We know from Corollary 4.5 that the number of iterations of the uniform grid method is at least $(\lfloor L/(2\varepsilon)\rfloor+2)^n$. Theorem 4.6 showed that any method which uses only function evaluations requires at least $(\lfloor L/(2\varepsilon)\rfloor)^n$ calls to have a better performance than ε. If for instance we take $\varepsilon = \mathcal{O}(L/n)$, these two bounds coincide up to a constant factor. In this sense, the uniform grid method is an *optimal method for the class of problems* $\mathcal{P}$.

4.3 Optimality Conditions for Smooth Optimization Problems

Let $f : \mathbb{R}^n \to \mathbb{R}$ be a differentiable function on $\mathbb{R}^n$, $\bar{\boldsymbol{x}} \in \mathbb{R}^n$, and $\boldsymbol{s}$ be a direction in $\mathbb{R}^n$ such that $\|\boldsymbol{s}\|_2 = 1$. Consider the local decrease (or increase) of $f(\boldsymbol{x})$ along $\boldsymbol{s}$:

$$f'(\boldsymbol{x};\boldsymbol{s}) = \lim_{\alpha\to 0} \frac{1}{\alpha}\left[f(\bar{\boldsymbol{x}}+\alpha\boldsymbol{s}) - f(\bar{\boldsymbol{x}})\right].$$

Since $f(\bar{\boldsymbol{x}}+\alpha\boldsymbol{s}) - f(\bar{\boldsymbol{x}}) = \alpha\langle\boldsymbol{\nabla} f(\bar{\boldsymbol{x}}),\boldsymbol{s}\rangle + o(\|\alpha\boldsymbol{s}\|_2)$, we have $f'(\bar{\boldsymbol{x}};\boldsymbol{s}) = \langle\boldsymbol{\nabla} f(\bar{\boldsymbol{x}}),\boldsymbol{s}\rangle$.
Using the Cauchy-Schwarz inequality $-\|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2 \le \langle\boldsymbol{x},\boldsymbol{y}\rangle \le \|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2$,

$$f'(\bar{\boldsymbol{x}};\boldsymbol{s}) = \langle\boldsymbol{\nabla} f(\bar{\boldsymbol{x}}),\boldsymbol{s}\rangle \ge -\|\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})\|_2.$$

Choosing in particular the direction $\bar{\boldsymbol{s}} = -\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})/\|\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})\|_2$,

$$f'(\bar{x};\boldsymbol{s}) = -\left\langle \boldsymbol{\nabla} f(\bar{\boldsymbol{x}}), \frac{\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})}{\|\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})\|_2}\right\rangle = -\|\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})\|_2.$$

Thus, the direction $-\boldsymbol{\nabla} f(\bar{\boldsymbol{x}})$ is the direction of the *fastest local decrease of* $f(\boldsymbol{x})$ *at point* $\bar{\boldsymbol{x}}$.

Theorem 4.8 (First-order necessary optimality condition) Let $\boldsymbol{x}^*$ be a local minimum of the differentiable function $f(\boldsymbol{x})$. Then

$$\boldsymbol{\nabla} f(\boldsymbol{x}^*) = \boldsymbol{0}.$$

Proof:
Let $\boldsymbol{x}^*$ be the local minimum of $f(\boldsymbol{x})$. Then, there is $r > 0$ such that for all $\boldsymbol{y}$ with $\|\boldsymbol{y}-\boldsymbol{x}^*\|_2 \le r$, $f(\boldsymbol{y}) \ge f(\boldsymbol{x}^*)$.
Since f is differentiable on $\mathbb{R}^n$,

$$f(\boldsymbol{y}) = f(\boldsymbol{x}^*) + \langle\boldsymbol{\nabla} f(\boldsymbol{x}^*),\boldsymbol{y}-\boldsymbol{x}^*\rangle + o(\|\boldsymbol{y}-\boldsymbol{x}^*\|_2) \ge f(\boldsymbol{x}^*).$$

