)$ such that

   $$\left|\frac{\beta_{k+1}}{\beta_k}\right| \leq \rho,$$

   for all k sufficiently large. *Q-superlinear* convergence occurs when we have

   $$\lim_{k\to\infty} \frac{\beta_{k+1}}{\beta_k} = 0,$$

   while the convergence is *Q-quadratic* if there is a constant C such that

   $$\frac{|\beta_{k+1}|}{\beta_k^2} \leq C$$

   for all k sufficiently large. *Q-superquadratic* convergence is indicated by

   $$\lim_{k\to\infty} \frac{\beta_{k+1}}{\beta_k^2} = 0.$$

   (a) Show that the following implications are valid: Q-superquadratic $\Rightarrow$ Q-quadratic $\Rightarrow$ Q-superlinear $\Rightarrow$ Q-linear.

   (b) Give examples of sequences which do not imply the opposite directions in the three cases above.

   A zero converging sequence $\{\beta_k\}_{k=0}^{\infty}$ is said to converge *R-linearly* if it is dominated by a Q-linearly converging sequence. That is, if there is a Q-linearly converging sequence $\{\hat{\beta}_k\}_{k=0}^{\infty}$ such that $0 \leq |\beta_k| \leq \hat{\beta}_k$.

   (c) Give a sequence which is R-linearly converging but not Q-linearly converging.

7. Let $f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^T\boldsymbol{Q}\boldsymbol{x}$ such that $\boldsymbol{Q}$ is symmetric, and indefinite. Apply the steepest descent method with constant step. Show that if the starting point $\boldsymbol{x}_0$ belongs to the space spanned by the negative eigenvectors, the sequence generated by the steepest descent method diverges.

8. Prove Lemma 4.26.

9. In light of Theorem 4.21, show that under Assumption 4.20, if we want to obtain $\|\boldsymbol{x}_k - \boldsymbol{x}^*\|_2 < \varepsilon$, we need an order of $\ln(\ln \varepsilon^{-1})$ iterations for the Newton method.

10. In the Section 4.4.3, show that $\mathcal{L}_k = \{\boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_{k-1}\}$.

11. In the same section, arrive at the expression (9) for a strictly convex quadratic function.

12. Show that the secant equation is valid for BFGS, DFP and symmetric-rank-one formulae.

13. Given $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$ and a non-singular matrix $\boldsymbol{M} \in \mathbb{R}^{n \times n}$, if $1 + \boldsymbol{v}^T \boldsymbol{M}^{-1} \boldsymbol{u} \neq 0$, then the following formula is valid:

$$(\boldsymbol{M} + \boldsymbol{u}\boldsymbol{v}^T)^{-1} = \boldsymbol{M}^{-1} - \frac{\boldsymbol{M}^{-1}\boldsymbol{u}\boldsymbol{v}^T\boldsymbol{M}^{-1}}{1 + \boldsymbol{v}^T\boldsymbol{M}^{-1}\boldsymbol{u}}. \quad \text{(Sherman-Morrison formula)}$$

Apply this formula to compute the inverses $\boldsymbol{B}_{k+1}$ of $\boldsymbol{H}_{k+1}$ for BFGS, DFP and symmetric-rank-one formulae.

14. Apply the quasi-Newton method with BFGS, DFP, and Symmetric-Rank-One updates for the strictly convex function $f(\boldsymbol{x}) = \alpha + \langle \boldsymbol{a}, \boldsymbol{x} \rangle + \frac{1}{2}\langle \boldsymbol{A}\boldsymbol{x}, \boldsymbol{x} \rangle$ with $\boldsymbol{A} \succ \boldsymbol{O}$.