Dividing by $\|\boldsymbol{y}-\boldsymbol{x}^*\|_2$, and taking the limit $\boldsymbol{y} \to \boldsymbol{x}^*$,

$$\langle\boldsymbol{\nabla} f(\boldsymbol{x}^*),\boldsymbol{s}\rangle \ge 0, \quad \forall \boldsymbol{s}\in\mathbb{R}^n, \quad \|\boldsymbol{s}\|_2 = 1.$$

Consider the opposite direction $-\boldsymbol{s}$, and then we conclude that

$$\langle\boldsymbol{\nabla} f(\boldsymbol{x}^*),\boldsymbol{s}\rangle = 0, \quad \forall \boldsymbol{s}\in\mathbb{R}^n, \quad \|\boldsymbol{s}\|_2 = 1.$$

Choosing $\boldsymbol{s} = \boldsymbol{e}_i \quad (i = 1,2,\ldots,n)$, we conclude that $\boldsymbol{\nabla} f(\boldsymbol{x}^*) = 0$. ∎

Remark 4.9 For the first-order sufficient optimality condition, we need convexity for the function $f(\boldsymbol{x})$.

Corollary 4.10 Let $\boldsymbol{x}^*$ be a local minimum of a differentiable function $f(\boldsymbol{x})$ subject to linear equality constraints

$$\boldsymbol{x} \in \mathcal{L} := \{\boldsymbol{x}\in\mathbb{R}^n \mid \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}\} \neq \emptyset,$$

where $\boldsymbol{A}\in\mathbb{R}^{m\times n}$, $\boldsymbol{b}\in\mathbb{R}^m$, $m < n$.
Then, there exists a vector of multipliers $\boldsymbol{\lambda}^*\in\mathbb{R}^m$ such that

$$\boldsymbol{\nabla} f(\boldsymbol{x}^*) = \boldsymbol{A}^T\boldsymbol{\lambda}^*.$$

Proof:
Consider the vectors $\boldsymbol{u}_i$ $(i = 1,2,\ldots,k)$ with $k \ge n-m$ which form an orthonormal basis of the null space of $\boldsymbol{A}$. Then, $\boldsymbol{x}\in\mathcal{L}$ can be represented as

$$\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{t}) := \boldsymbol{x}^* + \sum_{i=1}^{k} t_i \boldsymbol{u}_i, \quad \boldsymbol{t} \in \mathbb{R}^k.$$

Moreover, the point $\boldsymbol{t} = \boldsymbol{0}$ is the local minimal solution of the function $\phi(\boldsymbol{t}) = f(\boldsymbol{x}(\boldsymbol{t}))$.

From Theorem 4.8, $\phi'(\boldsymbol{0}) = \boldsymbol{0}$. That is,

$$\frac{d\phi}{dt_i}(\boldsymbol{0}) = \langle \boldsymbol{\nabla} \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{u}_i \rangle = 0, \quad i = 1, 2, \ldots, k.$$

Now there is $\boldsymbol{t}^* \in \mathbb{R}^k$ and $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that

$$\boldsymbol{\nabla} \boldsymbol{f}(\boldsymbol{x}^*) = \sum_{i=1}^{k} t_i^* \boldsymbol{u}_i + \boldsymbol{A}^T \boldsymbol{\lambda}^*.$$

For each $i = 1, 2, \ldots, k$,

$$\langle \boldsymbol{\nabla} \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{u}_i \rangle = t_i^* = 0.$$

Therefore, we have the result. ∎

The following type of result is called *theorems of the alternative*, and are closed related to duality theory in optimization.

Corollary 4.11 Given $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $\boldsymbol{b} \in \mathbb{R}^m$, $\boldsymbol{c} \in \mathbb{R}^n$, $\eta \in \mathbb{R}$, either

$$\begin{cases} \langle \boldsymbol{c}, \boldsymbol{x} \rangle < \eta \\ \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \end{cases} \text{ has a solution } \boldsymbol{x} \in \mathbb{R}^n, \tag{3}$$

or

$$\left(\begin{array}{c} \begin{cases} \langle \boldsymbol{b}, \boldsymbol{\lambda} \rangle > 0 \\ \boldsymbol{A}^T \boldsymbol{\lambda} = \boldsymbol{0} \end{cases} \\ \text{or} \\ \begin{cases} \langle \boldsymbol{b}, \boldsymbol{\lambda} \rangle \geq \eta \\ \boldsymbol{A}^T \boldsymbol{\lambda} = \boldsymbol{c} \end{cases} \end{array} \right) \text{ has a solution } \boldsymbol{\lambda} \in \mathbb{R}^m, \tag{4}$$

but never both

Proof:

Let us first show that if $\exists \boldsymbol{x} \in \mathbb{R}^n$ satisfying (3), $\nexists \boldsymbol{\lambda} \in \mathbb{R}^m$ satisfying (4). Let us assume by contradiction that $\exists \boldsymbol{\lambda}$. Then $\langle \boldsymbol{\lambda}, \boldsymbol{A}\boldsymbol{x} \rangle = \langle \boldsymbol{\lambda}, \boldsymbol{b} \rangle$ and in the homogeneous case it gives $0 = \langle \boldsymbol{\lambda}, \boldsymbol{b} \rangle > 0$ and in the non-homogeneous case it gives $\eta > \langle \boldsymbol{c}, \boldsymbol{x} \rangle = \langle \boldsymbol{\lambda}, \boldsymbol{b} \rangle \geq \eta$. Both of cases are impossible.

Now, let us assume that $\nexists \boldsymbol{x} \in \mathbb{R}^n$ satisfying (3). If additionally $\nexists \boldsymbol{x} \in \mathbb{R}^n$ such that $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, it means that the columns of the matrix $\boldsymbol{A}$ do not spam the vector $\boldsymbol{b}$. Therefore, there is $\boldsymbol{0} \neq \boldsymbol{\lambda} \in \mathbb{R}^m$ which is orthogonal to all of these columns and $\langle \boldsymbol{b}, \boldsymbol{\lambda} \rangle \neq 0$. Selecting the correct sign, we constructed a $\boldsymbol{\lambda}$ which satisfies the homogeneous system of (4). Now, if for all $\boldsymbol{x}$ such that $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ we have $\langle \boldsymbol{c}, \boldsymbol{x} \rangle \geq \eta$, it means that the minimization of the function $f(\boldsymbol{x}) = \langle \boldsymbol{c}, \boldsymbol{x} \rangle$ subject to $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ has an optimal solution $\boldsymbol{x}^*$ with $f(\boldsymbol{x}^*) \geq \eta$ (since $\exists \boldsymbol{x} \in \mathbb{R}^n$ such that $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, we can always assume that $m \leq n$ eliminating redundant linear constraints from the system. If $n = m$ and $\boldsymbol{A}$ is nonsingular, take $\boldsymbol{\lambda} = \boldsymbol{A}^{-T}\boldsymbol{c}$. Otherwise, we can eliminate again redundant linear constraint to have $n > m$). From Corollary 4.10, $\exists \boldsymbol{\lambda} \in \mathbb{R}^m$ such that $\boldsymbol{A}^T \boldsymbol{\lambda} = \boldsymbol{c}$, and $\langle \boldsymbol{b}, \boldsymbol{\lambda} \rangle = \langle \boldsymbol{x}^*, \boldsymbol{A}^T \boldsymbol{\lambda} \rangle = \langle \boldsymbol{x}^*, \boldsymbol{c} \rangle \geq \eta$. ∎

If $f(\boldsymbol{x})$ is twice differentiable at $\bar{\boldsymbol{x}} \in \mathbb{R}^n$, then for $\boldsymbol{y} \in \mathbb{R}^n$, we have

$$\boldsymbol{\nabla} \boldsymbol{f}(\boldsymbol{y}) = \boldsymbol{\nabla} \boldsymbol{f}(\bar{\boldsymbol{x}}) + \boldsymbol{\nabla}^2 \boldsymbol{f}(\bar{\boldsymbol{x}})(\boldsymbol{y} - \bar{\boldsymbol{x}}) + \boldsymbol{o}(\|\boldsymbol{y} - \bar{\boldsymbol{x}}\|_2),$$

where $\boldsymbol{o}(r)$ is such that $\lim_{r \to 0} \|\boldsymbol{o}(r)\|_2 / r = 0$ and $\boldsymbol{o}(0) = 0$.