5 Differentiable Convex Functions

5.1 Convex Functions

Definition 5.1 Let Q be a subset of $\mathbb{R}^n$. We denote by $\mathcal{F}^k(Q)$ the class of functions with the following properties:

- Any $f \in \mathcal{F}^k(Q)$ is k times continuously differentiable on Q;
- f is convex on Q, *i.e.*, given $\forall \boldsymbol{x}, \boldsymbol{y} \in Q$ and $\forall \alpha \in [0, 1]$,

$$f(\alpha \boldsymbol{x} + (1 - \alpha)\boldsymbol{y}) \leq \alpha f(\boldsymbol{x}) + (1 - \alpha) f(\boldsymbol{y}).$$

Theorem 5.2 $f \in \mathcal{F}(\mathbb{R}^n)$ if and only if its epigraph $E := \{(\boldsymbol{x}, y) \in \mathbb{R}^{n+1} \mid f(\boldsymbol{x}) \leq y\}$ is a convex.

Proof:

$\boxed{\Rightarrow}$ Let $(\boldsymbol{x}_1, y_1), (\boldsymbol{x}_2, y_2) \in E$. Then for any $0 \leq \alpha \leq 1$, we have

$$f(\alpha \boldsymbol{x}_1 + (1 - \alpha)\boldsymbol{x}_2) \leq \alpha f(\boldsymbol{x}_1) + (1 - \alpha) f(\boldsymbol{x}_2) \leq \alpha y_1 + (1 - \alpha) y_2$$

and therefore $(\alpha \boldsymbol{x}_1 + (1 - \alpha)\boldsymbol{x}_2, \alpha y_1 + (1 - \alpha) y_2) \in E$.

$\boxed{\Leftarrow}$ Let $(\boldsymbol{x}_1, f(\boldsymbol{x}_1)), (\boldsymbol{x}_2, f(\boldsymbol{x}_2)) \in E$. By the convexity of E, for any $0 \leq \alpha \leq 1$,

$$f(\alpha \boldsymbol{x}_1 + (1 - \alpha)\boldsymbol{x}_2) \leq \alpha f(\boldsymbol{x}_1) + (1 - \alpha) f(\boldsymbol{x}_2)$$

and therefore, $f \in \mathcal{F}(\mathbb{R}^n)$. ∎

Theorem 5.3 If $f \in \mathcal{F}(\mathbb{R}^n)$, then its λ-*level set* $L_\lambda := \{\boldsymbol{x} \in \mathbb{R}^n \mid f(\boldsymbol{x}) \leq \lambda\}$ is convex for each $\lambda \in \mathbb{R}$. But the converse is not true.

Proof:

For any $\lambda \in \mathbb{R}$, let $\boldsymbol{x}, \boldsymbol{y} \in L_\lambda$. Then for $\forall \alpha \in (0, 1)$, since $f \in \mathcal{F}(\mathbb{R}^n)$, $f(\alpha \boldsymbol{x} + (1 - \alpha)\boldsymbol{y}) \leq \alpha f(\boldsymbol{x}) + (1 - \alpha) f(\boldsymbol{y}) \leq \alpha \lambda + (1 - \alpha)\lambda = \lambda$. Therefore, $\alpha \boldsymbol{x} + (1 - \alpha)\boldsymbol{y} \in L_\lambda$.

For the converse, $L_\lambda = \{x \in \mathbb{R} \mid f(x) = x^3 \leq \lambda\}$ is convex for all $\lambda \in \mathbb{R}$, but $f \notin \mathcal{F}(\mathbb{R})$. ∎

Example 5.4 The function $-\log x$ is convex on $(0, +\infty)$. Let $a, b \in (0, +\infty)$ and $0 \leq \theta \leq 1$. Then, from the definition of the convexity, we have

$$-\log(\theta a + (1-\theta)b) \leq -\theta \log a - (1-\theta)\log b.$$

If we take the exponential of both sides, we obtain

$$a^\theta b^{1-\theta} \leq \theta a + (1-\theta)b.$$

For $\theta = \frac{1}{2}$, we have the arithmetic-geometric mean inequality: $\sqrt{ab} \leq \dfrac{a+b}{2}$.

Let $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n \backslash \{\mathbf{0}\}$, $p > 1$, and q such that $\frac{1}{p} + \frac{1}{q} = 1$. Consider

$$a = \frac{|[\boldsymbol{x}]_i|^p}{\displaystyle\sum_{j=1}^n |[\boldsymbol{x}]_j|^p}, \ b = \frac{|[\boldsymbol{y}]_i|^q}{\displaystyle\sum_{j=1}^n |[\boldsymbol{y}]_j|^q}, \ \theta = \frac{1}{p}, \text{ and } (1-\theta) = \frac{1}{q}.$$

Then we have

$$\left(\frac{|[\boldsymbol{x}]_i|^p}{\displaystyle\sum_{j=1}^n |[\boldsymbol{x}]_j|^p}\right)^{\frac{1}{p}} \left(\frac{|[\boldsymbol{y}]_i|^q}{\displaystyle\sum_{j=1}^n |[\boldsymbol{y}]_j|^q}\right)^{\frac{1}{q}} \leq \frac{|[\boldsymbol{x}]_i|^p}{p\displaystyle\sum_{j=1}^n |[\boldsymbol{x}]_j|^p} + \frac{|[\boldsymbol{y}]_i|^q}{q\displaystyle\sum_{j=1}^n |[\boldsymbol{y}]_j|^q}.$$

and summing over i, we obtain the Hölder inequality:

$$|\langle \boldsymbol{x}, \boldsymbol{y} \rangle| \leq \|\boldsymbol{x}\|_p \|\boldsymbol{y}\|_q$$

where $\|\boldsymbol{x}\|_p := \left(\displaystyle\sum_{i=1}^n |[\boldsymbol{x}]_i|^p\right)^{\frac{1}{p}}$.

Theorem 5.5 (Jensen's inequality) A function $f : \mathbb{R}^n \to \mathbb{R}$ is convex if and only if for any positive integer m, the following condition is valid

$$\left.\begin{array}{l} \boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_m \in \mathbb{R}^n \\ \alpha_1, \alpha_2, \ldots, \alpha_m \geq 0 \\ \displaystyle\sum_{i=1}^m \alpha_i = 1 \end{array}\right\} \Rightarrow f\left(\sum_{i=1}^m \alpha_i \boldsymbol{x}_i\right) \leq \sum_{i=1}^m \alpha_i f(\boldsymbol{x}_i).$$

Proof:

Left for exercise. ∎

Theorem 5.6 Let $\{f_i\}_{i \in I}$ be a family of (finite or infinite) functions which are bounded from above and $f_i \in \mathcal{F}(\mathbb{R}^n)$. Then, $f(\boldsymbol{x}) := \sup_{i \in I} f_i(\boldsymbol{x})$ is convex on $\mathbb{R}^n$.

Proof:

For each $i \in I$, since $f_i \in \mathcal{F}(\mathbb{R}^n)$, its epigraph $E_i = \{(\boldsymbol{x}, y) \in \mathbb{R}^{n+1} \mid f_i(\boldsymbol{x}) \leq y\}$ is convex on $\mathbb{R}^{n+1}$ by Theorem 5.2. Also their intersection

$$\bigcap_{i \in I} E_i = \bigcap_{i \in I} \left\{(\boldsymbol{x}, y) \in \mathbb{R}^{n+1} \mid f_i(\boldsymbol{x}) \leq y\right\} = \left\{(\boldsymbol{x}, y) \in \mathbb{R}^{n+1} \;\middle|\; \sup_{i \in I} f_i(\boldsymbol{x}) \leq y\right\}$$

is convex by Exercise 2 of Section 1, which is exactly the epigraph of $f(\boldsymbol{x})$. ∎

5.2 Differentiable Convex Functions

Theorem 5.7 Let f be a continuously differentiable function. The following conditions are equivalent:

1. $f \in \mathcal{F}^1(\mathbb{R}^n)$.
2. $f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \langle \boldsymbol{\nabla f}(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle, \quad \forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$.
3. $\langle \boldsymbol{\nabla f}(\boldsymbol{x}) - \boldsymbol{\nabla f}(\boldsymbol{y}), \boldsymbol{x} - \boldsymbol{y} \rangle \geq 0, \ \forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$.

Proof:
Left for exercise. ∎

Theorem 5.8 If $f \in \mathcal{F}^1(\mathbb{R}^n)$ and $\boldsymbol{\nabla f}(\boldsymbol{x}^*) = 0$, then $\boldsymbol{x}^*$ is the *global minimum* of $f(\boldsymbol{x})$ on $\mathbb{R}^n$.

Proof:
Left for exercise. ∎

Lemma 5.9 If $f \in \mathcal{F}^1(\mathbb{R}^m)$, $\boldsymbol{b} \in \mathbb{R}^m$, and $\boldsymbol{A} : \mathbb{R}^n \to \mathbb{R}^m$, then

$$\phi(\boldsymbol{x}) = f(\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}) \in \mathcal{F}^1(\mathbb{R}^n).$$

Proof:
Left for exercise. ∎

Example 5.10 The following functions are differentiable and convex:

1. $f(x) = e^x$
2. $f(x) = |x|^p, \quad p > 1$
3. $f(x) = \frac{x^2}{1+|x|}$
4. $f(x) = |x| - \ln(1 + |x|)$
5. $f(\boldsymbol{x}) = \sum_{i=1}^m e^{\alpha_i + \langle \boldsymbol{a}_i, \boldsymbol{x} \rangle}$
6. $f(\boldsymbol{x}) = \sum_{i=1}^m |\langle \boldsymbol{a}_i, \boldsymbol{x} \rangle - b_i|^p, \quad p > 1$

Theorem 5.11 Let f be a twice continuously differentiable function. Then $f \in \mathcal{F}^2(\mathbb{R}^n)$ if and only if

$$\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}) \succeq \boldsymbol{O}, \quad \forall \boldsymbol{x} \in \mathbb{R}^n.$$

Proof:
Let $f \in \mathcal{F}^2(\mathbb{R}^n)$, and denote $\boldsymbol{x}_\tau = \boldsymbol{x} + \tau \boldsymbol{s}$, $\tau > 0$. Then, from the previous result

$$\begin{aligned} 0 &\leq \frac{1}{\tau^2} \langle \boldsymbol{\nabla f}(\boldsymbol{x}_\tau) - \boldsymbol{\nabla f}(\boldsymbol{x}), \boldsymbol{x}_\tau - \boldsymbol{x} \rangle = \frac{1}{\tau} \langle \boldsymbol{\nabla f}(\boldsymbol{x}_\tau) - \boldsymbol{\nabla f}(\boldsymbol{x}), \boldsymbol{s} \rangle \\ &= \frac{1}{\tau} \int_0^\tau \langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x} + \lambda \boldsymbol{s}) \boldsymbol{s}, \boldsymbol{s} \rangle d\lambda \\ &= \frac{F(\tau) - F(0)}{\tau} \end{aligned}$$

where $F(\tau) = \int_0^\tau \langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x} + \lambda \boldsymbol{s}) \boldsymbol{s}, \boldsymbol{s} \rangle d\lambda$. Therefore, tending τ to 0, we get $0 \leq F'(0) = \langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}) \boldsymbol{s}, \boldsymbol{s} \rangle$, and we have the result.

Conversely, $\forall \boldsymbol{x} \in \mathbb{R}^n$,

$$\begin{aligned} f(\boldsymbol{y}) &= f(\boldsymbol{x}) + \langle \boldsymbol{\nabla f}(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle + \int_0^1 \int_0^\tau \langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x} + \lambda(\boldsymbol{y} - \boldsymbol{x}))(\boldsymbol{y} - \boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle d\lambda d\tau \\ &\geq f(\boldsymbol{x}) + \langle \boldsymbol{\nabla f}(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle. \end{aligned}$$

∎