Theorem 4.12 (Second-order necessary optimality condition) Let $\boldsymbol{x}^*$ be a local minimum of a twice continuously differentiable function $f(\boldsymbol{x})$. Then

$$\boldsymbol{\nabla f}(\boldsymbol{x}^*) = 0, \qquad \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*) \succeq \boldsymbol{O}.$$

Proof:

Since $\boldsymbol{x}^*$ is a local minimum of $f(\boldsymbol{x})$, $\exists r > 0$ such that for all $\boldsymbol{y} \in \mathbb{R}^n$ which satisfy $\|\boldsymbol{y} - \boldsymbol{x}^*\|_2 \leq r$, $f(\boldsymbol{y}) \geq f(\boldsymbol{x}^*)$.

From Theorem 4.8, $\boldsymbol{\nabla f}(\boldsymbol{x}^*) = 0$. Then

$$f(\boldsymbol{y}) = f(\boldsymbol{x}^*) + \frac{1}{2}\langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*)(\boldsymbol{y} - \boldsymbol{x}^*), \boldsymbol{y} - \boldsymbol{x}^* \rangle + o(\|\boldsymbol{y} - \boldsymbol{x}^*\|_2^2) \geq f(\boldsymbol{x}^*).$$

And $\langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*)\boldsymbol{s}, \boldsymbol{s} \rangle \geq 0, \ \forall \boldsymbol{s} \in \mathbb{R}^n$ with $\|\boldsymbol{s}\|_2 = 1$. ∎

Theorem 4.13 (Second-order sufficient optimality condition) Let the function $f(\boldsymbol{x})$ be twice continuously differentiable on $\mathbb{R}^n$, and let $\boldsymbol{x}^*$ satisfy the following conditions:

$$\boldsymbol{\nabla f}(\boldsymbol{x}^*) = 0, \quad \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*) \succ \boldsymbol{O}.$$

Then, $\boldsymbol{x}^*$ is a strict local minimum of $f(\boldsymbol{x})$.

Proof:

In a small neighborhood of $\boldsymbol{x}^*$, function $f(\boldsymbol{x}^*)$ can be represented as:

$$f(\boldsymbol{y}) = f(\boldsymbol{x}^*) + \frac{1}{2}\langle \boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*)(\boldsymbol{y} - \boldsymbol{x}^*), \boldsymbol{y} - \boldsymbol{x}^* \rangle + o(\|\boldsymbol{y} - \boldsymbol{x}^*\|_2^2).$$

Since $o(r)/r \to 0$, there is a $\bar{r} > 0$ such that for all $r \in [0, \bar{r}]$,

$$|o(r)| \leq \frac{r}{4}\lambda_1(\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*)),$$

where $\lambda_1(\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*))$ is the smallest eigenvalue of the symmetric matrix $\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*)$ which is positive. Then

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}^*) + \frac{1}{2}\lambda_1(\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*))\|\boldsymbol{y} - \boldsymbol{x}^*\|_2^2 + o(\|\boldsymbol{y} - \boldsymbol{x}^*\|_2^2).$$

W.L.O.G, considering that $\bar{r} < 1$, $|o(r^2)| \leq r^2\lambda_1(\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*))/4$ for $r \in [0, \bar{r}]$, finally we arrived at

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}^*) + \frac{1}{4}\lambda_1(\boldsymbol{\nabla}^2 \boldsymbol{f}(\boldsymbol{x}^*))\|\boldsymbol{y} - \boldsymbol{x}^*\|_2^2 > f(\boldsymbol{x}^*).$$

∎

4.4 Algorithms for Minimizing Smooth Functions

4.4.1 Steepest Descent Method

Consider $f : \mathbb{R}^n \to \mathbb{R}$ a differentiable function on its domain.

Steepest Descent Method	
Choose:	$\boldsymbol{x}_0 \in \mathbb{R}^n$
Iterate:	$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k - h_k \boldsymbol{\nabla f}(\boldsymbol{x}_k), \ \ k = 0, 1, \ldots$

We consider four strategies for the step-size h_k: