CHAPTER 6

BASICS OF SET-CONSTRAINED AND UNCONSTRAINED OPTIMIZATION

6.1 Introduction

In this chapter we consider the optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega. \end{array}$$

The function $f : \mathbb{R}^n \to \mathbb{R}$ that we wish to minimize is a real-valued function called the *objective function* or *cost function.* The vector $\boldsymbol{x}$ is an n-vector of independent variables: $\boldsymbol{x} = [x_1, x_2, \ldots, x_n]^\top \in \mathbb{R}^n$. The variables $x_1, \ldots, x_n$ are often referred to as *decision variables.* The set Ω is a subset of $\mathbb{R}^n$ called the *constraint set* or *feasible set.*

The optimization problem above can be viewed as a decision problem that involves finding the "best" vector $\boldsymbol{x}$ of the decision variables over all possible vectors in Ω. By the "best" vector we mean the one that results in the smallest value of the objective function. This vector is called the *minimizer* of f over Ω. It is possible that there may be many minimizers. In this case, finding any of the minimizers will suffice.

There are also optimization problems that require maximization of the objective function, in which case we seek *maximizers.* Minimizers and maximizers are also called *extremizers.* Maximization problems, however, can be represented equivalently in the minimization form above because maximizing f is equivalent to minimizing $-f$. Therefore, we can confine our attention to minimization problems without loss of generality.

The problem above is a general form of a *constrained optimization problem*, because the decision variables are constrained to be in the constraint set Ω. If $\Omega = \mathbb{R}^n$, then we refer to the problem as an *unconstrained optimization problem.* In this chapter we discuss basic properties of the general optimization problem above, which includes the unconstrained case. In the remaining chapters of this part, we deal with iterative algorithms for solving unconstrained optimization problems.

The constraint "$\boldsymbol{x} \in \Omega$" is called a *set constraint.* Often, the constraint set Ω takes the form $\Omega = \{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0},\ \boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}\}$, where $\boldsymbol{h}$ and $\boldsymbol{g}$ are given functions. We refer to such constraints as *functional constraints.* The remainder of this chapter deals with general set constraints, including the special case where $\Omega = \mathbb{R}^n$. The case where $\Omega = \mathbb{R}^n$ is called the *unconstrained* case. In Parts III and IV we consider constrained optimization problems with functional constraints.

In considering the general optimization problem above, we distinguish between two kinds of minimizers, as specified by the following definitions.

Definition 6.1 Suppose that $f : \mathbb{R}^n \to \mathbb{R}$ is a real-valued function defined on some set $\Omega \subset \mathbb{R}^n$. A point $\boldsymbol{x}^* \in \Omega$ is a *local minimizer* of f over Ω if there exists $\varepsilon > 0$ such that $f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*)$ for all $\boldsymbol{x} \in \Omega \setminus \{\boldsymbol{x}^*\}$ and $\|\boldsymbol{x} - \boldsymbol{x}^*\| < \varepsilon$. A point $\boldsymbol{x}^* \in \Omega$ is a *global minimizer* of f over Ω if $f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*)$ for all $\boldsymbol{x} \in \Omega \setminus \{\boldsymbol{x}^*\}$. ∎

If in the definitions above we replace "$\geq$" with "$>$," then we have a *strict local minimizer* and a *strict global minimizer*, respectively. In Figure 6.1, we illustrate the definitions for $n = 1$.

If $\boldsymbol{x}^*$ is a global minimizer of f over Ω, we write $f(\boldsymbol{x}^*) = \min_{\boldsymbol{x}\in\Omega} f(\boldsymbol{x})$ and $\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}\in\Omega} f(\boldsymbol{x})$. If the minimization is unconstrained, we simply write $\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} f(\boldsymbol{x})$ or $\boldsymbol{x}^* = \arg\min f(\boldsymbol{x})$. In other words, given a real-valued function f, the notation $\arg\min f(\boldsymbol{x})$ denotes the *argument* that minimizes the function f (a point in the domain of f), assuming that such a point is unique (if there is more than one such point, we pick one arbitrarily). For example, if $f : \mathbb{R} \to \mathbb{R}$ is given by $f(x) = (x+1)^2 + 3$, then $\arg\min f(x) = -1$. If we write $\arg\min_{x\in\Omega}$, then we treat "$x \in \Omega$" to be a constraint for the minimization. For example, for the function f above, $\arg\min_{x\geq 0} f(x) = 0$.

Strictly speaking, an optimization problem is solved only when a global minimizer is found. However, global minimizers are, in general, difficult to find. Therefore, in practice, we often have to be satisfied with finding local minimizers.

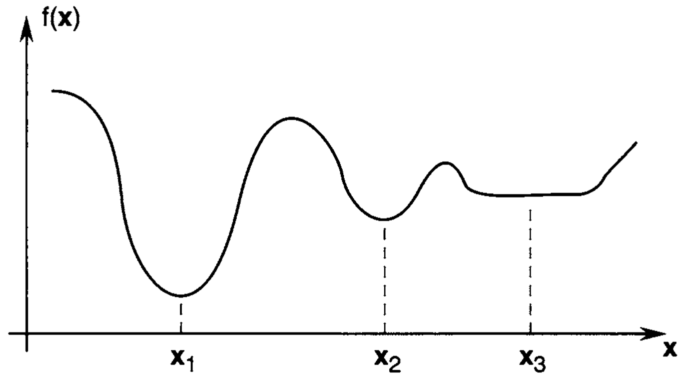

Figure 6.1 Examples of minimizers: $\boldsymbol{x}_1$: strict global minimizer; $\boldsymbol{x}_2$: strict local minimizer; $\boldsymbol{x}_3$: local (not strict) minimizer.

6.2 Conditions for Local Minimizers

In this section we derive conditions for a point $\boldsymbol{x}^*$ to be a local minimizer. We use derivatives of a function $f : \mathbb{R}^n \to \mathbb{R}$. Recall that the first-order derivative of f, denoted Df, is

$$Df \triangleq \left[\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \ldots, \frac{\partial f}{\partial x_n}\right].$$

Note that the gradient ∇f is just the transpose of Df; that is, $\nabla f = (Df)^\top$. The second derivative of $f : \mathbb{R}^n \to \mathbb{R}$ (also called the *Hessian* of f) is

$$\boldsymbol{F}(\boldsymbol{x}) \triangleq D^2 f(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2}(\boldsymbol{x}) & \cdots & \frac{\partial^2 f}{\partial x_n \partial x_1}(\boldsymbol{x}) \\ \vdots & & \vdots \\ \frac{\partial^2 f}{\partial x_1 \partial x_n}(\boldsymbol{x}) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(\boldsymbol{x}) \end{bmatrix}.$$

Example 6.1 Let $f(x_1, x_2) = 5x_1 + 8x_2 + x_1x_2 - x_1^2 - 2x_2^2$. Then,

$$Df(\boldsymbol{x}) = (\nabla f(\boldsymbol{x}))^\top = \left[\frac{\partial f}{\partial x_1}(\boldsymbol{x}), \frac{\partial f}{\partial x_2}(\boldsymbol{x})\right] = [5 + x_2 - 2x_1, 8 + x_1 - 4x_2]$$

and

$$\boldsymbol{F}(\boldsymbol{x}) = D^2 f(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2}(\boldsymbol{x}) & \frac{\partial^2 f}{\partial x_2 \partial x_1}(\boldsymbol{x}) \\ \frac{\partial^2 f}{\partial x_1 \partial x_2}(\boldsymbol{x}) & \frac{\partial^2 f}{\partial x_2^2}(\boldsymbol{x}) \end{bmatrix} = \begin{bmatrix} -2 & 1 \\ 1 & -4 \end{bmatrix}.$$

∎

Given an optimization problem with constraint set Ω, a minimizer may lie either in the interior or on the boundary of Ω. To study the case where it lies on the boundary, we need the notion of *feasible directions*.

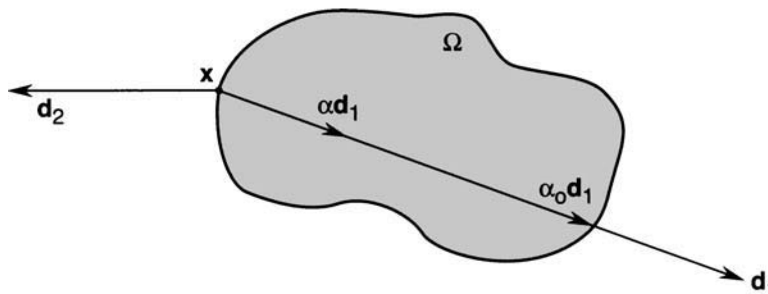

Figure 6.2 Two-dimensional illustration of feasible directions; $\boldsymbol{d}_1$ is a feasible direction, $\boldsymbol{d}_2$ is not a feasible direction.

Definition 6.2 A vector $\boldsymbol{d} \in \mathbb{R}^n$, $\boldsymbol{d} \neq \boldsymbol{0}$, is a *feasible direction* at $\boldsymbol{x} \in \Omega$ if there exists $\alpha_0 > 0$ such that $\boldsymbol{x} + \alpha \boldsymbol{d} \in \Omega$ for all $\alpha \in [0, \alpha_0]$. ∎

Figure 6.2 illustrates the notion of feasible directions.

Let $f : \mathbb{R}^n \to \mathbb{R}$ be a real-valued function and let $\boldsymbol{d}$ be a feasible direction at $\boldsymbol{x} \in \Omega$. The *directional derivative of f in the direction $\boldsymbol{d}$*, denoted $\partial f / \partial \boldsymbol{d}$, is the real-valued function defined by

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}) = \lim_{\alpha \to 0} \frac{f(\boldsymbol{x} + \alpha \boldsymbol{d}) - f(\boldsymbol{x})}{\alpha}.$$

If $\|\boldsymbol{d}\| = 1$, then $\partial f / \partial \boldsymbol{d}$ is the rate of increase of f at $\boldsymbol{x}$ in the direction $\boldsymbol{d}$. To compute the directional derivative above, suppose that $\boldsymbol{x}$ and $\boldsymbol{d}$ are given. Then, $f(\boldsymbol{x} + \alpha \boldsymbol{d})$ is a function of α, and

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}) = \left. \frac{d}{d\alpha} f(\boldsymbol{x} + \alpha \boldsymbol{d}) \right|_{\alpha=0}.$$

Applying the chain rule yields

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}) = \left. \frac{d}{d\alpha} f(\boldsymbol{x} + \alpha \boldsymbol{d}) \right|_{\alpha=0} = \nabla f(\boldsymbol{x})^\top \boldsymbol{d} = \langle \nabla f(\boldsymbol{x}), \boldsymbol{d} \rangle = \boldsymbol{d}^\top \nabla f(\boldsymbol{x}).$$

In summary, if $\boldsymbol{d}$ is a unit vector ($\|\boldsymbol{d}\| = 1$), then $\langle \nabla f(\boldsymbol{x}), \boldsymbol{d} \rangle$ is the rate of increase of f at the point $\boldsymbol{x}$ in the direction $\boldsymbol{d}$.

Example 6.2 Define $f : \mathbb{R}^3 \to \mathbb{R}$ by $f(\boldsymbol{x}) = x_1 x_2 x_3$, and let

$$\boldsymbol{d} = \left[\frac{1}{2}, \frac{1}{2}, \frac{1}{\sqrt{2}} \right]^\top.$$

The directional derivative of f in the direction $\boldsymbol{d}$ is

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}) = \nabla f(\boldsymbol{x})^\top \boldsymbol{d} = [x_2 x_3, x_1 x_3, x_1 x_2] \begin{bmatrix} 1/2 \\ 1/2 \\ 1/\sqrt{2} \end{bmatrix} = \frac{x_2 x_3 + x_1 x_3 + \sqrt{2} x_1 x_2}{2}.$$

Note that because $\|\boldsymbol{d}\| = 1$, the above is also the rate of increase of f at $\boldsymbol{x}$ in the direction $\boldsymbol{d}$. ∎

We are now ready to state and prove the following theorem.

Theorem 6.1 ***First-Order Necessary Condition (FONC).*** *Let Ω be a subset of $\mathbb{R}^n$ and $f \in \mathcal{C}^1$ a real-valued function on Ω. If $\boldsymbol{x}^*$ is a local minimizer of f over Ω, then for any feasible direction $\boldsymbol{d}$ at $\boldsymbol{x}^*$, we have*

$$\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq 0.$$

□

Proof. Define

$$\boldsymbol{x}(\alpha) = \boldsymbol{x}^* + \alpha \boldsymbol{d} \in \Omega.$$

Note that $\boldsymbol{x}(0) = \boldsymbol{x}^*$. Define the composite function

$$\phi(\alpha) = f(\boldsymbol{x}(\alpha)).$$

Then, by Taylor's theorem,

$$f(\boldsymbol{x}^* + \alpha \boldsymbol{d}) - f(\boldsymbol{x}^*) = \phi(\alpha) - \phi(0) = \phi'(0)\alpha + o(\alpha) = \alpha \boldsymbol{d}^\top \nabla f(\boldsymbol{x}(0)) + o(\alpha),$$

where $\alpha \geq 0$ [recall the definition of $o(\alpha)$ ("little-oh of α") in Part I]. Thus, if $\phi(\alpha) \geq \phi(0)$, that is, $f(\boldsymbol{x}^* + \alpha \boldsymbol{d}) \geq f(\boldsymbol{x}^*)$ for sufficiently small values of $\alpha > 0$ ($\boldsymbol{x}^*$ is a local minimizer), then we have to have $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq 0$ (see Exercise 5.8). ∎

Theorem 6.1 is illustrated in Figure 6.3.

An alternative way to express the FONC is

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}^*) \geq 0$$

for all feasible directions $\boldsymbol{d}$. In other words, if $\boldsymbol{x}^*$ is a local minimizer, then the rate of increase of f at $\boldsymbol{x}^*$ in any feasible direction $\boldsymbol{d}$ in Ω is nonnegative. Using directional derivatives, an alternative proof of Theorem 6.1 is as follows. Suppose that $\boldsymbol{x}^*$ is a local minimizer. Then, for any feasible direction $\boldsymbol{d}$, there exists $\bar{\alpha} > 0$ such that for all $\alpha \in (0, \bar{\alpha})$,

$$f(\boldsymbol{x}^*) \leq f(\boldsymbol{x}^* + \alpha \boldsymbol{d}).$$

Hence, for all $\alpha \in (0, \bar{\alpha})$, we have

$$\frac{f(\boldsymbol{x}^* + \alpha \boldsymbol{d}) - f(\boldsymbol{x}^*)}{\alpha} \geq 0.$$

Taking the limit as $\alpha \to 0$, we conclude that

$$\frac{\partial f}{\partial \boldsymbol{d}}(\boldsymbol{x}^*) \geq 0.$$

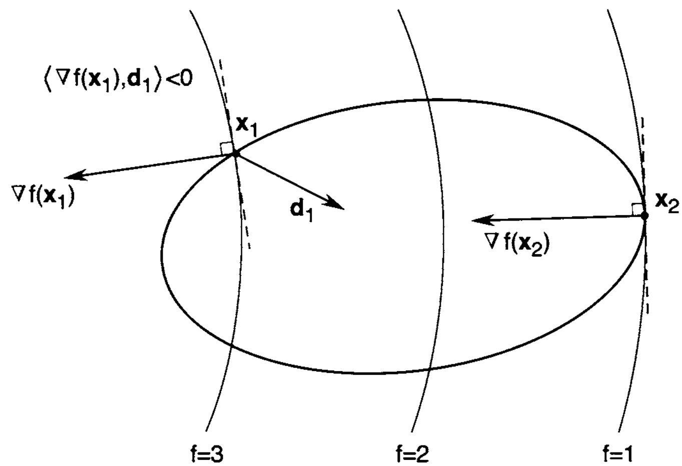

Figure 6.3 Illustration of the FONC for a constrained case; x_1 does not satisfy the FONC, whereas x_2 satisfies the FONC.

A special case of interest is when x^* is an interior point of Ω (see Section 4.4). In this case, any direction is feasible, and we have the following result.

Corollary 6.1 ***Interior Case.*** *Let Ω be a subset of $\mathbb{R}^n$ and $f \in \mathcal{C}^1$ a real-valued function on Ω. If x^* is a local minimizer of f over Ω and if x^* is an interior point of Ω, then*

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

□

Proof. Suppose that f has a local minimizer x^* that is an interior point of Ω. Because x^* is an interior point of Ω, the set of feasible directions at x^* is the whole of $\mathbb{R}^n$. Thus, for any $d \in \mathbb{R}^n$, $d^\top \nabla f(x^*) \geq 0$ and $-d^\top \nabla f(x^*) \geq 0$. Hence, $d^\top \nabla f(x^*) = 0$ for all $d \in \mathbb{R}^n$, which implies that $\nabla f(x^*) = \mathbf{0}$. ∎

Example 6.3 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_1^2 + 0.5x_2^2 + 3x_2 + 4.5 \\ \text{subject to} & x_1, x_2 \geq 0. \end{array}$$

a. Is the first-order necessary condition (FONC) for a local minimizer satisfied at $x = [1, 3]^\top$?

b. Is the FONC for a local minimizer satisfied at $x = [0, 3]^\top$?

c. Is the FONC for a local minimizer satisfied at $x = [1, 0]^\top$?

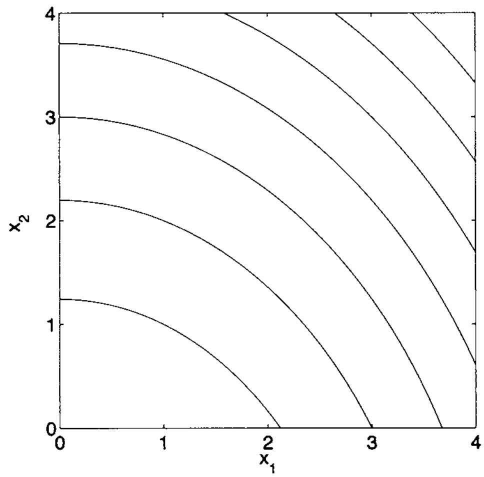

Figure 6.4 Level sets of the function in Example 6.3.

d. Is the FONC for a local minimizer satisfied at $\boldsymbol{x} = [0,0]^\top$?

Solution: First, let $f : \mathbb{R}^2 \to \mathbb{R}$ be defined by $f(\boldsymbol{x}) = x_1^2 + 0.5x_2^2 + 3x_2 + 4.5$, where $\boldsymbol{x} = [x_1, x_2]^\top$. A plot of the level sets of f is shown in Figure 6.4.

a. At $\boldsymbol{x} = [1,3]^\top$, we have $\nabla f(\boldsymbol{x}) = [2x_1, x_2 + 3]^\top = [2,6]^\top$. The point $\boldsymbol{x} = [1,3]^\top$ is an interior point of $\Omega = \{\boldsymbol{x} : x_1 \geq 0, x_2 \geq 0\}$. Hence, the FONC requires that $\nabla f(\boldsymbol{x}) = \boldsymbol{0}$. The point $\boldsymbol{x} = [1,3]^\top$ does not satisfy the FONC for a local minimizer.

b. At $\boldsymbol{x} = [0,3]^\top$, we have $\nabla f(\boldsymbol{x}) = [0,6]^\top$, and hence $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}) = 6d_2$, where $\boldsymbol{d} = [d_1, d_2]^\top$. For $\boldsymbol{d}$ to be feasible at $\boldsymbol{x}$, we need $d_1 \geq 0$, and d_2 can take an arbitrary value in $\mathbb{R}$. The point $\boldsymbol{x} = [0,3]^\top$ does not satisfy the FONC for a minimizer because d_2 is allowed to be less than zero. For example, $\boldsymbol{d} = [1,-1]^\top$ is a feasible direction, but $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}) = -6 < 0$.

c. At $\boldsymbol{x} = [1,0]^\top$, we have $\nabla f(\boldsymbol{x}) = [2,3]^\top$, and hence $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}) = 2d_1 + 3d_2$. For $\boldsymbol{d}$ to be feasible, we need $d_2 \geq 0$, and d_1 can take an arbitrary value in $\mathbb{R}$. For example, $\boldsymbol{d} = [-5,1]^\top$ is a feasible direction. But $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}) = -7 < 0$. Thus, $\boldsymbol{x} = [1,0]^\top$ does not satisfy the FONC for a local minimizer.

d. At $\boldsymbol{x} = [0,0]^\top$, we have $\nabla f(\boldsymbol{x}) = [0,3]^\top$, and hence $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}) = 3d_2$. For $\boldsymbol{d}$ to be feasible, we need $d_2 \geq 0$ and $d_1 \geq 0$. Hence, $\boldsymbol{x} = [0,0]^\top$ satisfies the FONC for a local minimizer. ∎

Example 6.4 Figure 6.5 shows a simplified model of a cellular wireless system (the distances shown have been scaled down to make the calculations

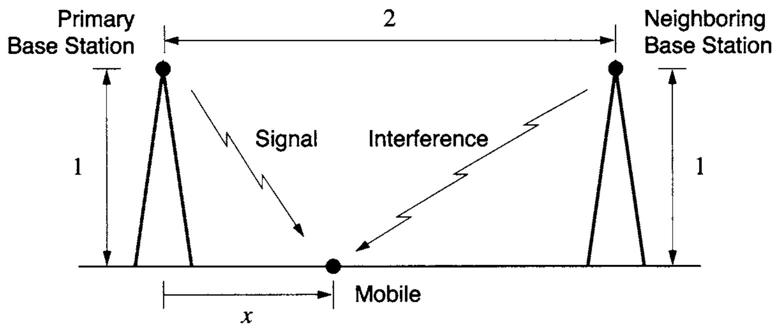

Figure 6.5 Simplified cellular wireless system in Example 6.4.

simpler). A mobile user (also called a *mobile*) is located at position x (see Figure 6.5).

There are two base station antennas, one for the primary base station and another for the neighboring base station. Both antennas are transmitting signals to the mobile user, at equal power. However, the power of the received signal as measured by the mobile is the reciprocal of the squared distance from the associated antenna (primary or neighboring base station). We are interested in finding the position of the mobile that maximizes the *signal-to-interference ratio*, which is the ratio of the signal power received from the primary base station to the signal power received from the neighboring base station.

We use the FONC to solve this problem. The squared distance from the mobile to the primary antenna is $1 + x^2$, while the squared distance from the mobile to the neighboring antenna is $1 + (2 - x)^2$. Therefore, the signal-to-interference ratio is

$$f(x) = \frac{1 + (2 - x)^2}{1 + x^2}.$$

We have

$$\begin{aligned} f'(x) &= \frac{-2(2 - x)(1 + x^2) - 2x(1 + (2 - x)^2)}{(1 + x^2)^2} \\ &= \frac{4(x^2 - 2x - 1)}{(1 + x^2)^2}. \end{aligned}$$

By the FONC, at the optimal position x^* we have $f'(x^*) = 0$. Hence, either $x^* = 1 - \sqrt{2}$ or $x^* = 1 + \sqrt{2}$. Evaluating the objective function at these two candidate points, it easy to see that $x^* = 1 - \sqrt{2}$ is the optimal position. ∎

The next example illustrates that in some problems the FONC is not helpful for eliminating candidate local minimizers. However, in such cases, there may be a recasting of the problem into an equivalent form that makes the FONC useful.

Example 6.5 Consider the set-constrained problem

$$\begin{aligned} &\text{minimize} && f(\boldsymbol{x}) \\ &\text{subject to} && \boldsymbol{x} \in \Omega, \end{aligned}$$

where $\Omega = \{[x_1, x_2]^\top : x_1^2 + x_2^2 = 1\}$.

a. Consider a point $\boldsymbol{x}^* \in \Omega$. Specify all feasible directions at $\boldsymbol{x}^*$.

b. Which points in Ω satisfy the FONC for this set-constrained problem?

c. Based on part b, is the FONC for this set-constrained problem useful for eliminating local-minimizer candidates?

d. Suppose that we use polar coordinates to parameterize points $\boldsymbol{x} \in \Omega$ in terms of a single parameter θ:

$$x_1 = \cos\theta \qquad x_2 = \sin\theta.$$

Now use the FONC for unconstrained problems (with respect to θ) to derive a necessary condition of this sort: If $\boldsymbol{x}^* \in \Omega$ is a local minimizer, then $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) = 0$ for all $\boldsymbol{d}$ satisfying a "certain condition." Specify what this certain condition is.

Solution:

a. There are no feasible directions at any $\boldsymbol{x}^*$.

b. Because of part a, *all* points in Ω satisfy the FONC for this set-constrained problem.

c. No, the FONC for this set-constrained problem is not useful for eliminating local-minimizer candidates.

d. Write $h(\theta) = f(g(\theta))$, where $g : \mathbb{R} \to \mathbb{R}^2$ is given by the equations relating θ to $\boldsymbol{x} = [x_1, x_2]^\top$. Note that $Dg(\theta) = [-\sin\theta, \cos\theta]^\top$. Hence, by the chain rule,

$$h'(\theta) = Df(g(\theta))Dg(\theta) = Dg(\theta)^\top \nabla f(g(\theta)).$$

Notice that $Dg(\theta)$ is tangent to Ω at $\boldsymbol{x} = g(\theta)$. Alternatively, we could say that $Dg(\theta)$ is orthogonal to $\boldsymbol{x} = g(\theta)$.

Suppose that $\boldsymbol{x}^* \in \Omega$ is a local minimizer. Write $\boldsymbol{x}^* = g(\theta^*)$. Then θ^* is an unconstrained minimizer of h. By the FONC for unconstrained problems, $h'(\theta^*) = 0$, which implies that $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) = 0$ for all $\boldsymbol{d}$ tangent to Ω at $\boldsymbol{x}^*$ (or, alternatively, for all $\boldsymbol{d}$ orthogonal to $\boldsymbol{x}^*$). ∎

We now derive a second-order necessary condition that is satisfied by a local minimizer.

Theorem 6.2 ***Second-Order Necessary Condition (SONC).*** *Let $\Omega \subset \mathbb{R}^n$, $f \in \mathcal{C}^2$ a function on Ω, $\boldsymbol{x}^*$ a local minimizer of f over Ω, and $\boldsymbol{d}$ a feasible direction at $\boldsymbol{x}^*$. If $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) = 0$, then*

$$\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} \geq 0,$$

where $\boldsymbol{F}$ is the Hessian of f. □

Proof. We prove the result by contradiction. Suppose that there is a feasible direction $\boldsymbol{d}$ at $\boldsymbol{x}^*$ such that $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) = 0$ and $\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} < 0$. Let $\boldsymbol{x}(\alpha) = \boldsymbol{x}^* + \alpha \boldsymbol{d}$ and define the composite function $\phi(\alpha) = f(\boldsymbol{x}^* + \alpha\boldsymbol{d}) = f(\boldsymbol{x}(\alpha))$. Then, by Taylor's theorem,

$$\phi(\alpha) = \phi(0) + \phi''(0)\frac{\alpha^2}{2} + o(\alpha^2),$$

where by assumption, $\phi'(0) = \boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) = 0$ and $\phi''(0) = \boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} < 0$. For sufficiently small α,

$$\phi(\alpha) - \phi(0) = \phi''(0)\frac{\alpha^2}{2} + o(\alpha^2) < 0,$$

that is,

$$f(\boldsymbol{x}^* + \alpha\boldsymbol{d}) < f(\boldsymbol{x}^*),$$

which contradicts the assumption that $\boldsymbol{x}^*$ is a local minimizer. Thus,

$$\phi''(0) = \boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} \geq 0.$$

■

Corollary 6.2 ***Interior Case.*** *Let $\boldsymbol{x}^*$ be an interior point of $\Omega \subset \mathbb{R}^n$. If $\boldsymbol{x}^*$ is a local minimizer of $f : \Omega \to \mathbb{R}$, $f \in \mathcal{C}^2$, then*

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

and $\boldsymbol{F}(\boldsymbol{x}^*)$ is positive semidefinite ($\boldsymbol{F}(\boldsymbol{x}^*) \geq 0$); that is, for all $\boldsymbol{d} \in \mathbb{R}^n$,

$$\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} \geq 0.$$

□

Proof. If $\boldsymbol{x}^*$ is an interior point, then all directions are feasible. The result then follows from Corollary 6.1 and Theorem 6.2. ■

In the examples below, we show that the necessary conditions are *not* sufficient.

Example 6.6 Consider a function of one variable $f(x) = x^3$, $f : \mathbb{R} \to \mathbb{R}$. Because $f'(0) = 0$, and $f''(0) = 0$, the point $x = 0$ satisfies both the FONC and SONC. However, $x = 0$ is not a minimizer (see Figure 6.6). ■

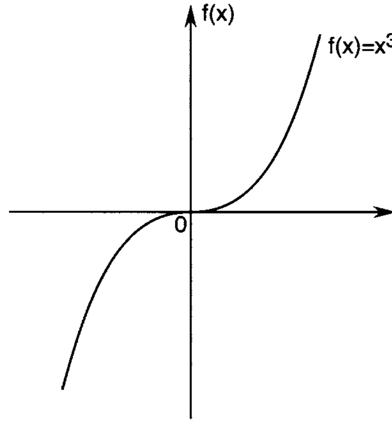

Figure 6.6 The point 0 satisfies the FONC and SONC but is not a minimizer.

Example 6.7 Consider a function $f : \mathbb{R}^2 \to \mathbb{R}$, where $f(\boldsymbol{x}) = x_1^2 - x_2^2$. The FONC requires that $\nabla f(\boldsymbol{x}) = [2x_1, -2x_2]^\top = \mathbf{0}$. Thus, $\boldsymbol{x} = [0, 0]^\top$ satisfies the FONC. The Hessian matrix of f is

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix}.$$

The Hessian matrix is indefinite; that is, for some $\boldsymbol{d}_1 \in \mathbb{R}^2$ we have $\boldsymbol{d}_1^\top \boldsymbol{F} \boldsymbol{d}_1 > 0$ (e.g., $\boldsymbol{d}_1 = [1, 0]^\top$) and for some $\boldsymbol{d}_2$ we have $\boldsymbol{d}_2^\top \boldsymbol{F} \boldsymbol{d}_2 < 0$ (e.g., $\boldsymbol{d}_2 = [0, 1]^\top$). Thus, $\boldsymbol{x} = [0, 0]^\top$ does not satisfy the SONC, and hence it is not a minimizer. The graph of $f(\boldsymbol{x}) = x_1^2 - x_2^2$ is shown in Figure 6.7. ∎

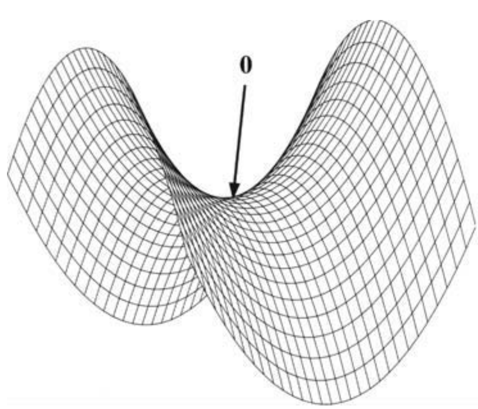

Figure 6.7 Graph of $f(\boldsymbol{x}) = x_1^2 - x_2^2$. The point $\mathbf{0}$ satisfies the FONC but not SONC; this point is not a minimizer.

We now derive sufficient conditions that imply that $\boldsymbol{x}^*$ is a local minimizer.

Theorem 6.3 ***Second-Order Sufficient Condition (SOSC), Interior Case.*** *Let $f \in \mathcal{C}^2$ be defined on a region in which $\boldsymbol{x}^*$ is an interior point. Suppose that*

1. $\nabla f(\boldsymbol{x}^*) = \boldsymbol{0}$.
2. $\boldsymbol{F}(\boldsymbol{x}^*) > 0$.

Then, $\boldsymbol{x}^*$ is a strict local minimizer of f. □

Proof. Because $f \in \mathcal{C}^2$, we have $\boldsymbol{F}(\boldsymbol{x}^*) = \boldsymbol{F}^\top(\boldsymbol{x}^*)$. Using assumption 2 and Rayleigh's inequality it follows that if $\boldsymbol{d} \neq \boldsymbol{0}$, then $0 < \lambda_{\min}(\boldsymbol{F}(\boldsymbol{x}^*))\|\boldsymbol{d}\|^2 \leq \boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d}$. By Taylor's theorem and assumption 1,

$$f(\boldsymbol{x}^* + \boldsymbol{d}) - f(\boldsymbol{x}^*) = \frac{1}{2}\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} + o(\|\boldsymbol{d}\|^2) \geq \frac{\lambda_{\min}(\boldsymbol{F}(\boldsymbol{x}^*))}{2}\|\boldsymbol{d}\|^2 + o(\|\boldsymbol{d}\|^2).$$

Hence, for all $\boldsymbol{d}$ such that $\|\boldsymbol{d}\|$ is sufficiently small,

$$f(\boldsymbol{x}^* + \boldsymbol{d}) > f(\boldsymbol{x}^*),$$

which completes the proof. ∎

Example 6.8 Let $f(\boldsymbol{x}) = x_1^2 + x_2^2$. We have $\nabla f(\boldsymbol{x}) = [2x_1, 2x_2]^\top = \boldsymbol{0}$ if and only if $\boldsymbol{x} = [0,0]^\top$. For all $\boldsymbol{x} \in \mathbb{R}^2$, we have

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} > 0.$$

The point $\boldsymbol{x} = [0,0]^\top$ satisfies the FONC, SONC, and SOSC. It is a strict local minimizer. Actually, $\boldsymbol{x} = [0,0]^\top$ is a strict global minimizer. Figure 6.8 shows the graph of $f(\boldsymbol{x}) = x_1^2 + x_2^2$. ∎

In this chapter we presented a theoretical basis for the solution of nonlinear unconstrained problems. In the following chapters we are concerned with iterative methods of solving such problems. Such methods are of great importance in practice. Indeed, suppose that one is confronted with a highly nonlinear function of 20 variables. Then, the FONC requires the solution of 20 nonlinear simultaneous equations for 20 variables. These equations, being nonlinear, will normally have multiple solutions. In addition, we would have to compute 210 second derivatives (provided that $f \in \mathcal{C}^2$) to use the SONC or SOSC. We begin our discussion of iterative methods in the next chapter with search methods for functions of one variable.

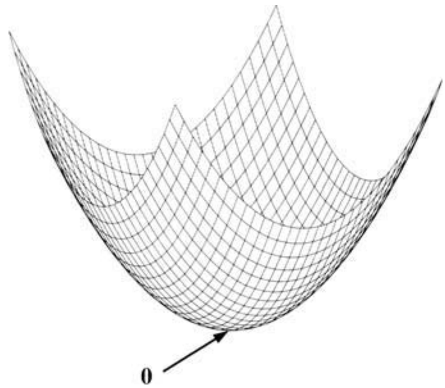

Figure 6.8 Graph of $f(\boldsymbol{x}) = x_1^2 + x_2^2$.

EXERCISES

6.1 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $f \in \mathcal{C}^2$. For each of the following specifications for Ω, $\boldsymbol{x}^*$, and f, determine if the given point $\boldsymbol{x}^*$ is: (i) definitely a local minimizer; (ii) definitely not a local minimizer; or (iii) possibly a local minimizer.

a. $f : \mathbb{R}^2 \to \mathbb{R}$, $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 \geq 1\}$, $\boldsymbol{x}^* = [1, 2]^\top$, and gradient $\nabla f(\boldsymbol{x}^*) = [1, 1]^\top$.

b. $f : \mathbb{R}^2 \to \mathbb{R}$, $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 \geq 1, x_2 \geq 2\}$, $\boldsymbol{x}^* = [1, 2]^\top$, and gradient $\nabla f(\boldsymbol{x}^*) = [1, 0]^\top$.

c. $f : \mathbb{R}^2 \to \mathbb{R}$, $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 \geq 0, x_2 \geq 0\}$, $\boldsymbol{x}^* = [1, 2]^\top$, gradient $\nabla f(\boldsymbol{x}^*) = [0, 0]^\top$, and Hessian $\boldsymbol{F}(\boldsymbol{x}^*) = \boldsymbol{I}$ (identity matrix).

d. $f : \mathbb{R}^2 \to \mathbb{R}$, $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 \geq 1, x_2 \geq 2\}$, $\boldsymbol{x}^* = [1, 2]^\top$, gradient $\nabla f(\boldsymbol{x}^*) = [1, 0]^\top$, and Hessian

$$\boldsymbol{F}(\boldsymbol{x}^*) = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

6.2 Find minimizers and maximizers of the function

$$f(x_1, x_2) = \frac{1}{3}x_1^3 - 4x_1 + \frac{1}{3}x_2^3 - 16x_2.$$

6.3 Show that if $\boldsymbol{x}^*$ is a global minimizer of f over Ω, and $\boldsymbol{x}^* \in \Omega' \subset \Omega$, then $\boldsymbol{x}^*$ is a global minimizer of f over Ω'.

6.4 Suppose that $\boldsymbol{x}^*$ is a local minimizer of f over Ω, and $\Omega \subset \Omega'$. Show that if $\boldsymbol{x}^*$ is an interior point of Ω, then $\boldsymbol{x}^*$ is a local minimizer of f over Ω'. Show that the same conclusion cannot be made if $\boldsymbol{x}^*$ is not an interior point of Ω.

6.5 Consider the problem of minimizing $f : \mathbb{R} \to \mathbb{R}$, $f \in \mathcal{C}^3$, over the constraint set Ω. Suppose that 0 is an *interior point* of Ω.

a. Suppose that 0 is a local minimizer. By the FONC we know that $f'(0) = 0$ (where f' is the first derivative of f). By the SONC we know that $f''(0) \geq 0$ (where f'' is the second derivative of f). State and prove a *third-order necessary condition* (*TONC*) involving the third derivative at 0, $f'''(0)$.

b. Give an example of f such that the FONC, SONC, and TONC (in part a) hold at the interior point 0, but 0 is not a local minimizer of f over Ω. (Show that your example is correct.)

c. Suppose that f is a third-order polynomial. If 0 satisfies the FONC, SONC, and TONC (in part a), then is this *sufficient* for 0 to be a local minimizer?

6.6 Consider the problem of minimizing $f : \mathbb{R} \to \mathbb{R}$, $f \in \mathcal{C}^3$, over the constraint set $\Omega = [0, 1]$. Suppose that $x^* = 0$ is a local minimizer.

a. By the FONC we know that $f'(0) \geq 0$ (where f' is the first derivative of f). By the SONC we know that if $f'(0) = 0$, then $f''(0) \geq 0$ (where f'' is the second derivative of f). State and prove a *third-order necessary condition* involving the third derivative at 0, $f'''(0)$.

b. Give an example of f such that the FONC, SONC, and TONC (in part a) hold at the point 0, but 0 is not a local minimizer of f over $\Omega = [0, 1]$.

6.7 Let $f : \mathbb{R}^n \to \mathbb{R}$, $\boldsymbol{x}_0 \in \mathbb{R}^n$, and $\Omega \subset \mathbb{R}^n$. Show that

$$\boldsymbol{x}_0 + \arg\min_{\boldsymbol{x} \in \Omega} f(\boldsymbol{x}) = \arg\min_{\boldsymbol{y} \in \Omega'} f(\boldsymbol{y}),$$

where $\Omega' = \{\boldsymbol{y} : \boldsymbol{y} - \boldsymbol{x}_0 \in \Omega\}$.

6.8 Consider the following function $f : \mathbb{R}^2 \to \mathbb{R}$:

$$f(\boldsymbol{x}) = \boldsymbol{x}^\top \begin{bmatrix} 1 & 2 \\ 4 & 7 \end{bmatrix} \boldsymbol{x} + \boldsymbol{x}^\top \begin{bmatrix} 3 \\ 5 \end{bmatrix} + 6.$$

a. Find the gradient and Hessian of f at the point $[1,1]^\top$.

b. Find the directional derivative of f at $[1,1]^\top$ with respect to a unit vector in the direction of maximal rate of increase.

c. Find a point that satisfies the FONC (interior case) for f. Does this point satisfy the SONC (for a minimizer)?

6.9 Consider the following function:

$$f(x_1, x_2) = x_1^2 x_2 + x_2^3 x_1.$$

a. In what direction does the function f *decrease* most rapidly at the point $\boldsymbol{x}^{(0)} = [2,1]^\top$?

b. What is the rate of increase of f at the point $\boldsymbol{x}^{(0)}$ in the direction of maximum decrease of f?

c. Find the rate of increase of f at the point $\boldsymbol{x}^{(0)}$ in the direction $\boldsymbol{d} = [3,4]^\top$.

6.10 Consider the following function $f : \mathbb{R}^2 \to \mathbb{R}$:

$$f(\boldsymbol{x}) = \boldsymbol{x}^\top \begin{bmatrix} 2 & 5 \\ -1 & 1 \end{bmatrix} \boldsymbol{x} + \boldsymbol{x}^\top \begin{bmatrix} 3 \\ 4 \end{bmatrix} + 7.$$

a. Find the directional derivative of f at $[0,1]^\top$ in the direction $[1,0]^\top$.

b. Find all points that satisfy the first-order necessary condition for f.

Does f have a minimizer? If it does, then find all minimizer(s); otherwise, explain why it does not.

6.11 Consider the problem

$$\begin{array}{rl} \text{minimize} & -x_2^2 \\ \text{subject to} & |x_2| \leq x_1^2 \\ & x_1 \geq 0, \end{array}$$

where $x_1, x_2 \in \mathbb{R}$.

a. Does the point $[x_1, x_2]^\top = \mathbf{0}$ satisfy the first-order necessary condition for a minimizer? That is, if f is the objective function, is it true that $\boldsymbol{d}^\top \nabla f(\mathbf{0}) \geq 0$ for all feasible directions $\boldsymbol{d}$ at $\mathbf{0}$?

b. Is the point $[x_1, x_2]^\top = \mathbf{0}$ a local minimizer, a strict local minimizer, a local maximizer, a strict local maximizer, or none of the above?

6.12 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ is given by $f(\boldsymbol{x}) = 5x_2$ with $\boldsymbol{x} = [x_1, x_2]^\top$, and $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1^2 + x_2 \geq 1\}$.

a. Does the point $\boldsymbol{x}^* = [0, 1]^\top$ satisfy the first-order necessary condition?

b. Does the point $\boldsymbol{x}^* = [0, 1]^\top$ satisfy the second-order necessary condition?

c. Is the point $\boldsymbol{x}^* = [0, 1]^\top$ a local minimizer?

6.13 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ is given by $f(\boldsymbol{x}) = -3x_1$ with $\boldsymbol{x} = [x_1, x_2]^\top$, and $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 + x_2^2 \leq 2\}$. Answer each of the following questions, showing complete justification.

a. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the first-order necessary condition?

b. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the second-order necessary condition?

c. Is the point $\boldsymbol{x}^* = [2, 0]^\top$ a local minimizer?

6.14 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $\Omega = \{\boldsymbol{x} \in \mathbb{R}^2 : x_1^2 + x_2^2 \geq 1\}$ and $f(\boldsymbol{x}) = x_2$.

a. Find all point(s) satisfying the FONC.

b. Which of the point(s) in part a satisfy the SONC?

c. Which of the point(s) in part a are local minimizers?

6.15 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega \end{array}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ is given by $f(\boldsymbol{x}) = 3x_1$ with $\boldsymbol{x} = [x_1, x_2]^\top$, and $\Omega = \{\boldsymbol{x} = [x_1, x_2]^\top : x_1 + x_2^2 \geq 2\}$. Answer each of the following questions, showing complete justification.

a. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the first-order necessary condition?

b. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the second-order necessary condition?

c. Is the point $\boldsymbol{x}^* = [2, 0]^\top$ a local minimizer?
Hint: Draw a picture with the constraint set and level sets of f.

6.16 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $\boldsymbol{x} = [x_1, x_2]^\top$, $f : \mathbb{R}^2 \to \mathbb{R}$ is given by $f(\boldsymbol{x}) = 4x_1^2 - x_2^2$, and $\Omega = \{\boldsymbol{x} : x_1^2 + 2x_1 - x_2 \geq 0, x_1 \geq 0, x_2 \geq 0\}$.

a. Does the point $\boldsymbol{x}^* = \boldsymbol{0} = [0, 0]^\top$ satisfy the first-order necessary condition?

b. Does the point $\boldsymbol{x}^* = \boldsymbol{0}$ satisfy the second-order necessary condition?

c. Is the point $\boldsymbol{x}^* = \boldsymbol{0}$ a local minimizer of the given problem?

6.17 Consider the problem

$$\begin{array}{ll} \text{maximize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $\Omega \subset \{\boldsymbol{x} \in \mathbb{R}^2 : x_1 > 0, x_2 > 0\}$ and $f : \Omega \to \mathbb{R}$ is given by $f(\boldsymbol{x}) = \log(x_1) + \log(x_2)$ with $\boldsymbol{x} = [x_1, x_2]^\top$, where "log" represents natural logarithm. Suppose that $\boldsymbol{x}^*$ is an optimal solution. Answer each of the following questions, showing complete justification.

a. Is it possible that $\boldsymbol{x}^*$ is an interior point of Ω?

b. At what point(s) (if any) is the second-order necessary condition satisfied?

6.18 Suppose that we are given n real numbers, $x_1, \ldots, x_n$. Find the number $\bar{x} \in \mathbb{R}$ such that the sum of the squared difference between $\bar{x}$ and the numbers above is minimized (assuming that the solution $\bar{x}$ exists).

6.19 An art collector stands at a distance of x feet from the wall, where a piece of art (picture) of height a feet is hung, b feet above his eyes, as shown in

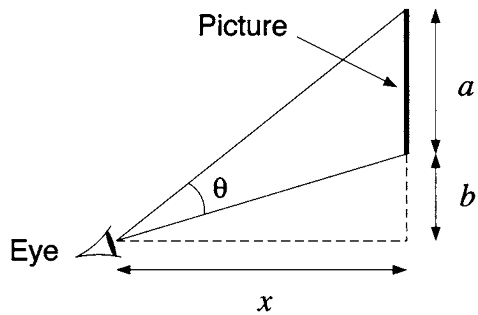

Figure 6.9 Art collector's eye position in Exercise 6.19.

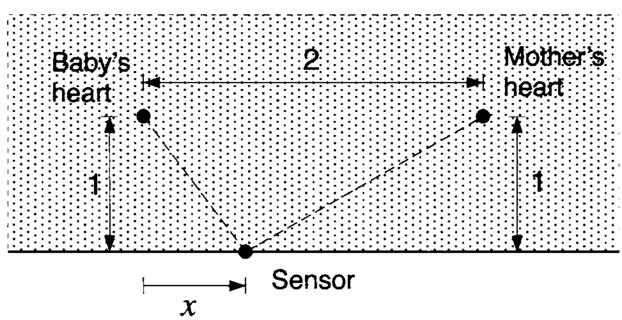

Figure 6.10 Simplified fetal heart monitoring system for Exercise 6.20.

Figure 6.9. Find the distance from the wall for which the angle θ subtended by the eye to the picture is maximized.
Hint: (1) Maximizing θ is equivalent to maximizing $\tan(\theta)$.
(2) If $\theta = \theta_2 - \theta_1$, then $\tan(\theta) = (\tan(\theta_2) - \tan(\theta_1))/(1 + \tan(\theta_2)\tan(\theta_1))$.

6.20 Figure 6.10 shows a simplified model of a fetal heart monitoring system (the distances shown have been scaled down to make the calculations simpler). A heartbeat sensor is located at position x (see Figure 6.10).

The energy of the heartbeat signal measured by the sensor is the reciprocal of the squared distance from the source (baby's heart or mother's heart). Find the position of the sensor that maximizes the *signal-to-interference ratio*, which is the ratio of the signal energy from the baby's heart to the signal energy from the mother's heart.

6.21 An amphibian vehicle needs to travel from point A (on land) to point B (in water), as illustrated in Figure 6.11. The speeds at which the vehicle travels on land and water are v_1 and v_2, respectively.

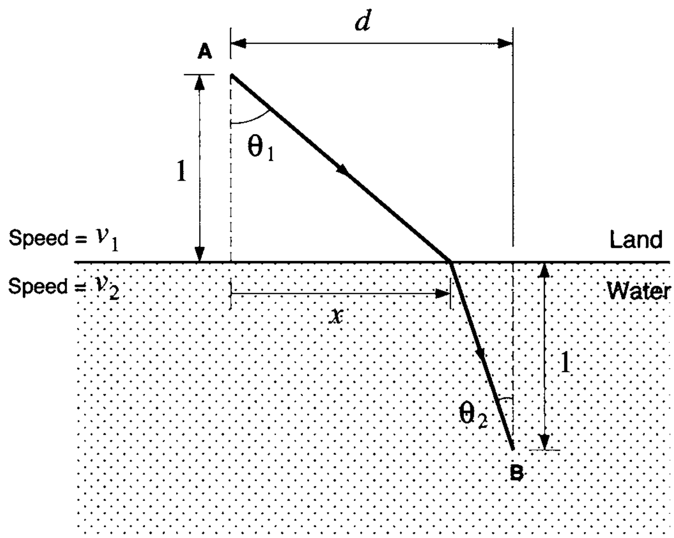

Figure 6.11 Path of amphibian vehicle in Exercise 6.21.

a. Suppose that the vehicle traverses a path that minimizes the total time taken to travel from A to B. Use the first-order necessary condition to show that for the optimal path above, the angles θ_1 and θ_2 in Figure 6.11 satisfy Snell's law:

$$\frac{\sin\theta_1}{\sin\theta_2} = \frac{v_1}{v_2}.$$

b. Does the minimizer for the problem in part a satisfy the second-order sufficient condition?

6.22 Suppose that you have a piece of land to sell and you have two buyers. If the first buyer receives a fraction x_1 of the piece of land, the buyer will pay you $U_1(x_1)$ dollars. Similarly, the second buyer will pay you $U_2(x_2)$ dollars for a fraction of x_2 of the land. Your goal is to sell parts of your land to the two buyers so that you maximize the total dollars you receive. (Other than the constraint that you can only sell whatever land you own, there are no restrictions on how much land you can sell to each buyer.)

a. Formulate the problem as an optimization problem of the kind

$$\begin{array}{ll} \text{maximize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega \end{array}$$

by specifying f and Ω. Draw a picture of the constraint set.

b. Suppose that $U_i(x_i) = a_i x_i$, $i = 1, 2$, where a_1 and a_2 are given positive constants such that $a_1 > a_2$. Find all feasible points that satisfy the first-order necessary condition, giving full justification.

c. Among those points in the answer of part b, find all that also satisfy the second-order necessary condition.

6.23 Let $f : \mathbb{R}^2 \to \mathbb{R}$ be defined by

$$f(\boldsymbol{x}) = (x_1 - x_2)^4 + x_1^2 - x_2^2 - 2x_1 + 2x_2 + 1,$$

where $\boldsymbol{x} = [x_1, x_2]^\top$. Suppose that we wish to minimize f over $\mathbb{R}^2$. Find all points satisfying the FONC. Do these points satisfy the SONC?

6.24 Show that if $\boldsymbol{d}$ is a feasible direction at a point $\boldsymbol{x} \in \Omega$, then for all $\beta > 0$, the vector $\beta\boldsymbol{d}$ is also a feasible direction at $\boldsymbol{x}$.

6.25 Let $\Omega = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}\}$. Show that $\boldsymbol{d} \in \mathbb{R}^n$ is a feasible direction at $\boldsymbol{x} \in \Omega$ if and only if $\boldsymbol{A}\boldsymbol{d} = \boldsymbol{0}$.

6.26 Let $f : \mathbb{R}^2 \to \mathbb{R}$. Consider the problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & x_1, x_2 \geq 0, \end{array}$$

where $\boldsymbol{x} = [x_1, x_2]^\top$. Suppose that $\nabla f(\boldsymbol{0}) \neq \boldsymbol{0}$, and

$$\frac{\partial f}{\partial x_1}(\boldsymbol{0}) \leq 0, \qquad \frac{\partial f}{\partial x_2}(\boldsymbol{0}) \leq 0.$$

Show that $\boldsymbol{0}$ cannot be a minimizer for this problem.

6.27 Let $\boldsymbol{c} \in \mathbb{R}^n$, $\boldsymbol{c} \neq \boldsymbol{0}$, and consider the problem of minimizing the function $f(\boldsymbol{x}) = \boldsymbol{c}^\top\boldsymbol{x}$ over a constraint set $\Omega \subset \mathbb{R}^n$. Show that we cannot have a solution lying in the interior of Ω.

6.28 Consider the problem

$$\begin{array}{ll} \text{maximize} & c_1 x_1 + c_2 x_2 \\ \text{subject to} & x_1 + x_2 \leq 1 \\ & x_1, x_2 \geq 0, \end{array}$$

where c_1 and c_2 are constants such that $c_1 > c_2 \geq 0$. This is a *linear programming* problem (see Part III). Assuming that the problem has an optimal feasible solution, use the first-order necessary condition to show that the *unique* optimal feasible solution $\boldsymbol{x}^*$ is $[1, 0]^\top$.

Hint: First show that $\boldsymbol{x}^*$ cannot lie in the interior of the constraint set. Then, show that $\boldsymbol{x}^*$ cannot lie on the line segments $L_1 = \{\boldsymbol{x} : x_1 = 0, 0 \leq x_2 < 1\}$, $L_2 = \{\boldsymbol{x} : 0 \leq x_1 < 1, x_2 = 0\}$, $L_3 = \{\boldsymbol{x} : 0 \leq x_1 < 1, x_2 = 1 - x_1\}$.

6.29 *Line Fitting.* Let $[x_1, y_1]^\top, \ldots, [x_n, y_n]^\top$, $n \geq 2$, be points on the $\mathbb{R}^2$ plane (each $x_i, y_i \in \mathbb{R}$). We wish to find the straight line of "best fit" through these points ("best" in the sense that the average squared error is minimized); that is, we wish to find $a, b \in \mathbb{R}$ to minimize

$$f(a,b) = \frac{1}{n}\sum_{i=1}^{n}(ax_i + b - y_i)^2 .$$

a. Let

$$\begin{aligned}
\overline{X} &= \frac{1}{n}\sum_{i=1}^{n} x_i, \\
\overline{Y} &= \frac{1}{n}\sum_{i=1}^{n} y_i, \\
\overline{X^2} &= \frac{1}{n}\sum_{i=1}^{n} x_i^2, \\
\overline{Y^2} &= \frac{1}{n}\sum_{i=1}^{n} y_i^2, \\
\overline{XY} &= \frac{1}{n}\sum_{i=1}^{n} x_i y_i.
\end{aligned}$$

Show that $f(a,b)$ can be written in the form $\boldsymbol{z}^\top \boldsymbol{Q}\boldsymbol{z} - 2\boldsymbol{c}^\top \boldsymbol{z} + d$, where $\boldsymbol{z} = [a,b]^\top$, $\boldsymbol{Q} = \boldsymbol{Q}^\top \in \mathbb{R}^{2\times 2}$, $\boldsymbol{c} \in \mathbb{R}^2$ and $d \in \mathbb{R}$, and find expressions for $\boldsymbol{Q}$, $\boldsymbol{c}$, and d in terms of $\overline{X}$, $\overline{Y}$, $\overline{X^2}$, $\overline{Y^2}$, and $\overline{XY}$.

b. Assume that the x_i, $i = 1, \ldots, n$, are not all equal. Find the parameters a^* and b^* for the line of best fit in terms of $\overline{X}$, $\overline{Y}$, $\overline{X^2}$, $\overline{Y^2}$, and $\overline{XY}$. Show that the point $[a^*, b^*]^\top$ is the only local minimizer of f.
Hint: $\overline{X^2} - (\overline{X})^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{X})^2$.

c. Show that if a^* and b^* are the parameters of the line of best fit, then $\overline{Y} = a^*\overline{X} + b^*$ (and hence once we have computed a^*, we can compute b^* using the formula $b^* = \overline{Y} - a^*\overline{X}$).

6.30 Suppose that we are given a set of vectors $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(p)}\}$, $\boldsymbol{x}^{(i)} \in \mathbb{R}^n$, $i = 1, \ldots, p$. Find the vector $\bar{\boldsymbol{x}} \in \mathbb{R}^n$ such that the average squared distance (norm) between $\bar{\boldsymbol{x}}$ and $\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(p)}$,

$$\frac{1}{p}\sum_{i=1}^{p}\|\bar{\boldsymbol{x}} - \boldsymbol{x}^{(i)}\|^2,$$

is minimized. Use the SOSC to prove that the vector $\bar{\boldsymbol{x}}$ found above is a strict local minimizer. How is $\bar{\boldsymbol{x}}$ related to the centroid (or center of gravity) of the given set of points $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(p)}\}$?

6.31 Consider a function $f : \Omega \to \mathbb{R}$, where $\Omega \subset \mathbb{R}^n$ is a convex set and $f \in \mathcal{C}^1$. Given $\boldsymbol{x}^* \in \Omega$, suppose that there exists $c > 0$ such that $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq c\|\boldsymbol{d}\|$ for all feasible directions $\boldsymbol{d}$ at $\boldsymbol{x}^*$. Show that $\boldsymbol{x}^*$ is a strict local minimizer of f over Ω.

6.32 Prove the following generalization of the second-order sufficient condition:

Theorem: Let Ω be a convex subset of $\mathbb{R}^n$, $f \in \mathcal{C}^2$ a real-valued function on Ω, and $\boldsymbol{x}^*$ a point in Ω. Suppose that there exists $c \in \mathbb{R}$, $c > 0$, such that for all feasible directions $\boldsymbol{d}$ at $\boldsymbol{x}^*$ ($\boldsymbol{d} \neq \boldsymbol{0}$), the following hold:

1. $\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq 0$.
2. $\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{d} \geq c\|\boldsymbol{d}\|^2$.

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

6.33 Consider the quadratic function $f : \mathbb{R}^n \to \mathbb{R}$ given by

$$f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - \boldsymbol{x}^\top \boldsymbol{b},$$

where $\boldsymbol{Q} = \boldsymbol{Q}^\top > 0$. Show that $\boldsymbol{x}^*$ minimizes f if and only if $\boldsymbol{x}^*$ satisfies the FONC.

6.34 Consider the linear system $x_{k+1} = ax_k + bu_{k+1}$, $k \geq 0$, where $x_i \in \mathbb{R}$, $u_i \in \mathbb{R}$, and the initial condition is $x_0 = 0$. Find the values of the control inputs $u_1, \ldots, u_n$ to minimize

$$-qx_n + r\sum_{i=1}^{n} u_i^2,$$

where $q, r > 0$ are given constants. This can be interpreted as desiring to make x_n as large as possible but at the same time desiring to make the total input energy $\sum_{i=1}^{n} u_i^2$ as small as possible. The constants q and r reflect the relative weights of these two objectives.

PART IV

NONLINEAR CONSTRAINED OPTIMIZATION

CHAPTER 20

PROBLEMS WITH EQUALITY CONSTRAINTS

20.1 Introduction

In this part we discuss methods for solving a class of nonlinear constrained optimization problems that can be formulated as

$$\begin{aligned} \text{minimize} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & h_i(\boldsymbol{x}) = 0, \qquad i = 1, \ldots, m \\ & g_j(\boldsymbol{x}) \leq 0, \qquad j = 1, \ldots, p, \end{aligned}$$

where $\boldsymbol{x} \in \mathbb{R}^n$, $f : \mathbb{R}^n \to \mathbb{R}$, $h_i : \mathbb{R}^n \to \mathbb{R}$, $g_j : \mathbb{R}^n \to \mathbb{R}$, and $m \leq n$. In vector notation, the problem above can be represented in the following *standard form*:

$$\begin{aligned} \text{minimize} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0} \\ & \boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}, \end{aligned}$$

where $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$ and $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^p$. As usual, we adopt the following terminology.

Definition 20.1 Any point satisfying the constraints is called a *feasible point.* The set of all feasible points,

$$\{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0},\ \boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}\},$$

is called a *feasible set.* ∎

Optimization problems of the above form are not new to us. Indeed, linear programming problems of the form

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \\ & \boldsymbol{x} \geq \mathbf{0}, \end{aligned}$$

which we studied in Part III, are of this type.

As we remarked in Part II, there is no loss of generality by considering only minimization problems. For if we are confronted with a maximization problem, it can easily be transformed into the minimization problem by observing that

$$\text{maximize}\, f(\boldsymbol{x}) = \text{minimize}\, -f(\boldsymbol{x}).$$

We illustrate the problems we study in this part by considering the following simple numerical example.

Example 20.1 Consider the following optimization problem:

$$\begin{aligned} \text{minimize} \quad & (x_1 - 1)^2 + x_2 - 2 \\ \text{subject to} \quad & x_2 - x_1 = 1, \\ & x_1 + x_2 \leq 2. \end{aligned}$$

This problem is already in the standard form given earlier, with $f(x_1, x_2) = (x_1 - 1)^2 + x_2 - 2$, $h(x_1, x_2) = x_2 - x_1 - 1$, and $g(x_1, x_2) = x_1 + x_2 - 2$. This problem turns out to be simple enough to be solved graphically (see Figure 20.1). In the figure the set of points that satisfy the constraints (the feasible set) is marked by the heavy solid line. The inverted parabolas represent level sets of the objective function f—the lower the level set, the smaller the objective function value. Therefore, the solution can be obtained by finding the lowest-level set that intersects the feasible set. In this case, the minimizer lies on the level set with $f = -1/4$. The minimizer of the objective function is $\boldsymbol{x}^* = [1/2, 3/2]^\top$. ∎

In the remainder of this chapter we discuss constrained optimization problems with only equality constraints. The general constrained optimization problem is discussed in the chapters to follow.

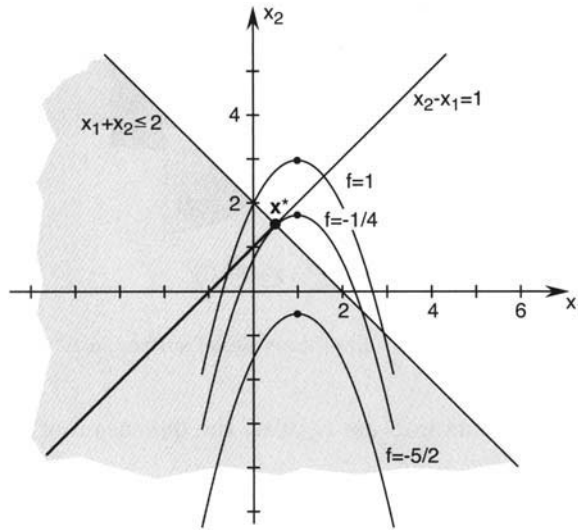

Figure 20.1 Graphical solution to the problem in Example 20.1.

20.2 Problem Formulation

The class of optimization problems we analyze in this chapter is

$$\begin{array}{rl} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}, \end{array}$$

where $\boldsymbol{x} \in \mathbb{R}^n$, $f : \mathbb{R}^n \to \mathbb{R}$, $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $\boldsymbol{h} = [h_1, \ldots, h_m]^\top$, and $m \leq n$. We assume that the function $\boldsymbol{h}$ is continuously differentiable, that is, $\boldsymbol{h} \in \mathcal{C}^1$.

We introduce the following definition.

Definition 20.2 A point $\boldsymbol{x}^*$ satisfying the constraints $h_1(\boldsymbol{x}^*) = 0, \ldots, h_m(\boldsymbol{x}^*) = 0$ is said to be a *regular point* of the constraints if the gradient vectors $\nabla h_1(\boldsymbol{x}^*), \ldots, \nabla h_m(\boldsymbol{x}^*)$ are linearly independent. ∎

Let $D\boldsymbol{h}(\boldsymbol{x}^*)$ be the Jacobian matrix of $\boldsymbol{h} = [h_1, \ldots, h_m]^\top$ at $\boldsymbol{x}^*$, given by

$$D\boldsymbol{h}(\boldsymbol{x}^*) = \begin{bmatrix} Dh_1(\boldsymbol{x}^*) \\ \vdots \\ Dh_m(\boldsymbol{x}^*) \end{bmatrix} = \begin{bmatrix} \nabla h_1(\boldsymbol{x}^*)^\top \\ \vdots \\ \nabla h_m(\boldsymbol{x}^*)^\top \end{bmatrix}.$$

Then, $\boldsymbol{x}^*$ is regular if and only if $\operatorname{rank} D\boldsymbol{h}(\boldsymbol{x}^*) = m$ (i.e., the Jacobian matrix is of full rank).

The set of equality constraints $h_1(\boldsymbol{x}) = 0, \ldots, h_m(\boldsymbol{x}) = 0$, $h_i : \mathbb{R}^n \to \mathbb{R}$, describes a surface

$$S = \{\boldsymbol{x} \in \mathbb{R}^n : h_1(\boldsymbol{x}) = 0, \ldots, h_m(\boldsymbol{x}) = 0\}.$$

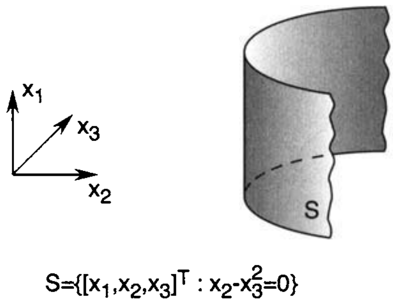

Figure 20.2 Two-dimensional surface in $\mathbb{R}^3$.

Assuming that the points in S are regular, the dimension of the surface S is $n - m$.

Example 20.2 Let $n = 3$ and $m = 1$ (i.e., we are operating in $\mathbb{R}^3$). Assuming that all points in S are regular, the set S is a two-dimensional surface. For example, let

$$h_1(\boldsymbol{x}) = x_2 - x_3^2 = 0.$$

Note that $\nabla h_1(\boldsymbol{x}) = [0, 1, -2x_3]^\top$, and hence for any $\boldsymbol{x} \in \mathbb{R}^3$, $\nabla h_1(\boldsymbol{x}) \neq \boldsymbol{0}$. In this case,

$$\dim S = \dim\{\boldsymbol{x} : h_1(\boldsymbol{x}) = 0\} = n - m = 2.$$

See Figure 20.2 for a graphical illustration. ■

Example 20.3 Let $n = 3$ and $m = 2$. Assuming regularity, the feasible set S is a one-dimensional object (i.e., a curve in $\mathbb{R}^3$). For example, let

$$\begin{aligned} h_1(\boldsymbol{x}) &= x_1, \\ h_2(\boldsymbol{x}) &= x_2 - x_3^2. \end{aligned}$$

In this case, $\nabla h_1(\boldsymbol{x}) = [1, 0, 0]^\top$ and $\nabla h_2(\boldsymbol{x}) = [0, 1, -2x_3]^\top$. Hence, the vectors $\nabla h_1(\boldsymbol{x})$ and $\nabla h_2(\boldsymbol{x})$ are linearly independent in $\mathbb{R}^3$. Thus,

$$\dim S = \dim\{\boldsymbol{x} : h_1(\boldsymbol{x}) = 0,\ h_2(\boldsymbol{x}) = 0\} = n - m = 1.$$

See Figure 20.3 for a graphical illustration. ■

20.3 Tangent and Normal Spaces

In this section we discuss the notion of a tangent space and normal space at a point on a surface. We begin by defining a *curve* on a surface S.

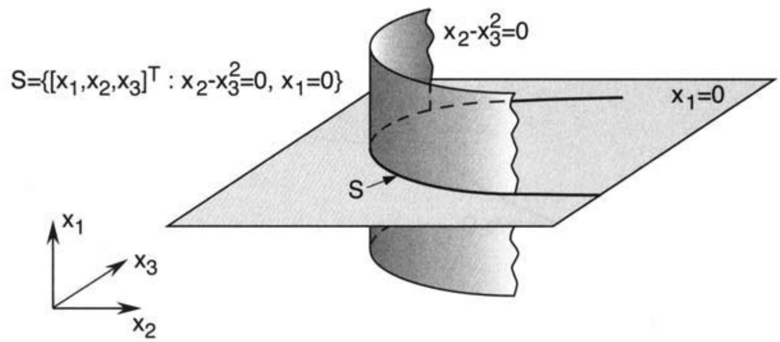

Figure 20.3 One-dimensional surface in $\mathbb{R}^3$.

Definition 20.3 A *curve* C on a surface S is a set of points $\{\boldsymbol{x}(t) \in S : t \in (a,b)\}$, continuously parameterized by $t \in (a,b)$; that is, $\boldsymbol{x} : (a,b) \to S$ is a continuous function. ∎

A graphical illustration of the definition of a curve is given in Figure 20.4. The definition of a curve implies that all the points on the curve satisfy the equation describing the surface. The curve C passes through a point $\boldsymbol{x}^*$ if there exists $t^* \in (a,b)$ such that $\boldsymbol{x}(t^*) = \boldsymbol{x}^*$.

Intuitively, we can think of a curve $C = \{\boldsymbol{x}(t) : t \in (a,b)\}$ as the path traversed by a point $\boldsymbol{x}$ traveling on the surface S. The position of the point at time t is given by $\boldsymbol{x}(t)$.

Definition 20.4 The curve $C = \{\boldsymbol{x}(t) : t \in (a,b)\}$ is *differentiable* if

$$\dot{\boldsymbol{x}}(t) = \frac{d\boldsymbol{x}}{dt}(t) = \begin{bmatrix} \dot{x}_1(t) \\ \vdots \\ \dot{x}_n(t) \end{bmatrix}$$

exists for all $t \in (a,b)$.

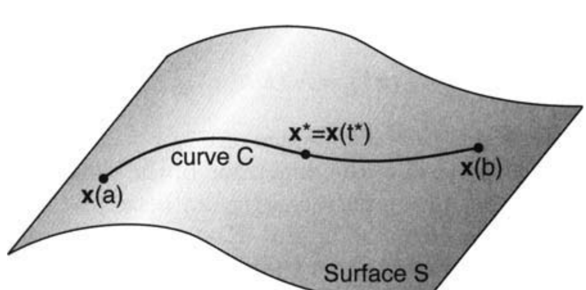

Figure 20.4 Curve on a surface.

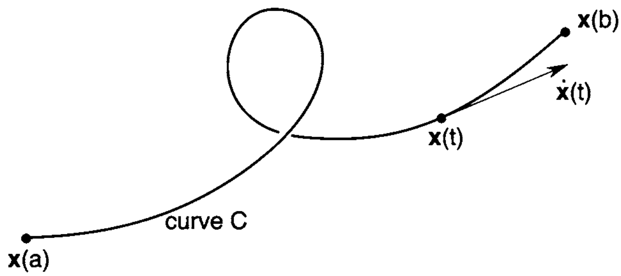

Figure 20.5 Geometric interpretation of the differentiability of a curve.

The curve $C = \{\boldsymbol{x}(t) : t \in (a, b)\}$ is *twice differentiable* if

$$\ddot{\boldsymbol{x}}(t) = \frac{d^2\boldsymbol{x}}{dt^2}(t) = \begin{bmatrix} \ddot{x}_1(t) \\ \vdots \\ \ddot{x}_n(t) \end{bmatrix}$$

exists for all $t \in (a, b)$. ∎

Note that both $\dot{\boldsymbol{x}}(t)$ and $\ddot{\boldsymbol{x}}(t)$ are n-dimensional vectors. We can think of $\dot{\boldsymbol{x}}(t)$ and $\ddot{\boldsymbol{x}}(t)$ as the velocity and acceleration, respectively, of a point traversing the curve C with position $\boldsymbol{x}(t)$ at time t. The vector $\dot{\boldsymbol{x}}(t)$ points in the direction of the instantaneous motion of $\boldsymbol{x}(t)$. Therefore, the vector $\dot{\boldsymbol{x}}(t^*)$ is *tangent* to the curve C at $\boldsymbol{x}^*$ (see Figure 20.5).

We are now ready to introduce the notions of a tangent space. For this recall the set

$$S = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\},$$

where $\boldsymbol{h} \in \mathcal{C}^1$. We think of S as a surface in $\mathbb{R}^n$.

Definition 20.5 The *tangent space* at a point $\boldsymbol{x}^*$ on the surface $S = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\}$ is the set $T(\boldsymbol{x}^*) = \{\boldsymbol{y} : D\boldsymbol{h}(\boldsymbol{x}^*)\boldsymbol{y} = \boldsymbol{0}\}$. ∎

Note that the tangent space $T(\boldsymbol{x}^*)$ is the nullspace of the matrix $D\boldsymbol{h}(\boldsymbol{x}^*)$:

$$T(\boldsymbol{x}^*) = \mathcal{N}(D\boldsymbol{h}(\boldsymbol{x}^*)).$$

The tangent space is therefore a subspace of $\mathbb{R}^n$.

Assuming that $\boldsymbol{x}^*$ is regular, the dimension of the tangent space is $n - m$, where m is the number of equality constraints $h_i(\boldsymbol{x}^*) = 0$. Note that the tangent space passes through the origin. However, it is often convenient to picture the tangent space as a plane that passes through the point $\boldsymbol{x}^*$. For this, we define the *tangent plane* at $\boldsymbol{x}^*$ to be the set

$$TP(\boldsymbol{x}^*) = T(\boldsymbol{x}^*) + \boldsymbol{x}^* = \{\boldsymbol{x} + \boldsymbol{x}^* : \boldsymbol{x} \in T(\boldsymbol{x}^*)\}.$$

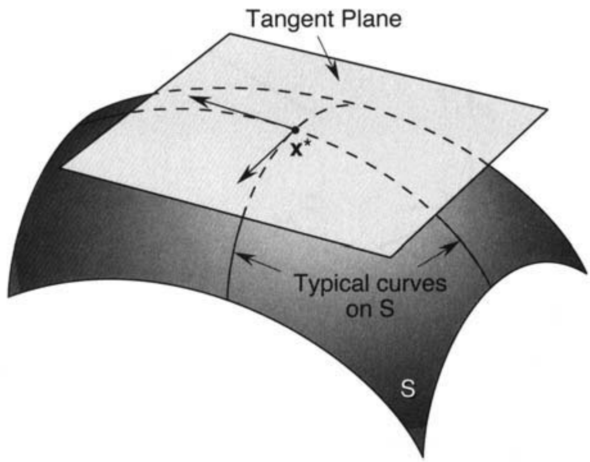

Figure 20.6 Tangent plane to the surface S at the point $\boldsymbol{x}^*$.

Figure 20.6 illustrates the notion of a tangent plane, and Figure 20.7, the relationship between the tangent plane and the tangent space.

Example 20.4 Let

$$S = \{\boldsymbol{x} \in \mathbb{R}^3 : h_1(\boldsymbol{x}) = x_1 = 0,\ h_2(\boldsymbol{x}) = x_1 - x_2 = 0\}.$$

Then, S is the x_3-axis in $\mathbb{R}^3$ (see Figure 20.8). We have

$$D\boldsymbol{h}(\boldsymbol{x}) = \begin{bmatrix} \nabla h_1(\boldsymbol{x})^\top \\ \nabla h_2(\boldsymbol{x})^\top \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & -1 & 0 \end{bmatrix}.$$

Because ∇h_1 and ∇h_2 are linearly independent when evaluated at any $\boldsymbol{x} \in S$, all the points of S are regular. The tangent space at an arbitrary point of S is

$$\begin{aligned} T(\boldsymbol{x}) &= \{\boldsymbol{y} : \nabla h_1(\boldsymbol{x})^\top \boldsymbol{y} = 0,\ \nabla h_2(\boldsymbol{x})^\top \boldsymbol{y} = 0\} \\ &= \left\{ \boldsymbol{y} : \begin{bmatrix} 1 & 0 & 0 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \boldsymbol{0} \right\} \\ &= \{[0, 0, \alpha]^\top : \alpha \in \mathbb{R}\} \\ &= \text{the } x_3\text{-axis in } \mathbb{R}^3. \end{aligned}$$

In this example, the tangent space $T(\boldsymbol{x})$ at any point $\boldsymbol{x} \in S$ is a one-dimensional subspace of $\mathbb{R}^3$. ∎

Intuitively, we would expect the definition of the tangent space at a point on a surface to be the collection of all "tangent vectors" to the surface at that

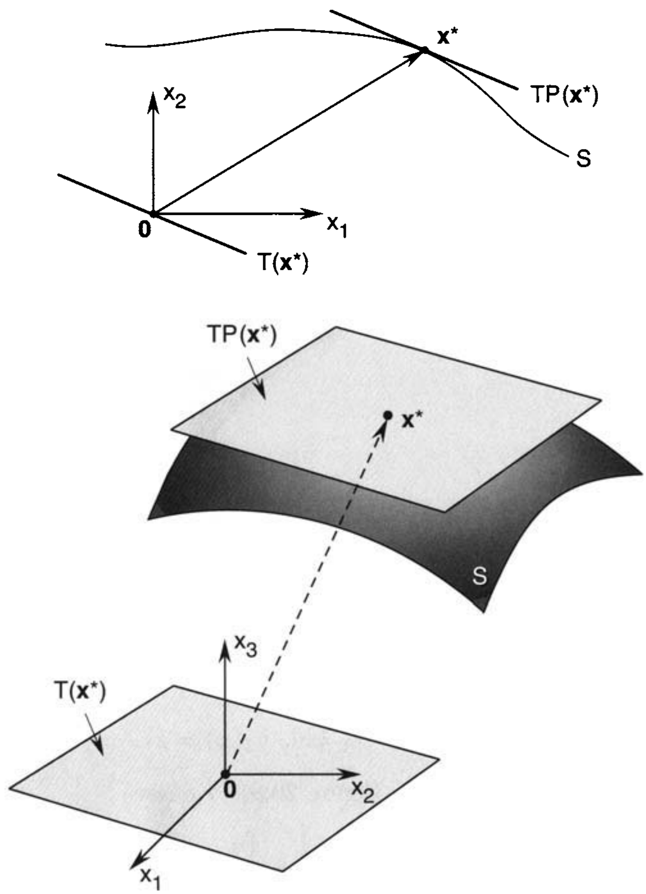

Figure 20.7 Tangent spaces and planes in $\mathbb{R}^2$ and $\mathbb{R}^3$.

point. We have seen that the derivative of a curve on a surface at a point is a tangent vector to the curve, and hence to the surface. The intuition above agrees with our definition whenever $\boldsymbol{x}^*$ is regular, as stated in the theorem below.

Theorem 20.1 *Suppose that $\boldsymbol{x}^* \in S$ is a regular point and $T(\boldsymbol{x}^*)$ is the tangent space at $\boldsymbol{x}^*$. Then, $\boldsymbol{y} \in T(\boldsymbol{x}^*)$ if and only if there exists a differentiable curve in S passing through $\boldsymbol{x}^*$ with derivative $\boldsymbol{y}$ at $\boldsymbol{x}^*$.* □

Proof. $\Leftarrow$: Suppose that there exists a curve $\{\boldsymbol{x}(t) : t \in (a,b)\}$ in S such that $\boldsymbol{x}(t^*) = \boldsymbol{x}^*$ and $\dot{\boldsymbol{x}}(t^*) = \boldsymbol{y}$ for some $t^* \in (a,b)$. Then,

$$\boldsymbol{h}(\boldsymbol{x}(t)) = \boldsymbol{0}$$

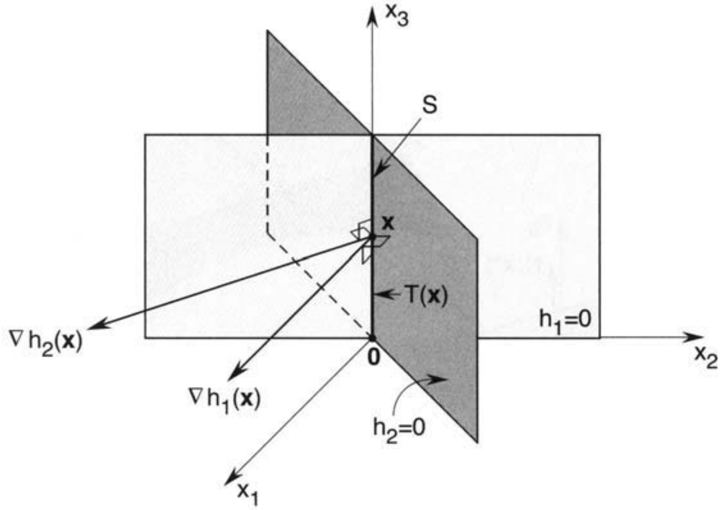

Figure 20.8 The surface $S = \{\boldsymbol{x} \in \mathbb{R}^3 : x_1 = 0,\ x_1 - x_2 = 0\}$.

for all $t \in (a, b)$. If we differentiate the function $\boldsymbol{h}(\boldsymbol{x}(t))$ with respect to t using the chain rule, we obtain

$$\frac{d}{dt}\boldsymbol{h}(\boldsymbol{x}(t)) = D\boldsymbol{h}(\boldsymbol{x}(t))\dot{\boldsymbol{x}}(t) = \mathbf{0}$$

for all $t \in (a, b)$. Therefore, at t^* we get

$$D\boldsymbol{h}(\boldsymbol{x}^*)\boldsymbol{y} = \mathbf{0},$$

and hence $\boldsymbol{y} \in T(\boldsymbol{x}^*)$.

$\Rightarrow$: To prove this, we need to use the implicit function theorem. We refer the reader to [88, p. 325]. ∎

We now introduce the notion of a normal space.

Definition 20.6 The *normal space* $N(\boldsymbol{x}^*)$ at a point $\boldsymbol{x}^*$ on the surface $S = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0}\}$ is the set $N(\boldsymbol{x}^*) = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} = D\boldsymbol{h}(\boldsymbol{x}^*)^\top \boldsymbol{z},\ \boldsymbol{z} \in \mathbb{R}^m\}$. ∎

We can express the normal space $N(\boldsymbol{x}^*)$ as

$$N(\boldsymbol{x}^*) = \mathcal{R}\left(D\boldsymbol{h}(\boldsymbol{x}^*)^\top\right),$$

that is, the range of the matrix $D\boldsymbol{h}(\boldsymbol{x}^*)^\top$. Note that the normal space $N(\boldsymbol{x}^*)$ is the subspace of $\mathbb{R}^n$ spanned by the vectors $\nabla h_1(\boldsymbol{x}^*), \ldots, \nabla h_m(\boldsymbol{x}^*)$; that is,

$$\begin{aligned} N(\boldsymbol{x}^*) &= \operatorname{span}[\nabla h_1(\boldsymbol{x}^*), \ldots, \nabla h_m(\boldsymbol{x}^*)] \\ &= \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} = z_1 \nabla h_1(\boldsymbol{x}^*) + \cdots + z_m \nabla h_m(\boldsymbol{x}^*),\ z_1, \ldots, z_m \in \mathbb{R}\}. \end{aligned}$$

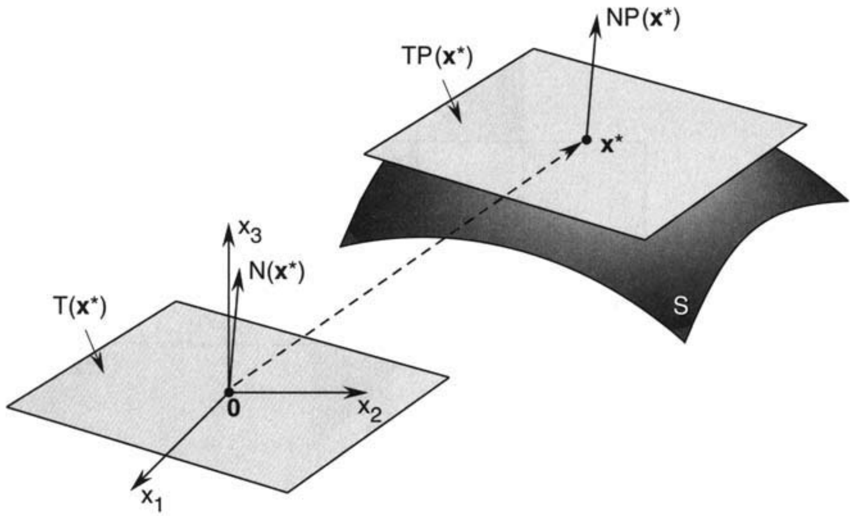

Figure 20.9 Normal space in $\mathbb{R}^3$.

Note that the normal space contains the zero vector. Assuming that $\boldsymbol{x}^*$ is regular, the dimension of the normal space $N(\boldsymbol{x}^*)$ is m. As in the case of the tangent space, it is often convenient to picture the normal space $N(\boldsymbol{x}^*)$ as passing through the point $\boldsymbol{x}^*$ (rather than through the origin of $\mathbb{R}^n$). For this, we define the *normal plane* at $\boldsymbol{x}^*$ as the set

$$NP(\boldsymbol{x}^*) = N(\boldsymbol{x}^*) + \boldsymbol{x}^* = \{\boldsymbol{x} + \boldsymbol{x}^* \in \mathbb{R}^n : \boldsymbol{x} \in N(\boldsymbol{x}^*)\}.$$

Figure 20.9 illustrates the normal space and plane in $\mathbb{R}^3$ (i.e., $n = 3$ and $m = 1$).

We now show that the tangent space and normal space are orthogonal complements of each other (see Section 3.3).

Lemma 20.1 *We have* $T(\boldsymbol{x}^*) = N(\boldsymbol{x}^*)^\perp$ *and* $T(\boldsymbol{x}^*)^\perp = N(\boldsymbol{x}^*)$. □

Proof. By definition of $T(\boldsymbol{x}^*)$, we may write

$$T(\boldsymbol{x}^*) = \{\boldsymbol{y} \in \mathbb{R}^n : \boldsymbol{x}^\top \boldsymbol{y} = 0 \text{ for all } \boldsymbol{x} \in N(\boldsymbol{x}^*)\}.$$

Hence, by definition of $N(\boldsymbol{x}^*)$, we have $T(\boldsymbol{x}^*) = N(\boldsymbol{x}^*)^\perp$. By Exercise 3.11 we also have $T(\boldsymbol{x}^*)^\perp = N(\boldsymbol{x}^*)$. ■

By Lemma 20.1, we can write $\mathbb{R}^n$ as the direct sum decomposition (see Section 3.3):

$$\mathbb{R}^n = N(\boldsymbol{x}^*) \oplus T(\boldsymbol{x}^*);$$

that is, given any vector $\boldsymbol{v} \in \mathbb{R}^n$, there are unique vectors $\boldsymbol{w} \in N(\boldsymbol{x}^*)$ and $\boldsymbol{y} \in T(\boldsymbol{x}^*)$ such that

$$\boldsymbol{v} = \boldsymbol{w} + \boldsymbol{y}.$$

20.4 Lagrange Condition

In this section we present a first-order necessary condition for extremum problems with constraints. The result is the well-known *Lagrange's theorem.* To better understand the idea underlying this theorem, we first consider functions of two variables and only one equality constraint. Let $h : \mathbb{R}^2 \to \mathbb{R}$ be the constraint function. Recall that at each point $\boldsymbol{x}$ of the domain, the gradient vector $\nabla h(\boldsymbol{x})$ is orthogonal to the level set that passes through that point. Indeed, let us choose a point $\boldsymbol{x}^* = [x_1^*, x_2^*]^\top$ such that $h(\boldsymbol{x}^*) = 0$, and assume that $\nabla h(\boldsymbol{x}^*) \neq \boldsymbol{0}$. The level set through the point $\boldsymbol{x}^*$ is the set $\{\boldsymbol{x} : h(\boldsymbol{x}) = 0\}$. We then parameterize this level set in a neighborhood of $\boldsymbol{x}^*$ by a curve $\{\boldsymbol{x}(t)\}$, that is, a continuously differentiable vector function $\boldsymbol{x} : \mathbb{R} \to \mathbb{R}^2$ such that

$$\boldsymbol{x}(t) = \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}, \quad t \in (a,b), \quad \boldsymbol{x}^* = \boldsymbol{x}(t^*), \quad \dot{\boldsymbol{x}}(t^*) \neq \boldsymbol{0}, \quad t^* \in (a,b).$$

We can now show that $\nabla h(\boldsymbol{x}^*)$ is orthogonal to $\dot{\boldsymbol{x}}(t^*)$. Indeed, because h is constant on the curve $\{\boldsymbol{x}(t) : t \in (a,b)\}$, we have that for all $t \in (a,b)$,

$$h(\boldsymbol{x}(t)) = 0.$$

Hence, for all $t \in (a,b)$,

$$\frac{d}{dt} h(\boldsymbol{x}(t)) = 0.$$

Applying the chain rule, we get

$$\frac{d}{dt} h(\boldsymbol{x}(t)) = \nabla h(\boldsymbol{x}(t))^\top \dot{\boldsymbol{x}}(t) = 0.$$

Therefore, $\nabla h(\boldsymbol{x}^*)$ is orthogonal to $\dot{\boldsymbol{x}}(t^*)$.

Now suppose that $\boldsymbol{x}^*$ is a minimizer of $f : \mathbb{R}^2 \to \mathbb{R}$ on the set $\{\boldsymbol{x} : h(\boldsymbol{x}) = 0\}$. We claim that $\nabla f(\boldsymbol{x}^*)$ is orthogonal to $\dot{\boldsymbol{x}}(t^*)$. To see this, it is enough to observe that the composite function of t given by

$$\phi(t) = f(\boldsymbol{x}(t))$$

achieves a minimum at t^*. Consequently, the first-order necessary condition for the unconstrained extremum problem implies that

$$\frac{d\phi}{dt}(t^*) = 0.$$

Applying the chain rule yields

$$0 = \frac{d}{dt}\phi(t^*) = \nabla f(\boldsymbol{x}(t^*))^\top \dot{\boldsymbol{x}}(t^*) = \nabla f(\boldsymbol{x}^*)^\top \dot{\boldsymbol{x}}(t^*).$$

Thus, $\nabla f(\boldsymbol{x}^*)$ is orthogonal to $\dot{\boldsymbol{x}}(t^*)$. The fact that $\dot{\boldsymbol{x}}(t^*)$ is tangent to the curve $\{\boldsymbol{x}(t)\}$ at $\boldsymbol{x}^*$ means that $\nabla f(\boldsymbol{x}^*)$ is orthogonal to the curve at $\boldsymbol{x}^*$ (see Figure 20.10).

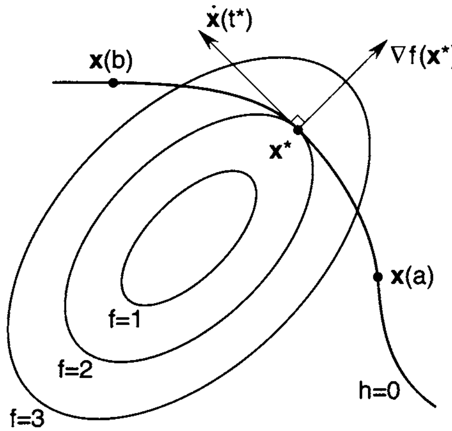

Figure 20.10 The gradient $\nabla f(\boldsymbol{x}^*)$ is orthogonal to the curve $\{\boldsymbol{x}(t)\}$ at the point $\boldsymbol{x}^*$ that is a minimizer of f on the curve.

Recall that $\nabla h(\boldsymbol{x}^*)$ is also orthogonal to $\dot{\boldsymbol{x}}(t^*)$. Therefore, the vectors $\nabla h(\boldsymbol{x}^*)$ and $\nabla f(\boldsymbol{x}^*)$ are parallel; that is, $\nabla f(\boldsymbol{x}^*)$ is a scalar multiple of $\nabla h(\boldsymbol{x}^*)$. The observations above allow us now to formulate *Lagrange's theorem* for functions of two variables with one constraint.

Theorem 20.2 ***Lagrange's Theorem for*** $n = 2$, $m = 1$. *Let the point* $\boldsymbol{x}^*$ *be a minimizer of* $f : \mathbb{R}^2 \to \mathbb{R}$ *subject to the constraint* $h(\boldsymbol{x}) = 0$, $h : \mathbb{R}^2 \to \mathbb{R}$. *Then,* $\nabla f(\boldsymbol{x}^*)$ *and* $\nabla h(\boldsymbol{x}^*)$ *are parallel. That is, if* $\nabla h(\boldsymbol{x}^*) \neq \mathbf{0}$, *then there exists a scalar* λ^* *such that*

$$\nabla f(\boldsymbol{x}^*) + \lambda^* \nabla h(\boldsymbol{x}^*) = \mathbf{0}.$$

□

In Theorem 20.2, we refer to λ^* as the *Lagrange multiplier.* Note that the theorem also holds for maximizers. Figure 20.11 gives an illustration of Lagrange's theorem for the case where $\boldsymbol{x}^*$ is a maximizer of f over the set $\{\boldsymbol{x} : h(\boldsymbol{x}) = 0\}$.

Lagrange's theorem provides a first-order necessary condition for a point to be a local minimizer. This condition, which we call the *Lagrange condition,* consists of two equations:

$$\begin{aligned} \nabla f(\boldsymbol{x}^*) + \lambda^* \nabla h(\boldsymbol{x}^*) &= \mathbf{0} \\ h(\boldsymbol{x}^*) &= 0. \end{aligned}$$

Note that the Lagrange condition is necessary but not sufficient. In Figure 20.12 we illustrate a variety of points where the Lagrange condition is

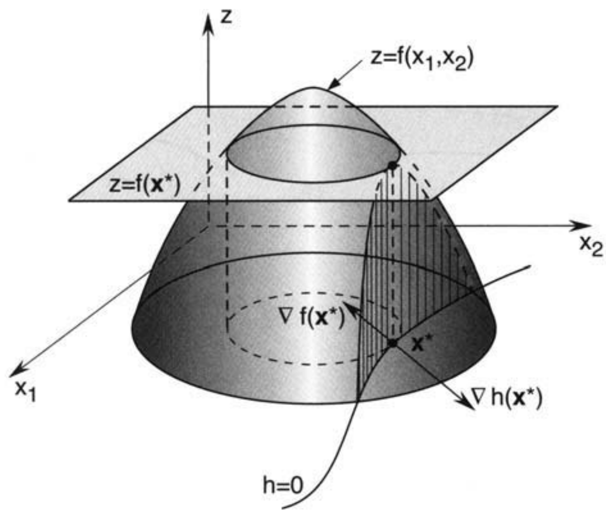

Figure 20.11 Lagrange's theorem for $n = 2$, $m = 1$.

satisfied, including a case where the point is not an extremizer (neither a maximizer nor a minimizer).

We now generalize Lagrange's theorem for the case when $f : \mathbb{R}^n \to \mathbb{R}$ and $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $m \le n$.

Theorem 20.3 ***Lagrange's Theorem.*** *Let $\boldsymbol{x}^*$ be a local minimizer (or maximizer) of $f : \mathbb{R}^n \to \mathbb{R}$, subject to $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $m \le n$. Assume that $\boldsymbol{x}^*$ is a regular point. Then, there exists $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that*

$$Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top.$$

□

Proof. We need to prove that

$$\nabla f(\boldsymbol{x}^*) = -D\boldsymbol{h}(\boldsymbol{x}^*)^\top \boldsymbol{\lambda}^*$$

for some $\boldsymbol{\lambda}^* \in \mathbb{R}^m$; that is, $\nabla f(\boldsymbol{x}^*) \in \mathcal{R}(D\boldsymbol{h}(\boldsymbol{x}^*)^\top) = N(\boldsymbol{x}^*)$. But by Lemma 20.1, $N(\boldsymbol{x}^*) = T(\boldsymbol{x}^*)^\perp$. Therefore, it remains to show that $\nabla f(\boldsymbol{x}^*) \in T(\boldsymbol{x}^*)^\perp$.

We proceed as follows. Suppose that

$$\boldsymbol{y} \in T(\boldsymbol{x}^*).$$

Then, by Theorem 20.1, there exists a differentiable curve $\{\boldsymbol{x}(t) : t \in (a, b)\}$ such that for all $t \in (a, b)$,

$$\boldsymbol{h}(\boldsymbol{x}(t)) = \boldsymbol{0},$$

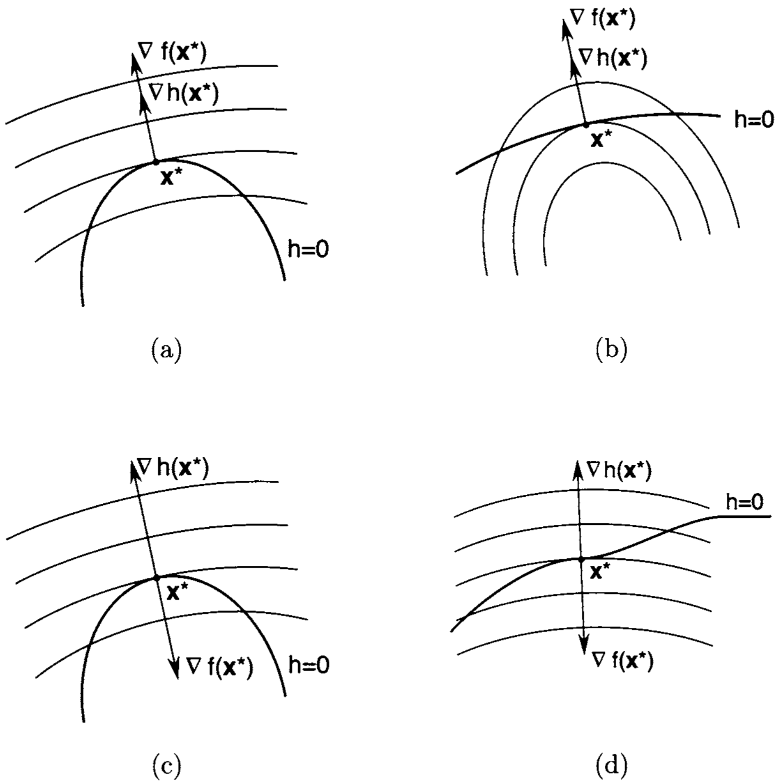

Figure 20.12 Four examples where the Lagrange condition is satisfied: (a) maximizer, (b) minimizer, (c) minimizer, (d) not an extremizer. (Adapted from [120].)

and there exists $t^* \in (a, b)$ satisfying

$$\boldsymbol{x}(t^*) = \boldsymbol{x}^*, \quad \dot{\boldsymbol{x}}(t^*) = \boldsymbol{y}.$$

Now consider the composite function $\phi(t) = f(\boldsymbol{x}(t))$. Note that t^* is a local minimizer of this function. By the first-order necessary condition for unconstrained local minimizers (see Theorem 6.1),

$$\frac{d\phi}{dt}(t^*) = 0.$$

Applying the chain rule yields

$$\frac{d\phi}{dt}(t^*) = Df(\boldsymbol{x}^*)\dot{\boldsymbol{x}}(t^*) = Df(\boldsymbol{x}^*)\boldsymbol{y} = \nabla f(\boldsymbol{x}^*)^\top \boldsymbol{y} = 0.$$

So all $\boldsymbol{y} \in T(\boldsymbol{x}^*)$ satisfy

$$\nabla f(\boldsymbol{x}^*)^\top \boldsymbol{y} = 0;$$

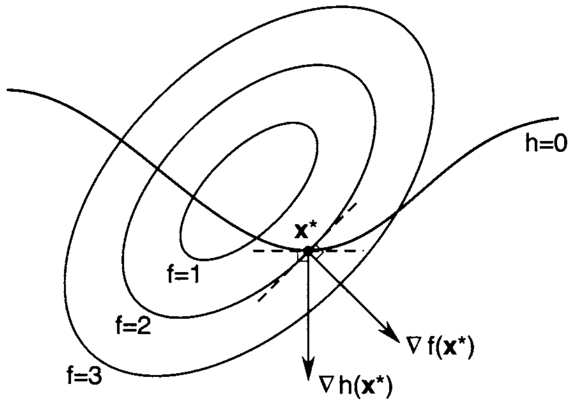

Figure 20.13 Example where the Lagrange condition does not hold.

that is,

$$\nabla f(\boldsymbol{x}^*) \in T(\boldsymbol{x}^*)^{\perp}.$$

This completes the proof. ∎

Lagrange's theorem states that if $\boldsymbol{x}^*$ is an extremizer, then the gradient of the objective function f can be expressed as a linear combination of the gradients of the constraints. We refer to the vector $\boldsymbol{\lambda}^*$ in Theorem 20.3 as the *Lagrange multiplier vector*, and its components as *Lagrange multipliers*.

From the proof of Lagrange's theorem, we see that a compact way to write the necessary condition is $\nabla f(\boldsymbol{x}^*) \in N(\boldsymbol{x}^*)$. If this condition fails, then $\boldsymbol{x}^*$ cannot be an extremizer. This situation is illustrated in Figure 20.13.

Notice that regularity is stated as an assumption in Lagrange's theorem. This assumption plays an essential role, as illustrated in the following example.

Example 20.5 Consider the following problem:

$$\begin{array}{ll} \text{minimize} & f(x) \\ \text{subject to} & h(x) = 0, \end{array}$$

where $f(x) = x$ and

$$h(x) = \begin{cases} x^2 & \text{if } x < 0 \\ 0 & \text{if } 0 \leq x \leq 1 \\ (x-1)^2 & \text{if } x > 1. \end{cases}$$

The feasible set is evidently $[0, 1]$. Clearly, $x^* = 0$ is a local minimizer. However, $f'(x^*) = 1$ and $h'(x^*) = 0$. Therefore, x^* does not satisfy the necessary condition in Lagrange's theorem. Note, however, that x^* is not a regular point, which is why Lagrange's theorem does not apply here. ∎

It is convenient to introduce the *Lagrangian function* $l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, given by

$$l(\boldsymbol{x}, \boldsymbol{\lambda}) \triangleq f(\boldsymbol{x}) + \boldsymbol{\lambda}^\top \boldsymbol{h}(\boldsymbol{x}).$$

The Lagrange condition for a local minimizer $\boldsymbol{x}^*$ can be represented using the Lagrangian function as

$$Dl(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) = \mathbf{0}^\top$$

for some $\boldsymbol{\lambda}^*$, where the derivative operation D is with respect to the entire argument $[\boldsymbol{x}^\top, \boldsymbol{\lambda}^\top]^\top$. In other words, the necessary condition in Lagrange's theorem is equivalent to the first-order necessary condition for unconstrained optimization applied to the Lagrangian function.

To see the above, denote the derivative of l with respect to $\boldsymbol{x}$ as $D_x l$ and the derivative of l with respect to $\boldsymbol{\lambda}$ as $D_\lambda l$. Then,

$$Dl(\boldsymbol{x}, \boldsymbol{\lambda}) = [D_x l(\boldsymbol{x}, \boldsymbol{\lambda}), D_\lambda l(\boldsymbol{x}, \boldsymbol{\lambda})].$$

Note that $D_x l(\boldsymbol{x}, \boldsymbol{\lambda}) = Df(\boldsymbol{x}) + \boldsymbol{\lambda}^\top D\boldsymbol{h}(\boldsymbol{x})$ and $D_\lambda l(\boldsymbol{x}, \boldsymbol{\lambda}) = \boldsymbol{h}(\boldsymbol{x})^\top$. Therefore, Lagrange's theorem for a local minimizer $\boldsymbol{x}^*$ can be stated as

$$\begin{aligned} D_x l(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) &= \mathbf{0}^\top, \\ D_\lambda l(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) &= \mathbf{0}^\top \end{aligned}$$

for some $\boldsymbol{\lambda}^*$, which is equivalent to

$$Dl(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) = \mathbf{0}^\top.$$

In other words, the Lagrange condition can be expressed as $Dl(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) = \mathbf{0}^\top$.

The Lagrange condition is used to find possible extremizers. This entails solving the equations

$$\begin{aligned} D_x l(\boldsymbol{x}, \boldsymbol{\lambda}) &= \mathbf{0}^\top, \\ D_\lambda l(\boldsymbol{x}, \boldsymbol{\lambda}) &= \mathbf{0}^\top. \end{aligned}$$

The above represents $n + m$ equations in $n + m$ unknowns. Keep in mind that the Lagrange condition is necessary but not sufficient; that is, a point $\boldsymbol{x}^*$ satisfying the equations above need not be an extremizer.

Example 20.6 Given a fixed area of cardboard, we wish to construct a closed cardboard box with maximum volume. We can formulate and solve this problem using the Lagrange condition. Denote the dimensions of the box with maximum volume by x_1, x_2, and x_3, and let the given fixed area of cardboard be A. The problem can then be formulated as

$$\begin{aligned} \text{maximize} \quad & x_1 x_2 x_3 \\ \text{subject to} \quad & x_1 x_2 + x_2 x_3 + x_3 x_1 = \frac{A}{2}. \end{aligned}$$

We denote $f(\boldsymbol{x}) = -x_1x_2x_3$ and $h(\boldsymbol{x}) = x_1x_2+x_2x_3+x_3x_1-A/2$. We have $\nabla f(\boldsymbol{x}) = -[x_2x_3, x_1x_3, x_1x_2]^\top$ and $\nabla h(\boldsymbol{x}) = [x_2+x_3, x_1+x_3, x_1+x_2]^\top$. Note that all feasible points are regular in this case. By the Lagrange condition, the dimensions of the box with maximum volume satisfies

$$\begin{aligned} x_2x_3 - \lambda(x_2+x_3) &= 0 \\ x_1x_3 - \lambda(x_1+x_3) &= 0 \\ x_1x_2 - \lambda(x_1+x_2) &= 0 \\ x_1x_2 + x_2x_3 + x_3x_1 &= \frac{A}{2}, \end{aligned}$$

where $\lambda \in \mathbb{R}$.

We now solve these equations. First, we show that that x_1, x_2, x_3, and λ are all nonzero. Suppose that $x_1 = 0$. By the constraints, we have $x_2x_3 = A/2$. However, the second and third equations in the Lagrange condition yield $\lambda x_2 = \lambda x_3 = 0$, which together with the first equation implies that $x_2x_3 = 0$. This contradicts the constraints. A similar argument applies to x_2 and x_3.

Next, suppose that $\lambda = 0$. Then, the sum of the three Lagrange equations gives $x_2x_3 + x_1x_3 + x_1x_2 = 0$, which contradicts the constraints.

We now solve for x_1, x_2, and x_3 in the Lagrange equations. First, multiply the first equation by x_1 and the second by x_2, and subtract one from the other. We arrive at $x_3\lambda(x_1 - x_2) = 0$. Because neither x_3 nor λ can be zero (by part b), we conclude that $x_1 = x_2$. We similarly deduce that $x_2 = x_3$. From the constraint equation, we obtain $x_1 = x_2 = x_3 = \sqrt{A/6}$.

Notice that we have ignored the constraints that x_1, x_2, and x_3 are positive so that we can solve the problem using Lagrange's theorem. However, there is only one solution to the Lagrange equations, and the solution is positive. Therefore, if a solution exists for the problem with positivity constraints on the variables x_1, x_2, and x_3, then this solution must necessarily be equal to the solution above obtained by ignoring the positivity constraints. ∎

Next we provide an example with a quadratic objective function and a quadratic constraint.

Example 20.7 Consider the problem of extremizing the objective function

$$f(\boldsymbol{x}) = x_1^2 + x_2^2$$

on the ellipse

$$\{[x_1, x_2]^\top : h(\boldsymbol{x}) = x_1^2 + 2x_2^2 - 1 = 0\}.$$

We have

$$\begin{aligned} \nabla f(\boldsymbol{x}) &= [2x_1, 2x_2]^\top, \\ \nabla h(\boldsymbol{x}) &= [2x_1, 4x_2]^\top. \end{aligned}$$

Thus,

$$D_x l(\boldsymbol{x},\lambda) = D_x[f(\boldsymbol{x}) + \lambda h(\boldsymbol{x})] = [2x_1 + 2\lambda x_1, 2x_2 + 4\lambda x_2]$$

and

$$D_\lambda l(\boldsymbol{x},\lambda) = h(\boldsymbol{x}) = x_1^2 + 2x_2^2 - 1.$$

Setting $D_x l(\boldsymbol{x},\lambda) = \mathbf{0}^\top$ and $D_\lambda l(\boldsymbol{x},\lambda) = 0$, we obtain three equations in three unknowns

$$\begin{aligned} 2x_1 + 2\lambda x_1 &= 0, \\ 2x_2 + 4\lambda x_2 &= 0, \\ x_1^2 + 2x_2^2 &= 1. \end{aligned}$$

All feasible points in this problem are regular. From the first of the equations above, we get either $x_1 = 0$ or $\lambda = -1$. For the case where $x_1 = 0$, the second and third equations imply that $\lambda = -1/2$ and $x_2 = \pm 1/\sqrt{2}$. For the case where $\lambda = -1$, the second and third equations imply that $x_1 = \pm 1$ and $x_2 = 0$. Thus, the points that satisfy the Lagrange condition for extrema are

$$\boldsymbol{x}^{(1)} = \begin{bmatrix} 0 \\ 1/\sqrt{2} \end{bmatrix}, \quad \boldsymbol{x}^{(2)} = \begin{bmatrix} 0 \\ -1/\sqrt{2} \end{bmatrix}, \quad \boldsymbol{x}^{(3)} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \boldsymbol{x}^{(4)} = \begin{bmatrix} -1 \\ 0 \end{bmatrix}.$$

Because

$$f(\boldsymbol{x}^{(1)}) = f(\boldsymbol{x}^{(2)}) = \frac{1}{2}$$

and

$$f(\boldsymbol{x}^{(3)}) = f(\boldsymbol{x}^{(4)}) = 1$$

we conclude that if there are minimizers, then they are located at $\boldsymbol{x}^{(1)}$ and $\boldsymbol{x}^{(2)}$, and if there are maximizers, then they are located at $\boldsymbol{x}^{(3)}$ and $\boldsymbol{x}^{(4)}$. It turns out that, indeed, $\boldsymbol{x}^{(1)}$ and $\boldsymbol{x}^{(2)}$ are minimizers and $\boldsymbol{x}^{(3)}$ and $\boldsymbol{x}^{(4)}$ are maximizers. This problem can be solved graphically, as illustrated in Figure 20.14. ∎

In the example above, both the objective function f and the constraint function h are quadratic functions. In the next example we take a closer look at a class of problems where both the objective function f and the constraint h are quadratic functions of n variables.

Example 20.8 Consider the following problem:

$$\text{maximize} \quad \frac{\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}},$$

where $\boldsymbol{Q} = \boldsymbol{Q}^\top \geq 0$ and $\boldsymbol{P} = \boldsymbol{P}^\top > 0$. Note that if a point $\boldsymbol{x} = [x_1, \ldots, x_n]^\top$ is a solution to the problem, then so is any nonzero scalar multiple of it,

$$t\boldsymbol{x} = [tx_1, \ldots, tx_n]^\top, \quad t \neq 0.$$

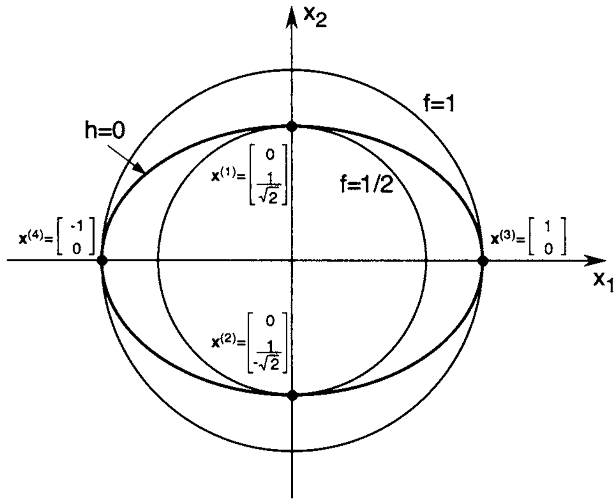

Figure 20.14 Graphical solution of the problem in Example 20.7.

Indeed,

$$\frac{(t\boldsymbol{x})^\top \boldsymbol{Q}(t\boldsymbol{x})}{(t\boldsymbol{x})^\top \boldsymbol{P}(t\boldsymbol{x})} = \frac{t^2\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}}{t^2\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}} = \frac{\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}}.$$

Therefore, to avoid the multiplicity of solutions, we further impose the constraint

$$\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 1.$$

The optimization problem becomes

$$\begin{array}{ll} \text{maximize} & \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 1. \end{array}$$

Let us write

$$\begin{aligned} f(\boldsymbol{x}) &= \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}, \\ h(\boldsymbol{x}) &= 1 - \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}. \end{aligned}$$

Any feasible point for this problem is regular (see Exercise 20.13). We now apply Lagrange's method. We first form the Lagrangian function

$$l(\boldsymbol{x}, \lambda) = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + \lambda(1 - \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}).$$

Applying the Lagrange condition yields

$$\begin{aligned} D_x l(\boldsymbol{x}, \lambda) &= 2\boldsymbol{x}^\top \boldsymbol{Q} - 2\lambda \boldsymbol{x}^\top \boldsymbol{P} = \boldsymbol{0}^\top, \\ D_\lambda l(\boldsymbol{x}, \lambda) &= 1 - \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 0. \end{aligned}$$

The first of the equations above can be represented as

$$\boldsymbol{Q}\boldsymbol{x} - \lambda \boldsymbol{P}\boldsymbol{x} = \mathbf{0}$$

or

$$(\lambda \boldsymbol{P} - \boldsymbol{Q})\boldsymbol{x} = \mathbf{0}.$$

This representation is possible because $\boldsymbol{P} = \boldsymbol{P}^\top$ and $\boldsymbol{Q} = \boldsymbol{Q}^\top$. By assumption $\boldsymbol{P} > 0$, hence $\boldsymbol{P}^{-1}$ exists. Premultiplying $(\lambda \boldsymbol{P} - \boldsymbol{Q})\boldsymbol{x} = \mathbf{0}$ by $\boldsymbol{P}^{-1}$, we obtain

$$(\lambda \boldsymbol{I}_n - \boldsymbol{P}^{-1}\boldsymbol{Q})\boldsymbol{x} = \mathbf{0}$$

or, equivalently,

$$\boldsymbol{P}^{-1}\boldsymbol{Q}\boldsymbol{x} = \lambda \boldsymbol{x}.$$

Therefore, the solution, if it exists, is an eigenvector of $\boldsymbol{P}^{-1}\boldsymbol{Q}$, and the Lagrange multiplier is the corresponding eigenvalue. As usual, let $\boldsymbol{x}^*$ and λ^* be the optimal solution. Because $\boldsymbol{x}^{*\top}\boldsymbol{P}\boldsymbol{x}^* = 1$ and $\boldsymbol{P}^{-1}\boldsymbol{Q}\boldsymbol{x}^* = \lambda^*\boldsymbol{x}^*$, we have

$$\lambda^* = \boldsymbol{x}^{*\top}\boldsymbol{Q}\boldsymbol{x}^*.$$

Hence, λ^* is the maximum of the objective function, and therefore is, in fact, the maximal eigenvalue of $\boldsymbol{P}^{-1}\boldsymbol{Q}$. It is also called the maximal *generalized eigenvalue.* ∎

In the problems above, we are able to find points that are candidates for extremizers of the given objective function subject to equality constraints. These critical points are the only candidates because they are the only points that satisfy the Lagrange condition. To classify such critical points as minimizers, maximizers, or neither, we need a stronger condition—possibly a necessary and sufficient condition. In the next section we discuss a second-order necessary condition and a second-order sufficient condition for minimizers.

20.5 Second-Order Conditions

We assume that $f : \mathbb{R}^n \to \mathbb{R}$ and $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$ are twice continuously differentiable: $f, \boldsymbol{h} \in \mathcal{C}^2$. Let

$$l(\boldsymbol{x}, \boldsymbol{\lambda}) = f(\boldsymbol{x}) + \boldsymbol{\lambda}^\top \boldsymbol{h}(\boldsymbol{x}) = f(\boldsymbol{x}) + \lambda_1 h_1(\boldsymbol{x}) + \cdots + \lambda_m h_m(\boldsymbol{x})$$

be the Lagrangian function. Let $\boldsymbol{L}(\boldsymbol{x}, \boldsymbol{\lambda})$ be the Hessian matrix of $l(\boldsymbol{x}, \boldsymbol{\lambda})$ with respect to $\boldsymbol{x}$:

$$\boldsymbol{L}(\boldsymbol{x}, \boldsymbol{\lambda}) = \boldsymbol{F}(\boldsymbol{x}) + \lambda_1 \boldsymbol{H}_1(\boldsymbol{x}) + \cdots + \lambda_m \boldsymbol{H}_m(\boldsymbol{x}),$$

where $\boldsymbol{F}(\boldsymbol{x})$ is the Hessian matrix of f at $\boldsymbol{x}$ and $\boldsymbol{H}_k(\boldsymbol{x})$ is the Hessian matrix of h_k at $\boldsymbol{x}$, $k = 1, \ldots, m$, given by

$$\boldsymbol{H}_k(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial^2 h_k}{\partial x_1^2}(\boldsymbol{x}) & \cdots & \frac{\partial^2 h_k}{\partial x_n \partial x_1}(\boldsymbol{x}) \\ \vdots & & \vdots \\ \frac{\partial^2 h_k}{\partial x_1 \partial x_n}(\boldsymbol{x}) & \cdots & \frac{\partial^2 h_k}{\partial^2 x_n}(\boldsymbol{x}) \end{bmatrix}.$$

We introduce the notation $[\boldsymbol{\lambda H}(\boldsymbol{x})]$:

$$[\boldsymbol{\lambda H}(\boldsymbol{x})] = \lambda_1 \boldsymbol{H}_1(\boldsymbol{x}) + \cdots + \lambda_m \boldsymbol{H}_m(\boldsymbol{x}).$$

Using the notation above, we can write

$$\boldsymbol{L}(\boldsymbol{x}, \boldsymbol{\lambda}) = \boldsymbol{F}(\boldsymbol{x}) + [\boldsymbol{\lambda H}(\boldsymbol{x})].$$

Theorem 20.4 ***Second-Order Necessary Conditions.*** *Let $\boldsymbol{x}^*$ be a local minimizer of $f : \mathbb{R}^n \to \mathbb{R}$ subject to $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $m \le n$, and $f, \boldsymbol{h} \in \mathcal{C}^2$. Suppose that $\boldsymbol{x}^*$ is regular. Then, there exists $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that:*

1. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$.
2. *For all $\boldsymbol{y} \in T(\boldsymbol{x}^*)$, we have $\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)\boldsymbol{y} \ge 0$.* □

Proof. The existence of $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$ follows from Lagrange's theorem. It remains to prove the second part of the result. Suppose that $\boldsymbol{y} \in T(\boldsymbol{x}^*)$; that is, $\boldsymbol{y}$ belongs to the tangent space to $S = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\}$ at $\boldsymbol{x}^*$. Because $\boldsymbol{h} \in \mathcal{C}^2$, following the argument of Theorem 20.1, there exists a twice-differentiable curve $\{\boldsymbol{x}(t) : t \in (a, b)\}$ on S such that

$$\boldsymbol{x}(t^*) = \boldsymbol{x}^*, \quad \dot{\boldsymbol{x}}(t^*) = \boldsymbol{y}$$

for some $t^* \in (a, b)$. Observe that by assumption, t^* is a local minimizer of the function $\phi(t) = f(\boldsymbol{x}(t))$. From the second-order necessary condition for unconstrained minimization (see Theorem 6.2), we obtain

$$\frac{d^2\phi}{dt^2}(t^*) \ge 0.$$

Using the formula

$$\frac{d}{dt}(\boldsymbol{y}(t)^\top \boldsymbol{z}(t)) = \boldsymbol{z}(t)^\top \frac{d\boldsymbol{y}}{dt}(t) + \boldsymbol{y}(t)^\top \frac{d\boldsymbol{z}}{dt}(t)$$

and applying the chain rule yields

$$\begin{aligned}
\frac{d^2\phi}{dt^2}(t^*) &= \frac{d}{dt}[Df(\boldsymbol{x}(t^*))\dot{\boldsymbol{x}}(t^*)] \\
&= \dot{\boldsymbol{x}}(t^*)^\top \boldsymbol{F}(\boldsymbol{x}^*)\dot{\boldsymbol{x}}(t^*) + Df(\boldsymbol{x}^*)\ddot{\boldsymbol{x}}(t^*) \\
&= \boldsymbol{y}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{y} + Df(\boldsymbol{x}^*)\ddot{\boldsymbol{x}}(t^*) \ge 0.
\end{aligned}$$

Because $\boldsymbol{h}(\boldsymbol{x}(t)) = \boldsymbol{0}$ for all $t \in (a, b)$, we have

$$\frac{d^2}{dt^2}\boldsymbol{\lambda}^{*\top}\boldsymbol{h}(\boldsymbol{x}(t)) = 0.$$

Thus, for all $t \in (a, b)$,

$$\begin{aligned}
\frac{d^2}{dt^2}\boldsymbol{\lambda}^{*\top}\boldsymbol{h}(\boldsymbol{x}(t)) &= \frac{d}{dt}\left[\boldsymbol{\lambda}^{*\top}\frac{d}{dt}\boldsymbol{h}(\boldsymbol{x}(t))\right] \\
&= \frac{d}{dt}\left[\sum_{k=1}^{m}\lambda_k^*\frac{d}{dt}h_k(\boldsymbol{x}(t))\right] \\
&= \frac{d}{dt}\left[\sum_{k=1}^{m}\lambda_k^* Dh_k(\boldsymbol{x}(t))\dot{\boldsymbol{x}}(t)\right] \\
&= \sum_{k=1}^{m}\lambda_k^*\frac{d}{dt}(Dh_k(\boldsymbol{x}(t))\dot{\boldsymbol{x}}(t)) \\
&= \sum_{k=1}^{m}\lambda_k^*\left[\dot{\boldsymbol{x}}(t)^\top \boldsymbol{H}_k(\boldsymbol{x}(t))\dot{\boldsymbol{x}}(t) + Dh_k(\boldsymbol{x}(t))\ddot{\boldsymbol{x}}(t)\right] \\
&= \dot{\boldsymbol{x}}^\top(t)[\boldsymbol{\lambda}^*\boldsymbol{H}(\boldsymbol{x}(t))]\dot{\boldsymbol{x}}(t) + \boldsymbol{\lambda}^{*\top}D\boldsymbol{h}(\boldsymbol{x}(t))\ddot{\boldsymbol{x}}(t) \\
&= 0.
\end{aligned}$$

In particular, the above is true for $t = t^*$; that is,

$$\boldsymbol{y}^\top[\boldsymbol{\lambda}^*\boldsymbol{H}(\boldsymbol{x}^*)]\boldsymbol{y} + \boldsymbol{\lambda}^{*\top}D\boldsymbol{h}(\boldsymbol{x}^*)\ddot{\boldsymbol{x}}(t^*) = 0.$$

Adding this equation to the inequality

$$\boldsymbol{y}^\top \boldsymbol{F}(\boldsymbol{x}^*)\boldsymbol{y} + Df(\boldsymbol{x}^*)\ddot{\boldsymbol{x}}(t^*) \geq 0$$

yields

$$\boldsymbol{y}^\top\left(\boldsymbol{F}(\boldsymbol{x}^*) + [\boldsymbol{\lambda}^*\boldsymbol{H}(\boldsymbol{x}^*)]\right)\boldsymbol{y} + (Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top}D\boldsymbol{h}(\boldsymbol{x}^*))\ddot{\boldsymbol{x}}(t^*) \geq 0.$$

But, by Lagrange's theorem, $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top}D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$. Therefore,

$$\boldsymbol{y}^\top\left(\boldsymbol{F}(\boldsymbol{x}^*) + [\boldsymbol{\lambda}^*\boldsymbol{H}(\boldsymbol{x}^*)]\right)\boldsymbol{y} = \boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)\boldsymbol{y} \geq 0,$$

which proves the result. ∎

Observe that $\boldsymbol{L}(\boldsymbol{x}, \boldsymbol{\lambda})$ plays a similar role as the Hessian matrix $\boldsymbol{F}(\boldsymbol{x})$ of the objective function f did in the unconstrained minimization case. However, we now require that $\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) \geq 0$ only on $T(\boldsymbol{x}^*)$ rather than on $\mathbb{R}^n$.

The conditions above are necessary, but not sufficient, for a point to be a local minimizer. We now present, without a proof, sufficient conditions for a point to be a strict local minimizer.

Theorem 20.5 ***Second-Order Sufficient Conditions.*** *Suppose that $f, \boldsymbol{h} \in \mathcal{C}^2$ and there exists a point $\boldsymbol{x}^* \in \mathbb{R}^n$ and $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that:*

1. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top}D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$.

2. For all $\boldsymbol{y} \in T(\boldsymbol{x}^*)$, $\boldsymbol{y} \neq \mathbf{0}$, we have $\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)\boldsymbol{y} > 0$.

Then, $\boldsymbol{x}^*$ is a strict local minimizer of f subject to $\boldsymbol{h}(\boldsymbol{x}) = \mathbf{0}$. □

Proof. The interested reader can consult [88, p. 334] for a proof of this result. ■

Theorem 20.5 states that if an $\boldsymbol{x}^*$ satisfies the Lagrange condition, and $\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)$ is positive definite on $T(\boldsymbol{x}^*)$, then $\boldsymbol{x}^*$ is a strict local minimizer. A similar result to Theorem 20.5 holds for a strict local maximizer, the only difference being that $\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)$ be negative definite on $T(\boldsymbol{x}^*)$. We illustrate this condition in the following example.

Example 20.9 Consider the following problem:

$$\text{maximize} \quad \frac{\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}}{\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}},$$

where

$$\boldsymbol{Q} = \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \boldsymbol{P} = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}.$$

As pointed out earlier, we can represent this problem in the equivalent form

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 1. \end{aligned}$$

The Lagrangian function for the transformed problem is given by

$$l(\boldsymbol{x}, \lambda) = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + \lambda(1 - \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x}).$$

The Lagrange condition yields

$$(\lambda \boldsymbol{I} - \boldsymbol{P}^{-1}\boldsymbol{Q})\boldsymbol{x} = \mathbf{0},$$

where

$$\boldsymbol{P}^{-1}\boldsymbol{Q} = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}.$$

There are only two values of λ that satisfy $(\lambda \boldsymbol{I} - \boldsymbol{P}^{-1}\boldsymbol{Q})\boldsymbol{x} = \mathbf{0}$, namely, the eigenvalues of $\boldsymbol{P}^{-1}\boldsymbol{Q}$: $\lambda_1 = 2$, $\lambda_2 = 1$. We recall from our previous discussion of this problem that the Lagrange multiplier corresponding to the solution is the maximum eigenvalue of $\boldsymbol{P}^{-1}\boldsymbol{Q}$, namely, $\lambda^* = \lambda_1 = 2$. The corresponding eigenvector is the maximizer—the solution to the problem. The eigenvector corresponding to the eigenvalue $\lambda^* = 2$ satisfying the constraint $\boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 1$ is $\pm\boldsymbol{x}^*$, where

$$\boldsymbol{x}^* = \left[\frac{1}{\sqrt{2}}, 0\right]^\top.$$

At this point, all we have established is that the pairs $(\pm\boldsymbol{x}^*, \lambda^*)$ satisfy the Lagrange condition. We now show that the points $\pm\boldsymbol{x}^*$ are, in fact, strict local maximizers. We do this for the point $\boldsymbol{x}^*$. A similar procedure applies to $-\boldsymbol{x}^*$. We first compute the Hessian matrix of the Lagrangian function. We have

$$\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*) = 2\boldsymbol{Q} - 2\lambda\boldsymbol{P} = \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix}.$$

The tangent space $T(\boldsymbol{x}^*)$ to $\{\boldsymbol{x} : 1 - \boldsymbol{x}^\top \boldsymbol{P}\boldsymbol{x} = 0\}$ is

$$\begin{aligned} T(\boldsymbol{x}^*) &= \{\boldsymbol{y} \in \mathbb{R}^2 : \boldsymbol{x}^{*\top}\boldsymbol{P}\boldsymbol{y} = 0\} \\ &= \{\boldsymbol{y} : [\sqrt{2}, 0]\boldsymbol{y} = 0\} \\ &= \{\boldsymbol{y} : \boldsymbol{y} = [0, a]^\top,\ a \in \mathbb{R}\}. \end{aligned}$$

Note that for each $\boldsymbol{y} \in T(\boldsymbol{x}^*)$, $\boldsymbol{y} \neq \boldsymbol{0}$,

$$\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \lambda^*)\boldsymbol{y} = [0, a] \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix} \begin{bmatrix} 0 \\ a \end{bmatrix} = -2a^2 < 0.$$

Hence, $\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*) < 0$ on $T(\boldsymbol{x}^*)$, and thus $\boldsymbol{x}^* = [1/\sqrt{2}, 0]^\top$ is a strict local maximizer. The same is true for the point $-\boldsymbol{x}^*$. Note that

$$\frac{\boldsymbol{x}^{*\top}\boldsymbol{Q}\boldsymbol{x}^*}{\boldsymbol{x}^{*\top}\boldsymbol{P}\boldsymbol{x}^*} = 2,$$

which, as expected, is the value of the maximal eigenvalue of $\boldsymbol{P}^{-1}\boldsymbol{Q}$. Finally, we point out that any scalar multiple $t\boldsymbol{x}^*$ of $\boldsymbol{x}^*$, $t \neq 0$, is a solution to the original problem of maximizing $\boldsymbol{x}^\top\boldsymbol{Q}\boldsymbol{x}/\boldsymbol{x}^\top\boldsymbol{P}\boldsymbol{x}$. ∎

20.6 Minimizing Quadratics Subject to Linear Constraints

Consider the problem

$$\begin{aligned} &\text{minimize} && \frac{1}{2}\boldsymbol{x}^\top\boldsymbol{Q}\boldsymbol{x} \\ &\text{subject to} && \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}, \end{aligned}$$

where $\boldsymbol{Q} > 0$, $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $m < n$, $\operatorname{rank}\boldsymbol{A} = m$. This problem is a special case of what is called a *quadratic programming problem* (the general form of a quadratic programming problem includes the constraint $\boldsymbol{x} \geq \boldsymbol{0}$). Note that the constraint set contains an infinite number of points (see Section 2.3). We now show, using Lagrange's theorem, that there is a unique solution to the optimization problem above. Following that, we provide an example illustrating the application of this solution to an optimal control problem.

To solve the problem, we first form the Lagrangian function

$$l(\boldsymbol{x}, \boldsymbol{\lambda}) = \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + \boldsymbol{\lambda}^\top(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}).$$

The Lagrange condition yields

$$D_x l(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) = \boldsymbol{x}^{*\top}\boldsymbol{Q} - \boldsymbol{\lambda}^{*\top}\boldsymbol{A} = \boldsymbol{0}^\top.$$

Rewriting, we get

$$\boldsymbol{x}^* = \boldsymbol{Q}^{-1}\boldsymbol{A}^\top\boldsymbol{\lambda}^*.$$

Premultiplying both sides of the above by $\boldsymbol{A}$ gives

$$\boldsymbol{A}\boldsymbol{x}^* = \boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^\top\boldsymbol{\lambda}^*.$$

Using the fact that $\boldsymbol{A}\boldsymbol{x}^* = \boldsymbol{b}$, and noting that $\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^\top$ is invertible because $\boldsymbol{Q} > 0$ and $\operatorname{rank} \boldsymbol{A} = m$, we can solve for $\boldsymbol{\lambda}^*$ to obtain

$$\boldsymbol{\lambda}^* = (\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^\top)^{-1}\boldsymbol{b}.$$

Therefore, we obtain

$$\boldsymbol{x}^* = \boldsymbol{Q}^{-1}\boldsymbol{A}^\top(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^\top)^{-1}\boldsymbol{b}.$$

The point $\boldsymbol{x}^*$ is the only candidate for a minimizer. To establish that $\boldsymbol{x}^*$ is indeed a minimizer, we verify that $\boldsymbol{x}^*$ satisfies the second-order sufficient conditions. For this, we first find the Hessian matrix of the Lagrangian function at $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)$. We have

$$\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) = \boldsymbol{Q},$$

which is positive definite. Thus, the point $\boldsymbol{x}^*$ is a strict local minimizer. We will see in Chapter 22 that $\boldsymbol{x}^*$ is, in fact, a global minimizer.

The special case where $\boldsymbol{Q} = \boldsymbol{I}_n$, the $n \times n$ identity matrix, reduces to the problem considered in Section 12.3. Specifically, the problem in Section 12.3 is to minimize the norm $\|\boldsymbol{x}\|$ subject to $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$. The objective function here is $f(\boldsymbol{x}) = \|\boldsymbol{x}\|$, which is not differentiable at $\boldsymbol{x} = \boldsymbol{0}$. This precludes the use of Lagrange's theorem because the theorem requires differentiability of the objective function. We can overcome this difficulty by considering an equivalent optimization problem:

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|\boldsymbol{x}\|^2 \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}. \end{array}$$

The objective function $\|\boldsymbol{x}\|^2/2$ has the same minimizer as the previous objective function $\|\boldsymbol{x}\|$. Indeed, if $\boldsymbol{x}^*$ is such that for all $\boldsymbol{x} \in \mathbb{R}^n$ satisfying $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, $\|\boldsymbol{x}^*\| \leq \|\boldsymbol{x}\|$, then $\|\boldsymbol{x}^*\|^2/2 \leq \|\boldsymbol{x}\|^2/2$. The same is true for the converse. Because the problem of minimizing $\|\boldsymbol{x}\|^2/2$ subject to $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$ is simply the problem considered above with $\boldsymbol{Q} = \boldsymbol{I}_n$, we easily deduce the solution to be $\boldsymbol{x}^* = \boldsymbol{A}^\top(\boldsymbol{A}\boldsymbol{A}^\top)^{-1}\boldsymbol{b}$, which agrees with the solution in Section 12.3.

Example 20.10 Consider the discrete-time linear system model

$$x_k = ax_{k-1} + bu_k, \qquad k \geq 1,$$

with initial condition x_0 given. We can think of $\{x_k\}$ as a discrete-time signal that is controlled by an external input signal $\{u_k\}$. In the control literature, x_k is called the *state* at time k. For a given x_0, our goal is to choose the control signal $\{u_k\}$ so that the state remains "small" over a time interval $[1, N]$, but at the same time the control signal is "not too large." To express the desire to keep the state $\{x_k\}$ small, we choose the control sequence to minimize

$$\frac{1}{2}\sum_{i=1}^{N} x_i^2.$$

On the other hand, maintaining a control signal that is not too large, we minimize

$$\frac{1}{2}\sum_{i=1}^{N} u_i^2.$$

The two objectives above are conflicting in the sense that they cannot, in general, be achieved simultaneously—minimizing the first may result in a large control effort, while minimizing the second may result in large states. This is clearly a problem that requires compromise. One way to approach the problem is to minimize a weighted sum of the two functions above. Specifically, we can formulate the problem as

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2}\sum_{i=1}^{N}\left(qx_i^2 + ru_i^2\right) \\ \text{subject to} \quad & x_k = ax_{k-1} + bu_k,\ k = 1, \ldots, N,\ x_0 \text{ given,} \end{aligned}$$

where the parameters q and r reflect the relative importance of keeping the state small versus keeping the control effort not too large. This problem is an instance of the *linear quadratic regulator (LQR) problem* (see, e.g., [15], [20], [85], [86], or [99]). Combining the two conflicting objectives of keeping the state small while keeping the control effort small is an instance of the *weighted sum* approach (see Section 24.4).

To solve the problem above, we can rewrite it as a quadratic programming problem. Define

$$\boldsymbol{Q} = \begin{bmatrix} q\boldsymbol{I}_N & \boldsymbol{O} \\ \boldsymbol{O} & r\boldsymbol{I}_N \end{bmatrix},$$

$$\boldsymbol{A} = \begin{bmatrix} 1 & & \cdots & 0 & -b & & \cdots & 0 \\ -a & 1 & & \vdots & & -b & & \vdots \\ & \ddots & \ddots & & \vdots & & \ddots & \\ 0 & & -a & 1 & 0 & \cdots & & -b \end{bmatrix},$$

$$\boldsymbol{b} = \begin{bmatrix} ax_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \qquad \boldsymbol{z} = [x_1, \ldots, x_N, u_1, \ldots, u_N]^\top.$$

With these definitions, the problem reduces to the previously considered quadratic programming problem,

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\boldsymbol{z}^\top \boldsymbol{Q}\boldsymbol{z} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{z} = \boldsymbol{b}, \end{array}$$

where $\boldsymbol{Q}$ is $2N \times 2N$, $\boldsymbol{A}$ is $N \times 2N$, and $\boldsymbol{b} \in \mathbb{R}^N$. The solution is

$$\boldsymbol{z}^* = \boldsymbol{Q}^{-1}\boldsymbol{A}^\top(\boldsymbol{A}\boldsymbol{Q}^{-1}\boldsymbol{A}^\top)^{-1}\boldsymbol{b}.$$

The first N components of $\boldsymbol{z}^*$ represent the optimal state signal in the interval $[1, N]$, whereas the second N components represent the optimal control signal.

In practice, computation of the matrix inverses in the formula for $\boldsymbol{z}^*$ above may be too costly. There are other ways to tackle the problem by exploiting its special structure. This is the study of *optimal control* (see, e.g., [15], [20], [85], [86], or [99]). ∎

The following example illustrates an application of the above discussion.

Example 20.11 *Credit-Card Holder Dilemma.* Suppose that we currently have a credit-card debt of \$10,000. Credit-card debts are subject to a monthly interest rate of 2%, and the account balance is increased by the interest amount every month. Each month we have the option of reducing the account balance by contributing a payment to the account. Over the next 10 months, we plan to contribute a payment every month in such a way as to minimize the overall debt level while minimizing the hardship of making monthly payments.

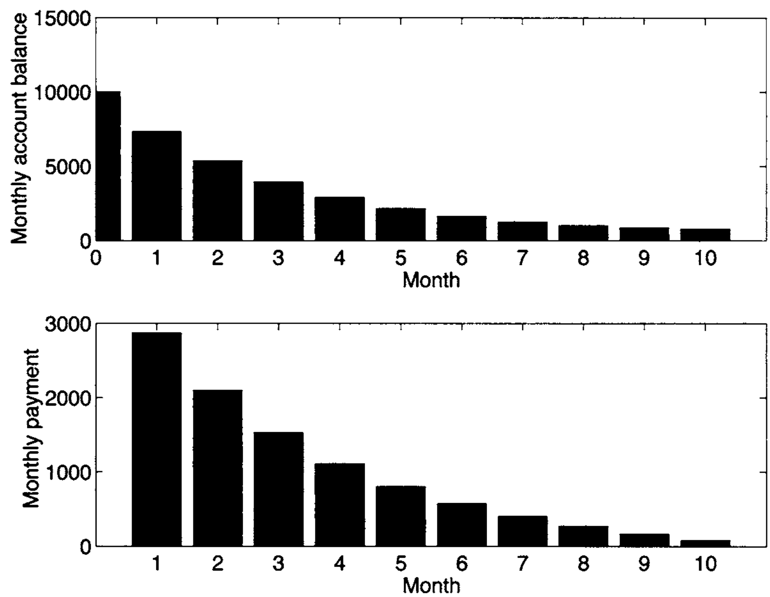

Figure 20.15 Plots for Example 20.11 with $q = 1$ and $r = 10$.

We solve our problem using the LQR framework described in Example 20.10. Let the current time be 0, x_k the account balance at the end of month k, and u_k our payment in month k. We have

$$x_k = 1.02x_{k-1} - u_k, \qquad k = 1, \dots, 10;$$

that is, the account balance in a given month is equal to the account balance in the previous month plus the monthly interest on that balance minus our payment that month. Our optimization problem is then

$$\begin{array}{ll} \text{minimize} & \dfrac{1}{2}\displaystyle\sum_{i=1}^{10} \left(qx_i^2 + ru_i^2\right) \\ \text{subject to} & x_k = 1.02x_{k-1} - u_k,\ k = 1, \dots, 10,\ x_0 = 10{,}000, \end{array}$$

which is an instance of the LQR problem. The parameters q and r reflect our priority in trading off between debt reduction and hardship in making payments. The more anxious we are to reduce our debt, the larger the value of q relative to r. On the other hand, the more reluctant we are to make payments, the larger the value of r relative to q.

The solution to the problem above is given by the formula derived in Example 20.10. In Figure 20.15 we plot the monthly account balances and payments over the next 10 months using $q = 1$ and $r = 10$. We can see here that our debt has been reduced to less than \$1000 after 10 months, but with a first

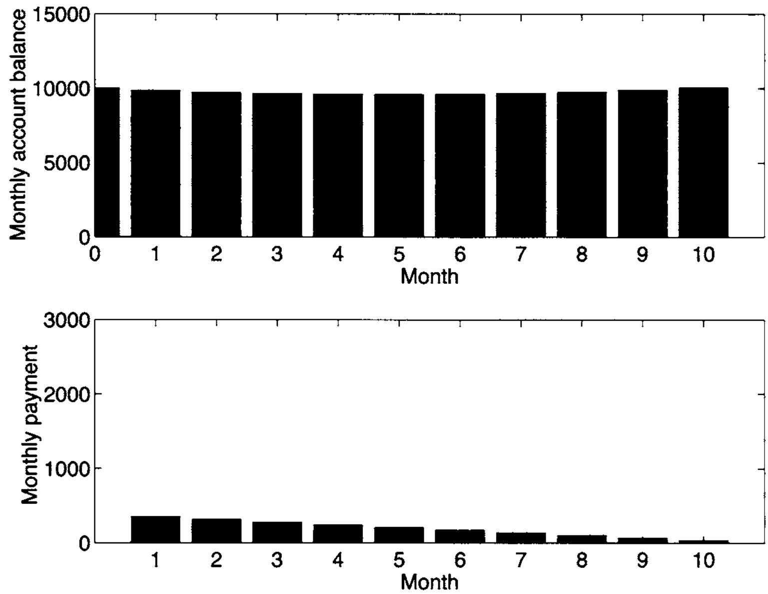

Figure 20.16 Plots for Example 20.11 with $q = 1$ and $r = 300$.

payment close to \$3000. If we feel that a payment of \$3000 is too high, then we can try to reduce this amount by increasing the value of r relative to q. However, going too far along these lines can lead to trouble. Indeed, if we use $q = 1$ and $r = 300$ (see Figure 20.16), although the monthly payments do not exceed \$400, the account balance is never reduced by much below \$10,000. In this case, the interest on the account balance eats up a significant portion of our monthly payments. In fact, our debt after 10 months will be higher than \$10,000. ■

For a treatment of optimization problems with quadratic objective functions, subject to linear or quadratic constraints, arising in communication and signal processing, see [105] and [106].

EXERCISES

20.1 Consider the following constraints on $\mathbb{R}^2$:

$$h(x_1, x_2) = (x_1 - 2)^2 = 0 \quad \text{and} \quad g(x_1, x_2) = (x_2 + 1)^3 \le 0.$$

Find the set of feasible points. Are the feasible points regular? Justify your answer.

20.2 Find local extremizers for the following optimization problems:

a.

$$\begin{aligned} \text{Minimize} \quad & x_1^2 + 2x_1x_2 + 3x_2^2 + 4x_1 + 5x_2 + 6x_3 \\ \text{subject to} \quad & x_1 + 2x_2 = 3 \\ & 4x_1 + 5x_3 = 6. \end{aligned}$$

b.

$$\begin{aligned} \text{Maximize} \quad & 4x_1 + x_2^2 \\ \text{subject to} \quad & x_1^2 + x_2^2 = 9. \end{aligned}$$

c.

$$\begin{aligned} \text{Maximize} \quad & x_1x_2 \\ \text{subject to} \quad & x_1^2 + 4x_2^2 = 1. \end{aligned}$$

20.3 Find minimizers and maximizers of the function

$$f(\boldsymbol{x}) = (\boldsymbol{a}^\top \boldsymbol{x})(\boldsymbol{b}^\top \boldsymbol{x}), \quad \boldsymbol{x} \in \mathbb{R}^3,$$

subject to

$$\begin{aligned} x_1 + x_2 &= 0 \\ x_2 + x_3 &= 0, \end{aligned}$$

where

$$\boldsymbol{a} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \boldsymbol{b} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}.$$

20.4 Consider the problem

$$\begin{aligned} \text{minimize} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & h(\boldsymbol{x}) = 0, \end{aligned}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$, $h : \mathbb{R}^2 \to \mathbb{R}$, and $\nabla f(\boldsymbol{x}) = [x_1, x_1 + 4]^\top$. Suppose that $\boldsymbol{x}^*$ is an optimal solution and $\nabla h(\boldsymbol{x}^*) = [1, 4]^\top$. Find $\nabla f(\boldsymbol{x}^*)$.

20.5 Consider the problem

$$\begin{aligned} \text{minimize} \quad & \|\boldsymbol{x} - \boldsymbol{x}_0\|^2 \\ \text{subject to} \quad & \|\boldsymbol{x}\|^2 = 9, \end{aligned}$$

where $\boldsymbol{x}_0 = [1, \sqrt{3}]^\top$.

a. Find all points satisfying the Lagrange condition for the problem.

b. Using second-order conditions, determine whether or not each of the points in part a is a local minimizer.

20.6 We wish to construct a closed box with minimum surface area that encloses a volume of V cubic feet, where $V > 0$.

a. Let a, b, and c denote the dimensions of the box with minimum surface area (with volume V). Derive the Lagrange condition that must be satisfied by a, b, and c.

b. What does it mean for a point $\boldsymbol{x}^*$ to be a *regular* point in this problem? Is the point $\boldsymbol{x}^* = [a, b, c]^\top$ a regular point?

c. Find a, b, and c.

d. Does the point $\boldsymbol{x}^* = [a, b, c]^\top$ found in part c satisfy the second-order sufficient condition?

20.7 Find local extremizers of

a. $f(x_1, x_2, x_3) = x_1^2 + 3x_2^2 + x_3$ subject to $x_1^2 + x_2^2 + x_3^2 = 16$.

b. $f(x_1, x_2) = x_1^2 + x_2^2$ subject to $3x_1^2 + 4x_1x_2 + 6x_2^2 = 140$.

20.8 Consider the problem

$$\begin{array}{rll} \text{minimize} & 2x_1 + 3x_2 - 4, & x_1, x_2 \in \mathbb{R} \\ \text{subject to} & x_1 x_2 = 6. & \end{array}$$

a. Use Lagrange's theorem to find all possible local minimizers and maximizers.

b. Use the second-order sufficient conditions to specify which points are strict local minimizers and which are strict local maximizers.

c. Are the points in part b global minimizers or maximizers? Explain.

20.9 Find all maximizers of the function

$$f(x_1, x_2) = \frac{18x_1^2 - 8x_1x_2 + 12x_2^2}{2x_1^2 + 2x_2^2}.$$

20.10 Find all solutions to the problem

$$\begin{array}{rl} \text{maximize} & \boldsymbol{x}^\top \begin{bmatrix} 3 & 4 \\ 0 & 3 \end{bmatrix} \boldsymbol{x} \\ \text{subject to} & \|\boldsymbol{x}\|^2 = 1. \end{array}$$

20.11 Consider a matrix $\boldsymbol{A}$ with the property that $\boldsymbol{A}^\top\boldsymbol{A}$ has eigenvalues ranging from 1 to 20 (i.e., the smallest eigenvalue is 1 and the largest is 20). Let $\boldsymbol{x}$ be a vector such that $\|\boldsymbol{x}\| = 1$, and let $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$. Use Lagrange multiplier methods to find the range of values that $\|\boldsymbol{y}\|$ can take.
Hint: What is the largest value that $\|\boldsymbol{y}\|$ can take? What is the smallest value that $\|\boldsymbol{y}\|$ can take?

20.12 Consider a matrix $\boldsymbol{A} \in \mathbb{R}^{m\times n}$. Define the *induced* 2-*norm* of $\boldsymbol{A}$, denoted $\|\boldsymbol{A}\|_2$, to be the number

$$\|\boldsymbol{A}\|_2 = \max\{\|\boldsymbol{A}\boldsymbol{x}\| : \boldsymbol{x} \in \mathbb{R}^n,\ \|\boldsymbol{x}\| = 1\},$$

where the norm $\|\cdot\|$ on the right-hand side above is the usual Euclidean norm.

Suppose that the eigenvalues of $\boldsymbol{A}^\top\boldsymbol{A}$ are $\lambda_1, \ldots, \lambda_n$ (ordered from largest to smallest). Use Lagrange's theorem to express $\|\boldsymbol{A}\|_2$ in terms of the eigenvalues above (cf. Theorem 3.8).

20.13 Let $\boldsymbol{P} = \boldsymbol{P}^\top$ be a positive definite matrix. Show that any point $\boldsymbol{x}$ satisfying $1 - \boldsymbol{x}^\top\boldsymbol{P}\boldsymbol{x} = 0$ is a regular point.

20.14 Consider the problem

$$\begin{array}{ll} \text{maximize} & ax_1 + bx_2, \qquad x_1, x_2 \in \mathbb{R} \\ \text{subject to} & x_1^2 + x_2^2 = 2, \end{array}$$

where $a, b \in \mathbb{R}$. Show that if $[1,1]^\top$ is a solution to the problem, then $a = b$.

20.15 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_1x_2 - 2x_1, \qquad x_1, x_2 \in \mathbb{R} \\ \text{subject to} & x_1^2 - x_2^2 = 0. \end{array}$$

a. Apply Lagrange's theorem directly to the problem to show that if a solution exists, it must be either $[1,1]^\top$ or $[-1,1]^\top$.

b. Use the second-order necessary conditions to show that $[-1,1]^\top$ cannot possibly be the solution.

c. Use the second-order sufficient conditions to show that $[1,1]^\top$ is a strict local minimizer.

20.16 Let $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $m \le n$, rank $\boldsymbol{A} = m$, and $\boldsymbol{x}_0 \in \mathbb{R}^n$. Let $\boldsymbol{x}^*$ be the point on the nullspace of $\boldsymbol{A}$ that is closest to $\boldsymbol{x}_0$ (in the sense of Euclidean norm).

a. Show that $\boldsymbol{x}^*$ is orthogonal to $\boldsymbol{x}^* - \boldsymbol{x}_0$.

b. Find a formula for $\boldsymbol{x}^*$ in terms of $\boldsymbol{A}$ and $\boldsymbol{x}_0$.

20.17 Consider the problem

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 \\ \text{subject to} \quad & \boldsymbol{C}\boldsymbol{x} = \boldsymbol{d}, \end{aligned}$$

where $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $m > n$, $\boldsymbol{C} \in \mathbb{R}^{p \times n}$, $p < n$, and both $\boldsymbol{A}$ and $\boldsymbol{C}$ are of full rank. We wish to find an expression for the solution (in terms of $\boldsymbol{A}$, $\boldsymbol{b}$, $\boldsymbol{C}$, and $\boldsymbol{d}$).

a. Apply Lagrange's theorem to solve this problem.

b. As an alternative, rewrite the given optimization problem in the form of a quadratic programming problem and apply the formula in Section 20.6 to obtain the solution.

20.18 Consider the problem of minimizing a general quadratic function subject to a linear constraint:

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - \boldsymbol{c}^\top \boldsymbol{x} + d \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}, \end{aligned}$$

where $\boldsymbol{Q} = \boldsymbol{Q}^\top > 0$, $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $m < n$, $\operatorname{rank} \boldsymbol{A} = m$, and d is a constant. Derive a closed-form solution to the problem.

20.19 Let $\boldsymbol{L}$ be an $n \times n$ real symmetric matrix, and let $\mathcal{M}$ be a subspace of $\mathbb{R}^n$ with dimension $m < n$. Let $\{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m\} \subset \mathbb{R}^n$ be a basis for $\mathcal{M}$, and let $\boldsymbol{B}$ be the $n \times m$ matrix with $\boldsymbol{b}_i$ as the ith column. Let $\boldsymbol{L}_{\mathcal{M}}$ be the $m \times m$ matrix defined by $\boldsymbol{L}_{\mathcal{M}} = \boldsymbol{B}^\top \boldsymbol{L}\boldsymbol{B}$. Show that $\boldsymbol{L}$ is positive semidefinite (definite) on $\mathcal{M}$ if and only if $\boldsymbol{L}_{\mathcal{M}}$ is positive semidefinite (definite).
Note: This result is useful for checking that the Hessian of the Lagrangian function at a point is positive definite on the tangent space at that point.

20.20 Consider the sequence $\{x_k\}$, $x_k \in \mathbb{R}$, generated by the recursion

$$x_{k+1} = a x_k + b u_k, \quad k \geq 0 \qquad (a, b \in \mathbb{R},\ a, b \neq 0),$$

where $u_0, u_1, u_2, \ldots$ is a sequence of "control inputs," and the initial condition $x_0 \neq 0$ is given. The recursion above is also called a *discrete-time linear system.* We wish to find values of control inputs u_0 and u_1 such that $x_2 = 0$, and the average input energy $(u_0^2 + u_1^2)/2$ is minimized. Denote the optimal inputs by u_0^* and u_1^*.

a. Find expressions for u_0^* and u_1^* in terms of a, b, and x_0.

b. Use the second-order sufficient conditions to show that the point $\boldsymbol{u}^* = [u_0^*, u_1^*]^\top$ in part a is a strict local minimizer.

20.21 Consider the discrete-time linear system $x_k = 2x_{k-1} + u_k$, $k \geq 1$, with $x_0 = 1$. Find the values of the control inputs u_1 and u_2 to minimize

$$x_2^2 + \frac{1}{2}u_1^2 + \frac{1}{3}u_2^2.$$

20.22 Consider the discrete-time linear system $x_{k+1} = x_k + 2u_k$, $0 \leq k \leq 2$, with $x_0 = 3$. Use the Lagrange multiplier approach to calculate the optimal control sequence $\{u_0, u_1, u_2\}$ that transfers the initial state x_0 to $x_3 = 9$ while minimizing

$$\frac{1}{2}\sum_{k=0}^{2} u_k^2.$$

CHAPTER 21

PROBLEMS WITH INEQUALITY CONSTRAINTS

21.1 Karush-Kuhn-Tucker Condition

In Chapter 20 we analyzed constrained optimization problems involving only equality constraints. In this chapter we discuss extremum problems that also involve inequality constraints. The treatment in this chapter parallels that of Chapter 20. In particular, as we shall see, problems with inequality constraints can also be treated using Lagrange multipliers.

We consider the following problem:

$$\begin{aligned} \text{minimize} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}, \\ & \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}, \end{aligned}$$

where $f : \mathbb{R}^n \to \mathbb{R}$, $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $m \leq n$, and $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^p$. For the general problem above, we adopt the following definitions.

Definition 21.1 An inequality constraint $g_j(\boldsymbol{x}) \leq 0$ is said to be *active* at $\boldsymbol{x}^*$ if $g_j(\boldsymbol{x}^*) = 0$. It is *inactive* at $\boldsymbol{x}^*$ if $g_j(\boldsymbol{x}^*) < 0$. ∎

An Introduction to Optimization, Fourth Edition.

By convention, we consider an equality constraint $h_i(\boldsymbol{x}) = 0$ to be always active.

Definition 21.2 Let $\boldsymbol{x}^*$ satisfy $\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}$, $\boldsymbol{g}(\boldsymbol{x}^*) \leq \boldsymbol{0}$, and let $J(\boldsymbol{x}^*)$ be the index set of active inequality constraints:

$$J(\boldsymbol{x}^*) \triangleq \{j : g_j(\boldsymbol{x}^*) = 0\}.$$

Then, we say that $\boldsymbol{x}^*$ is a *regular point* if the vectors

$$\nabla h_i(\boldsymbol{x}^*), \quad \nabla g_j(\boldsymbol{x}^*), \quad 1 \leq i \leq m, \quad j \in J(\boldsymbol{x}^*)$$

are linearly independent. ∎

We now prove a first-order necessary condition for a point to be a local minimizer. We call this condition the *Karush-Kuhn-Tucker (KKT) condition.* In the literature, this condition is sometimes also called the Kuhn-Tucker condition.

Theorem 21.1 ***Karush-Kuhn-Tucker (KKT) Theorem.*** *Let $f, \boldsymbol{h}, \boldsymbol{g} \in \mathcal{C}^1$. Let $\boldsymbol{x}^*$ be a regular point and a local minimizer for the problem of minimizing f subject to $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}$. Then, there exist $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ and $\boldsymbol{\mu}^* \in \mathbb{R}^p$ such that:*

1. $\boldsymbol{\mu}^* \geq \boldsymbol{0}$.
2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$.
3. $\boldsymbol{\mu}^{*\top} \boldsymbol{g}(\boldsymbol{x}^*) = 0$.

□

In Theorem 21.1, we refer to $\boldsymbol{\lambda}^*$ as the Lagrange multiplier vector and $\boldsymbol{\mu}^*$ as the Karush-Kuhn-Tucker (KKT) multiplier vector. We refer to their components as Lagrange multipliers and Karush-Kuhn-Tucker (KKT) multipliers, respectively.

Before proving this theorem, let us first discuss its meaning. Observe that $\mu_j^* \geq 0$ (by condition 1) and $g_j(\boldsymbol{x}^*) \leq 0$. Therefore, the condition

$$\boldsymbol{\mu}^{*\top} \boldsymbol{g}(\boldsymbol{x}^*) = \mu_1^* g_1(\boldsymbol{x}^*) + \cdots + \mu_p^* g_p(\boldsymbol{x}^*) = 0$$

implies that if $g_j(\boldsymbol{x}^*) < 0$, then $\mu_j^* = 0$; that is, for all $j \notin J(\boldsymbol{x}^*)$, we have $\mu_j^* = 0$. In other words, the KKT multipliers μ_j^* corresponding to inactive constraints are zero. The other KKT multipliers, μ_i^*, $i \in J(\boldsymbol{x}^*)$, are nonnegative; they may or may not be equal to zero.

Example 21.1 A graphical illustration of the KKT theorem is given in Figure 21.1. In this two-dimensional example, we have only inequality constraints

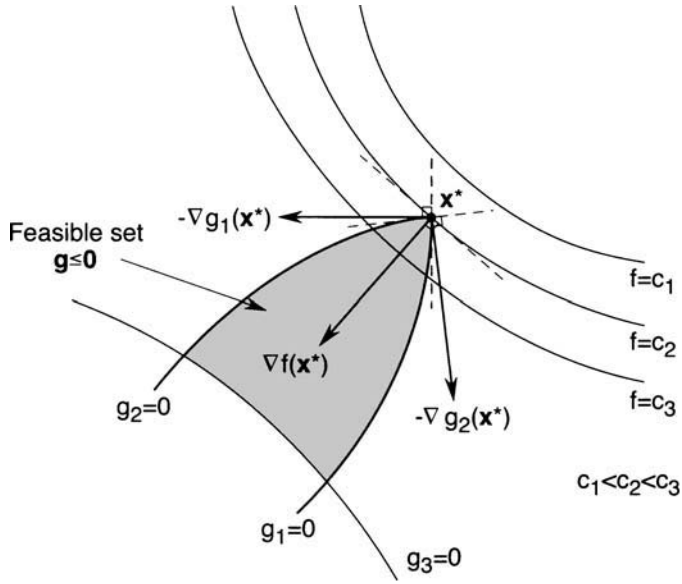

Figure 21.1 Illustration of the Karush-Kuhn-Tucker (KKT) theorem.

$g_j(\boldsymbol{x}) \leq 0$, $j = 1, 2, 3$. Note that the point $\boldsymbol{x}^*$ in the figure is indeed a minimizer. The constraint $g_3(\boldsymbol{x}) \leq 0$ is inactive: $g_3(\boldsymbol{x}^*) < 0$; hence $\mu_3^* = 0$. By the KKT theorem, we have

$$\nabla f(\boldsymbol{x}^*) + \mu_1^* \nabla g_1(\boldsymbol{x}^*) + \mu_2^* \nabla g_2(\boldsymbol{x}^*) = \mathbf{0},$$

or, equivalently,

$$\nabla f(\boldsymbol{x}^*) = -\mu_1^* \nabla g_1(\boldsymbol{x}^*) - \mu_2^* \nabla g_2(\boldsymbol{x}^*),$$

where $\mu_1^* > 0$ and $\mu_2^* > 0$. It is easy to interpret the KKT condition graphically for this example. Specifically, we can see from Figure 21.1 that $\nabla f(\boldsymbol{x}^*)$ must be a linear combination of the vectors $-\nabla g_1(\boldsymbol{x}^*)$ and $-\nabla g_2(\boldsymbol{x}^*)$ with positive coefficients. This is reflected exactly in the equation above, where the coefficients μ_1^* and μ_2^* are the KKT multipliers. ∎

We apply the KKT condition in the same way that we apply any necessary condition. Specifically, we search for points satisfying the KKT condition and treat these points as candidate minimizers. To summarize, the KKT condition consists of five parts (three equations and two inequalities):

1. $\boldsymbol{\mu}^* \geq \mathbf{0}$.
2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^{*\top} \boldsymbol{g}(\boldsymbol{x}^*) = 0$.

4. $\boldsymbol{h}(\boldsymbol{x}^*) = \mathbf{0}$.

5. $\boldsymbol{g}(\boldsymbol{x}^*) \leq \mathbf{0}$.

We now prove the KKT theorem.

Proof of the Karush-Kuhn-Tucker Theorem. Let $\boldsymbol{x}^*$ be a regular local minimizer of f on the set $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0}, \boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}\}$. Then, $\boldsymbol{x}^*$ is also a regular local minimizer of f on the set $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0}, g_j(\boldsymbol{x}) = 0, j \in J(\boldsymbol{x}^*)\}$ (see Exercise 21.16). Note that the latter constraint set involves only equality constraints. Therefore, from Lagrange's theorem, it follows that there exist vectors $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ and $\boldsymbol{\mu}^* \in \mathbb{R}^p$ such that

$$Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top,$$

where for all $j \notin J(\boldsymbol{x}^*)$, we have $\mu_j^* = 0$. To complete the proof it remains to show that for all $j \in J(\boldsymbol{x}^*)$, we have $\mu_j^* \geq 0$ (and hence for all $j = 1, \ldots, p$, we have $\mu_j^* \geq 0$, i.e., $\boldsymbol{\mu}^* \geq \mathbf{0}$). We use a proof by contradiction. So suppose that there exists $j \in J(\boldsymbol{x}^*)$ such that $\mu_j^* < 0$. Let $\hat{S}$ and $\hat{T}(\boldsymbol{x}^*)$ be the surface and tangent space defined by all other active constraints at $\boldsymbol{x}^*$. Specifically,

$$\hat{S} = \{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0}, g_i(\boldsymbol{x}) = 0, i \in J(\boldsymbol{x}^*), i \neq j\}$$

and

$$\hat{T}(\boldsymbol{x}^*) = \{\boldsymbol{y} : D\boldsymbol{h}(\boldsymbol{x}^*)\boldsymbol{y} = \mathbf{0}, Dg_i(\boldsymbol{x}^*)\boldsymbol{y} = 0, i \in J(\boldsymbol{x}^*), i \neq j\}.$$

We claim that by the regularity of $\boldsymbol{x}^*$, there exists $\boldsymbol{y} \in \hat{T}(\boldsymbol{x}^*)$ such that

$$Dg_j(\boldsymbol{x}^*)\boldsymbol{y} \neq 0.$$

To see this, suppose that for all $\boldsymbol{y} \in \hat{T}(\boldsymbol{x}^*)$, $\nabla g_j(\boldsymbol{x}^*)^\top \boldsymbol{y} = Dg_j(\boldsymbol{x}^*)\boldsymbol{y} = 0$. This implies that $\nabla g_j(\boldsymbol{x}^*) \in \hat{T}(\boldsymbol{x}^*)^\perp$. By Lemma 20.1, this, in turn, implies that

$$\nabla g_j(\boldsymbol{x}^*) \in \operatorname{span}[\nabla h_k(\boldsymbol{x}^*), k = 1, \ldots, m, \nabla g_i(\boldsymbol{x}^*), i \in J(\boldsymbol{x}^*), i \neq j].$$

But this contradicts the fact that $\boldsymbol{x}^*$ is a regular point, which proves our claim. Without loss of generality, we assume that we have $\boldsymbol{y}$ such that $Dg_j(\boldsymbol{x}^*)\boldsymbol{y} < 0$.

Consider the Lagrange condition, rewritten as

$$Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \mu_j^* Dg_j(\boldsymbol{x}^*) + \sum_{i \neq j} \mu_i^* Dg_i(\boldsymbol{x}^*) = \mathbf{0}^\top.$$

If we postmultiply the above by $\boldsymbol{y}$ and use the fact that $\boldsymbol{y} \in \hat{T}(\boldsymbol{x}^*)$, we get

$$Df(\boldsymbol{x}^*)\boldsymbol{y} = -\mu_j^* Dg_j(\boldsymbol{x}^*)\boldsymbol{y}.$$

Because $Dg_j(\boldsymbol{x}^*)\boldsymbol{y} < 0$ and we have assumed that $\mu_j^* < 0$, we have

$$Df(\boldsymbol{x}^*)\boldsymbol{y} < 0.$$

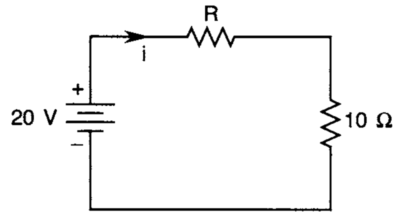

Figure 21.2 Circuit in Example 21.2.

Because $\boldsymbol{y} \in \hat{T}(\boldsymbol{x}^*)$, by Theorem 20.1 we can find a differentiable curve $\{\boldsymbol{x}(t) : t \in (a, b)\}$ on $\hat{S}$ such that there exists $t^* \in (a, b)$ with $\boldsymbol{x}(t^*) = \boldsymbol{x}^*$ and $\dot{\boldsymbol{x}}(t^*) = \boldsymbol{y}$. Now,

$$\frac{d}{dt} f(\boldsymbol{x}(t^*)) = Df(\boldsymbol{x}^*)\boldsymbol{y} < 0.$$

The above means that there is a $\delta > 0$ such that for all $t \in (t^*, t^* + \delta]$, we have

$$f(\boldsymbol{x}(t)) < f(\boldsymbol{x}(t^*)) = f(\boldsymbol{x}^*).$$

On the other hand,

$$\frac{d}{dt} g_j(\boldsymbol{x}(t^*)) = Dg_j(\boldsymbol{x}^*)\boldsymbol{y} < 0,$$

and for some $\varepsilon > 0$ and all $t \in [t^*, t^* + \varepsilon]$, we have that $g_j(\boldsymbol{x}(t)) \leq 0$. Therefore, for all $t \in (t^*, t^* + \min\{\delta, \varepsilon\}]$, we have that $g_j(\boldsymbol{x}(t)) \leq 0$ and $f(\boldsymbol{x}(t)) < f(\boldsymbol{x}^*)$. Because the points $\boldsymbol{x}(t)$, $t \in (t^*, t^* + \min\{\delta, \varepsilon\}]$, are in $\hat{S}$, they are feasible points with lower objective function values than $\boldsymbol{x}^*$. This contradicts the assumption that $\boldsymbol{x}^*$ is a local minimizer, which completes the proof. ∎

Example 21.2 Consider the circuit in Figure 21.2. Formulate and solve the KKT condition for the following problems.

a. Find the value of the resistor $R \geq 0$ such that the power absorbed by this resistor is maximized.

b. Find the value of the resistor $R \geq 0$ such that the power delivered to the 10-Ω resistor is maximized.

Solution:

a. The power absorbed by the resistor R is $p = i^2 R$, where $i = \frac{20}{10+R}$. The optimization problem can be represented as

$$\begin{aligned} \text{minimize} \quad & -\frac{400R}{(10+R)^2} \\ \text{subject to} \quad & -R \leq 0. \end{aligned}$$

The derivative of the objective function is

$$-\frac{400(10+R)^2 - 800R(10+R)}{(10+R)^4} = -\frac{400(10-R)}{(10+R)^3}.$$

Thus, the KKT condition is

$$\begin{aligned} -\frac{400(10-R)}{(10+R)^3} - \mu &= 0, \\ \mu &\geq 0, \\ \mu R &= 0, \\ -R &\leq 0. \end{aligned}$$

We consider two cases. In the first case, suppose that $\mu > 0$. Then, $R = 0$. But this contradicts the first condition above. Now suppose that $\mu = 0$. Then, by the first condition, we have $R = 10$. Therefore, the only solution to the KKT condition is $R = 10$, $\mu = 0$.

b. The power absorbed by the 10-Ω resistor is $p = i^2 10$, where $i = 20/(10 + R)$. The optimization problem can be represented as

$$\begin{array}{ll} \text{minimize} & -\dfrac{4000}{(10+R)^2} \\ \text{subject to} & -R \leq 0. \end{array}$$

The derivative of the objective function is

$$\frac{8000}{(10+R)^3}.$$

Thus, the KKT condition is

$$\begin{aligned} \frac{8000}{(10+R)^3} - \mu &= 0, \\ \mu &\geq 0, \\ \mu R &= 0, \\ -R &\leq 0. \end{aligned}$$

As before, we consider two cases. In the first case, suppose that $\mu > 0$. Then, $R = 0$, which is feasible. For the second case, suppose that $\mu = 0$. But this contradicts the first condition. Therefore, the only solution to the KKT condition is $R = 0$, $\mu = 8$. ∎

In the case when the objective function is to be maximized, that is, when the optimization problem has the form

$$\begin{array}{ll} \text{maximize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \\ & \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}, \end{array}$$

the KKT condition can be written as

1. $\boldsymbol{\mu}^* \geq \mathbf{0}$.
2. $-Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
4. $\boldsymbol{h}(\boldsymbol{x}^*) = \mathbf{0}$.
5. $\boldsymbol{g}(\boldsymbol{x}^*) \leq \mathbf{0}$.

The above is easily derived by converting the maximization problem above into a minimization problem, by multiplying the objective function by -1. It can be further rewritten as

1. $\boldsymbol{\mu}^* \leq \mathbf{0}$.
2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
4. $\boldsymbol{h}(\boldsymbol{x}^*) = \mathbf{0}$.
5. $\boldsymbol{g}(\boldsymbol{x}^*) \leq \mathbf{0}$.

The form shown above is obtained from the preceding one by changing the signs of $\boldsymbol{\mu}^*$ and $\boldsymbol{\lambda}^*$ and multiplying condition 2 by -1.

We can similarly derive the KKT condition for the case when the inequality constraint is of the form $\boldsymbol{g}(\boldsymbol{x}) \geq \mathbf{0}$. Specifically, consider the problem

$$\begin{array}{rl} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \mathbf{0}. \end{array}$$

We multiply the inequality constraint function by -1 to obtain $-\boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0}$. Thus, the KKT condition for this case is

1. $\boldsymbol{\mu}^* \geq \mathbf{0}$.
2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) - \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
4. $\boldsymbol{h}(\boldsymbol{x}^*) = \mathbf{0}$.
5. $\boldsymbol{g}(\boldsymbol{x}^*) \geq \mathbf{0}$.

Changing the sign of $\boldsymbol{\mu}^*$ as before, we obtain

1. $\boldsymbol{\mu}^* \leq \mathbf{0}$.

2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
4. $\boldsymbol{h}(\boldsymbol{x}^*) = \mathbf{0}$.
5. $\boldsymbol{g}(\boldsymbol{x}^*) \geq \mathbf{0}$.

For the problem

$$\begin{array}{ll} \text{maximize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \mathbf{0} \\ & \boldsymbol{g}(\boldsymbol{x}) \geq \mathbf{0}, \end{array}$$

the KKT condition is exactly the same as in Theorem 21.1, except for the reversal of the inequality constraint.

Example 21.3 In Figure 21.3, the two points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are feasible points; that is, $g(\boldsymbol{x}_1) \geq 0$ and $g(\boldsymbol{x}_2) \geq 0$, and they satisfy the KKT condition.

The point $\boldsymbol{x}_1$ is a maximizer. The KKT condition for this point (with KKT multiplier μ_1) is

1. $\mu_1 \geq 0$.
2. $\nabla f(\boldsymbol{x}_1) + \mu_1 \nabla g(\boldsymbol{x}_1) = \mathbf{0}$.
3. $\mu_1 g(\boldsymbol{x}_1) = 0$.
4. $g(\boldsymbol{x}_1) \geq 0$.

The point $\boldsymbol{x}_2$ is a minimizer of f. The KKT condition for this point (with KKT multiplier μ_2) is

1. $\mu_2 \leq 0$.
2. $\nabla f(\boldsymbol{x}_2) + \mu_2 \nabla g(\boldsymbol{x}_2) = \mathbf{0}$.
3. $\mu_2 g(\boldsymbol{x}_2) = 0$.
4. $g(\boldsymbol{x}_2) \geq 0$.

∎

Example 21.4 Consider the problem

$$\begin{array}{ll} \text{minimize} & f(x_1, x_2) \\ \text{subject to} & x_1, x_2 \geq 0, \end{array}$$

where

$$f(x_1, x_2) = x_1^2 + x_2^2 + x_1 x_2 - 3x_1.$$

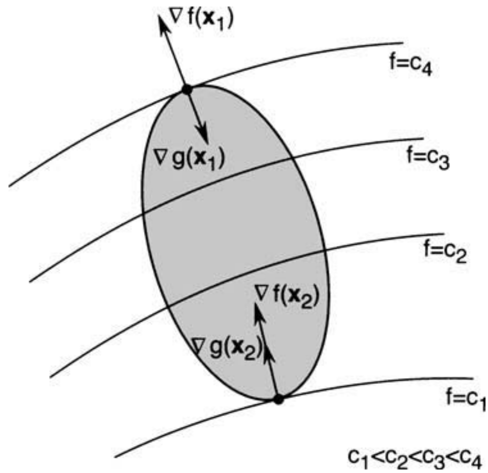

Figure 21.3 Points satisfying the KKT condition ($\boldsymbol{x}_1$ is a maximizer and $\boldsymbol{x}_2$ is a minimizer).

The KKT condition for this problem is

1. $\boldsymbol{\mu} = [\mu_1, \mu_2]^\top \leq \mathbf{0}$.
2. $Df(\boldsymbol{x}) + \boldsymbol{\mu}^\top = \mathbf{0}^\top$.
3. $\boldsymbol{\mu}^\top \boldsymbol{x} = 0$.
4. $\boldsymbol{x} \geq \mathbf{0}$.

We have

$$Df(\boldsymbol{x}) = [2x_1 + x_2 - 3,\ x_1 + 2x_2].$$

This gives

$$\begin{aligned} 2x_1 + x_2 + \mu_1 &= 3, \\ x_1 + 2x_2 + \mu_2 &= 0, \\ \mu_1 x_1 + \mu_2 x_2 &= 0. \end{aligned}$$

We now have four variables, three equations, and the inequality constraints on each variable. To find a solution $(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$, we first try

$$\mu_1^* = 0, \quad x_2^* = 0,$$

which gives

$$x_1^* = \frac{3}{2}, \quad \mu_2^* = -\frac{3}{2}.$$

The above satisfies all the KKT and feasibility conditions. In a similar fashion, we can try

$$\mu_2^* = 0, \quad x_1^* = 0,$$

which gives

$$x_2^* = 0, \quad \mu_1^* = 3.$$

This point clearly violates the nonpositivity constraint on μ_1^*.

The feasible point above satisfying the KKT condition is only a candidate for a minimizer. However, there is no guarantee that the point is indeed a minimizer, because the KKT condition is, in general, only necessary. A sufficient condition for a point to be a minimizer is given in the next section. ∎

Example 21.4 is a special case of a more general problem of the form

$$\begin{aligned} \text{minimize} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & \boldsymbol{x} \geq \mathbf{0}. \end{aligned}$$

The KKT condition for this problem has the form

$$\begin{aligned} \boldsymbol{\mu} &\leq \mathbf{0}, \\ \nabla f(\boldsymbol{x}) + \boldsymbol{\mu} &= \mathbf{0}, \\ \boldsymbol{\mu}^\top \boldsymbol{x} &= 0, \\ \boldsymbol{x} &\geq \mathbf{0}. \end{aligned}$$

From the above, we can eliminate $\boldsymbol{\mu}$ to obtain

$$\begin{aligned} \nabla f(\boldsymbol{x}) &\geq \mathbf{0}, \\ \boldsymbol{x}^\top \nabla f(\boldsymbol{x}) &= 0, \\ \boldsymbol{x} &\geq \mathbf{0}. \end{aligned}$$

Some possible points in $\mathbb{R}^2$ that satisfy these conditions are depicted in Figure 21.4.

For further results related to the KKT condition, we refer the reader to [90, Chapter 7].

21.2 Second-Order Conditions

As in the case of extremum problems with equality constraints, we can also give second-order necessary and sufficient conditions for extremum problems involving inequality constraints. For this, we need to define the following matrix:

$$\boldsymbol{L}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \boldsymbol{F}(\boldsymbol{x}) + [\boldsymbol{\lambda} \boldsymbol{H}(\boldsymbol{x})] + [\boldsymbol{\mu} \boldsymbol{G}(\boldsymbol{x})],$$

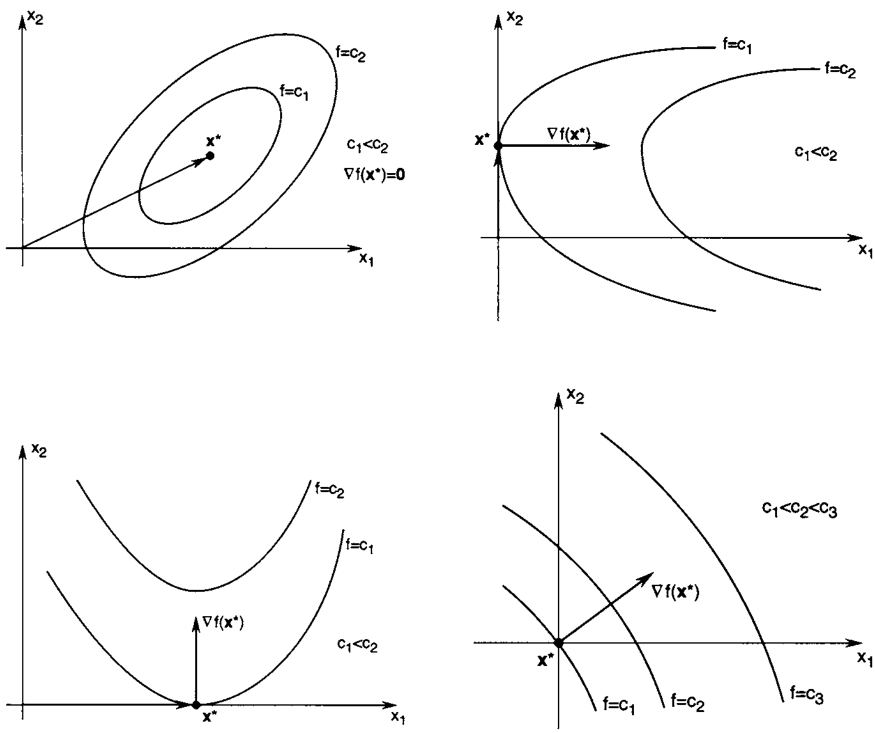

Figure 21.4 Some possible points satisfying the KKT condition for problems with positive constraints. (Adapted from [13].)

where $\boldsymbol{F}(\boldsymbol{x})$ is the Hessian matrix of f at $\boldsymbol{x}$, and the notation $[\boldsymbol{\lambda H}(\boldsymbol{x})]$ represents

$$[\boldsymbol{\lambda H}(\boldsymbol{x})] = \lambda_1 \boldsymbol{H}_1(\boldsymbol{x}) + \cdots + \lambda_m \boldsymbol{H}_m(\boldsymbol{x}),$$

as before. Similarly, the notation $[\boldsymbol{\mu G}(\boldsymbol{x})]$ represents

$$[\boldsymbol{\mu G}(\boldsymbol{x})] = \mu_1 \boldsymbol{G}_1(\boldsymbol{x}) + \cdots + \mu_p \boldsymbol{G}_p(\boldsymbol{x}),$$

where $\boldsymbol{G}_k(\boldsymbol{x})$ is the Hessian of g_k at $\boldsymbol{x}$, given by

$$\boldsymbol{G}_k(\boldsymbol{x}) = \begin{bmatrix} \frac{\partial^2 g_k}{\partial x_1^2}(\boldsymbol{x}) & \cdots & \frac{\partial^2 g_k}{\partial x_n \partial x_1}(\boldsymbol{x}) \\ \vdots & & \vdots \\ \frac{\partial^2 g_k}{\partial x_1 \partial x_n}(\boldsymbol{x}) & \cdots & \frac{\partial^2 g_k}{\partial^2 x_n}(\boldsymbol{x}) \end{bmatrix}.$$

In the following theorem, we use

$$T(\boldsymbol{x}^*) = \{\boldsymbol{y} \in \mathbb{R}^n : D\boldsymbol{h}(\boldsymbol{x}^*)\boldsymbol{y} = \boldsymbol{0},\ Dg_j(\boldsymbol{x}^*)\boldsymbol{y} = 0,\ j \in J(\boldsymbol{x}^*)\},$$

that is, the tangent space to the surface defined by active constraints.

Theorem 21.2 ***Second-Order Necessary Conditions.*** *Let $\boldsymbol{x}^*$ be a local minimizer of $f : \mathbb{R}^n \to \mathbb{R}$ subject to $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{g}(\boldsymbol{x}) \le \boldsymbol{0}$, $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $m \le n$, $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^p$, and $f, \boldsymbol{h}, \boldsymbol{g} \in \mathcal{C}^2$. Suppose that $\boldsymbol{x}^*$ is regular. Then, there exist $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ and $\boldsymbol{\mu}^* \in \mathbb{R}^p$ such that:*

1. $\boldsymbol{\mu}^* \ge \boldsymbol{0}$, $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$, $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
2. *For all* $\boldsymbol{y} \in T(\boldsymbol{x}^*)$ *we have* $\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)\boldsymbol{y} \ge 0$. □

Proof. Part 1 is simply a result of the KKT theorem. To prove part 2, we note that because the point $\boldsymbol{x}^*$ is a local minimizer over $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0},\ \boldsymbol{g}(\boldsymbol{x}) \le \boldsymbol{0}\}$, it is also a local minimizer over $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0},\ g_j(\boldsymbol{x}) = 0,\ j \in J(\boldsymbol{x}^*)\}$; that is, the point $\boldsymbol{x}^*$ is a local minimizer with active constraints taken as equality constraints (see Exercise 21.16). Hence, the second-order necessary conditions for equality constraints (Theorem 20.4) are applicable here, which completes the proof. ∎

We now state the second-order sufficient conditions for extremum problems involving inequality constraints. In the formulation of the result, we use the following set:

$$\tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*) = \{\boldsymbol{y} : D\boldsymbol{h}(\boldsymbol{x}^*)\boldsymbol{y} = \boldsymbol{0}, Dg_i(\boldsymbol{x}^*)\boldsymbol{y} = 0, i \in \tilde{J}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)\},$$

where $\tilde{J}(\boldsymbol{x}^*, \boldsymbol{\mu}^*) = \{i : g_i(\boldsymbol{x}^*) = 0, \mu_i^* > 0\}$. Note that $\tilde{J}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$ is a subset of $J(\boldsymbol{x}^*)$: $\tilde{J}(\boldsymbol{x}^*, \boldsymbol{\mu}^*) \subset J(\boldsymbol{x}^*)$. This, in turn, implies that $T(\boldsymbol{x}^*)$ is a subset of $\tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$: $T(\boldsymbol{x}^*) \subset \tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$.

Theorem 21.3 ***Second-Order Sufficient Conditions.*** *Suppose that $f, \boldsymbol{g}, \boldsymbol{h} \in \mathcal{C}^2$ and there exist a feasible point $\boldsymbol{x}^* \in \mathbb{R}^n$ and vectors $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ and $\boldsymbol{\mu}^* \in \mathbb{R}^p$ such that:*

1. $\boldsymbol{\mu}^* \ge \boldsymbol{0}$, $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$, $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$.
2. *For all* $\boldsymbol{y} \in \tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$, $\boldsymbol{y} \ne \boldsymbol{0}$, *we have* $\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)\boldsymbol{y} > 0$.

Then, $\boldsymbol{x}^*$ is a strict local minimizer of f subject to $\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}$, $\boldsymbol{g}(\boldsymbol{x}) \le \boldsymbol{0}$. □

Proof. For a proof of this theorem, we refer the reader to [88, p. 345]. ∎

A result similar to Theorem 21.3 holds for a strict local maximizer, the only difference being that we need $\boldsymbol{\mu}^* \le \boldsymbol{0}$ and that $\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*)$ be negative definite on $\tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$.

Example 21.5 Consider the following problem:

$$\begin{aligned} \text{minimize} \quad & x_1 x_2 \\ \text{subject to} \quad & x_1 + x_2 \ge 2 \\ & x_2 \ge x_1. \end{aligned}$$

a. Write down the KKT condition for this problem.

b. Find all points (and KKT multipliers) satisfying the KKT condition. In each case, determine if the point is regular.

c. Find all points in part b that also satisfy the SONC.

d. Find all points in part c that also satisfy the SOSC.

e. Find all points in part c that are local minimizers.

Solution:

a. Write $f(\boldsymbol{x}) = x_1x_2$, $g_1(\boldsymbol{x}) = 2 - x_1 - x_2$, and $g_2(\boldsymbol{x}) = x_1 - x_2$. The KKT condition is

$$\begin{aligned} x_2 - \mu_1 + \mu_2 &= 0, \\ x_1 - \mu_1 - \mu_2 &= 0, \\ \mu_1(2 - x_1 - x_2) + \mu_2(x_1 - x_2) &= 0, \\ \mu_1, \mu_2 &\geq 0, \\ 2 - x_1 - x_2 &\leq 0, \\ x_1 - x_2 &\leq 0. \end{aligned}$$

b. It is easy to check that $\mu_1 \neq 0$ and $\mu_2 \not> 0$. This leaves us with only one solution to the KKT condition: $x_1^* = x_2^* = 1$, $\mu_1^* = 1$, $\mu_2^* = 0$. For this point, we have $Dg_1(\boldsymbol{x}^*) = [-1, -1]$ and $Dg_2(\boldsymbol{x}^*) = [1, -1]$. Hence, $\boldsymbol{x}^*$ is regular.

c. Both constraints are active. Hence, because $\boldsymbol{x}^*$ is regular, $T(\boldsymbol{x}^*) = \{\mathbf{0}\}$. This implies that the SONC is satisfied.

d. Now,

$$\boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\mu}^*) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Moreover, $\tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*) = \{\boldsymbol{y} : [-1, -1]\boldsymbol{y} = 0\} = \{\boldsymbol{y} : y_1 = -y_2\}$. Pick $\boldsymbol{y} = [1, -1]^\top \in \tilde{T}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)$. We have $\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \boldsymbol{\mu}^*)\boldsymbol{y} = -2 < 0$, which means that the SOSC fails.

e. In fact, the point $\boldsymbol{x}^*$ is not a local minimizer. To see this, draw a picture of the constraint set and level sets of the objective function. Moving in the feasible direction $[1, 1]^\top$, the objective function increases; but moving in the feasible direction $[-1, 1]^\top$, the objective function decreases.

∎

We now solve analytically the problem in Example 20.1 that we solved graphically earlier.

Example 21.6 We wish to minimize $f(\boldsymbol{x}) = (x_1 - 1)^2 + x_2 - 2$ subject to

$$\begin{aligned} h(\boldsymbol{x}) &= x_2 - x_1 - 1 = 0, \\ g(\boldsymbol{x}) &= x_1 + x_2 - 2 \leq 0. \end{aligned}$$

For all $\boldsymbol{x} \in \mathbb{R}^2$, we have

$$Dh(\boldsymbol{x}) = [-1, 1], \quad Dg(\boldsymbol{x}) = [1, 1].$$

Thus, $\nabla h(\boldsymbol{x})$ and $\nabla g(\boldsymbol{x})$ are linearly independent and hence all feasible points are regular. We first write the KKT condition. Because $Df(\boldsymbol{x}) = [2x_1 - 2, 1]$, we have

$$\begin{aligned} Df(\boldsymbol{x}) + \lambda Dh(\boldsymbol{x}) + \mu Dg(\boldsymbol{x}) = [2x_1 - 2 - \lambda + \mu, 1 + \lambda + \mu] &= \boldsymbol{0}^\top, \\ \mu(x_1 + x_2 - 2) &= 0, \\ \mu &\geq 0, \\ x_2 - x_1 - 1 &= 0, \\ x_1 + x_2 - 2 &\leq 0. \end{aligned}$$

To find points that satisfy the conditions above, we first try $\mu > 0$, which implies that $x_1 + x_2 - 2 = 0$. Thus, we are faced with a system of four linear equations

$$\begin{aligned} 2x_1 - 2 - \lambda + \mu &= 0, \\ 1 + \lambda + \mu &= 0, \\ x_2 - x_1 - 1 &= 0, \\ x_1 + x_2 - 2 &= 0. \end{aligned}$$

Solving the system of equations above, we obtain

$$x_1 = \frac{1}{2}, \quad x_2 = \frac{3}{2}, \quad \lambda = -1, \quad \mu = 0.$$

However, the above is not a legitimate solution to the KKT condition, because we obtained $\mu = 0$, which contradicts the assumption that $\mu > 0$.

In the second try, we assume that $\mu = 0$. Thus, we have to solve the system of equations

$$\begin{aligned} 2x_1 - 2 - \lambda &= 0, \\ 1 + \lambda &= 0, \\ x_2 - x_1 - 1 &= 0, \end{aligned}$$

and the solutions must satisfy

$$g(x_1, x_2) = x_1 + x_2 - 2 \leq 0.$$

Solving the equations above, we obtain

$$x_1 = \frac{1}{2}, \quad x_2 = \frac{3}{2}, \quad \lambda = -1.$$

Note that $\boldsymbol{x}^* = [1/2, 3/2]^\top$ satisfies the constraint $g(\boldsymbol{x}^*) \leq 0$. The point $\boldsymbol{x}^*$ satisfying the KKT necessary condition is therefore the candidate for being a minimizer.

We now verify if $\boldsymbol{x}^* = [1/2, 3/2]^\top$, $\lambda^* = -1$, $\mu^* = 0$, satisfy the second-order sufficient conditions. For this, we form the matrix

$$\begin{aligned}
\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*, \mu^*) &= \boldsymbol{F}(\boldsymbol{x}^*) + \lambda^* \boldsymbol{H}(\boldsymbol{x}^*) + \mu^* \boldsymbol{G}(\boldsymbol{x}^*) \\
&= \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix} + (-1) \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} + (0) \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \\
&= \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}.
\end{aligned}$$

We then find the subspace

$$\tilde{T}(\boldsymbol{x}^*, \mu^*) = \{\boldsymbol{y} : Dh(\boldsymbol{x}^*)\boldsymbol{y} = 0\}.$$

Note that because $\mu^* = 0$, the active constraint $g(\boldsymbol{x}^*) = 0$ does not enter the computation of $\tilde{T}(\boldsymbol{x}^*, \mu^*)$. Note also that in this case, $T(\boldsymbol{x}^*) = \{\boldsymbol{0}\}$. We have

$$\tilde{T}(\boldsymbol{x}^*, \mu^*) = \{\boldsymbol{y} : [-1, 1]\boldsymbol{y} = 0\} = \{[a, a]^\top : a \in \mathbb{R}\}.$$

We then check for positive definiteness of $\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*, \mu^*)$ on $\tilde{T}(\boldsymbol{x}^*, \mu^*)$. We have

$$\boldsymbol{y}^\top \boldsymbol{L}(\boldsymbol{x}^*, \lambda^*, \mu^*)\boldsymbol{y} = [a, a] \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} a \\ a \end{bmatrix} = 2a^2.$$

Thus, $\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*, \mu^*)$ is positive definite on $\tilde{T}(\boldsymbol{x}^*, \mu^*)$. Observe that $\boldsymbol{L}(\boldsymbol{x}^*, \lambda^*, \mu^*)$ is, in fact, only positive semidefinite on $\mathbb{R}^2$.

By the second-order sufficient conditions, we conclude that $\boldsymbol{x}^* = [1/2, 3/2]^\top$ is a strict local minimizer. ∎

EXERCISES

21.1 Consider the optimization problem

$$\begin{aligned}
\text{minimize} \quad & x_1^2 + 4x_2^2 \\
\text{subject to} \quad & x_1^2 + 2x_2^2 \geq 4.
\end{aligned}$$

a. Find all the points that satisfy the KKT conditions.

b. Apply the SOSC to determine the nature of the critical points from the previous part.

21.2 Find local extremizers for:

a. $x_1^2 + x_2^2 - 2x_1 - 10x_2 + 26$ subject to $\frac{1}{5}x_2 - x_1^2 \leq 0$, $5x_1 + \frac{1}{2}x_2 \leq 5$.

b. $x_1^2 + x_2^2$ subject to $x_1 \geq 0$, $x_2 \geq 0$, $x_1 + x_2 \geq 5$.

c. $x_1^2 + 6x_1x_2 - 4x_1 - 2x_2$ subject to $x_1^2 + 2x_2 \leq 1$, $2x_1 - 2x_2 \leq 1$.

21.3 Find local minimizers for $x_1^2 + x_2^2$ subject to $x_1^2 + 2x_1x_2 + x_2^2 = 1$, $x_1^2 - x_2 \leq 0$.

21.4 Write down the Karush-Kuhn-Tucker condition for the optimization problem in Exercise 15.8.

21.5 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_2 - (x_1 - 2)^3 + 3 \\ \text{subject to} & x_2 \geq 1, \end{array}$$

where x_1 and x_2 are real variables. Answer each of the following questions, making sure that you give complete reasoning for your answers.

a. Write down the KKT condition for the problem, and find all points that satisfy the condition. Check whether or not each point is regular.

b. Determine whether or not the point(s) in part a satisfy the second-order necessary condition.

c. Determine whether or not the point(s) in part b satisfy the second-order sufficient condition.

21.6 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_2 \\ \text{subject to} & x_2 \geq -(x_1 - 1)^2 + 3. \end{array}$$

a. Find all points satisfying the KKT condition for the problem.

b. For each point $\boldsymbol{x}^*$ in part a, find $T(\boldsymbol{x}^*)$, $N(\boldsymbol{x}^*)$, and $\tilde{T}(\boldsymbol{x}^*)$.

c. Find the subset of points from part a that satisfy the second-order necessary condition.

21.7 Consider the problem of optimizing (either minimizing or maximizing) $(x_1 - 2)^2 + (x_2 - 1)^2$ subject to

$$\begin{aligned} x_2 - x_1^2 &\geq 0 \\ 2 - x_1 - x_2 &\geq 0 \\ x_1 &\geq 0. \end{aligned}$$

The point $\boldsymbol{x}^* = \mathbf{0}$ satisfies the KKT conditions.

a. Does $\boldsymbol{x}^*$ satisfy the FONC for minimization or maximization? What are the KKT multipliers?

b. Does $\boldsymbol{x}^*$ satisfy the SOSC? Carefully justify your answer.

21.8 Consider the optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $f(\boldsymbol{x}) = x_1 x_2^2$, where $\boldsymbol{x} = [x_1, x_2]^\top$, and $\Omega = \{\boldsymbol{x} \in \mathbb{R}^2 : x_1 = x_2,\ x_1 \geq 0\}$.

a. Find all points satisfying the KKT condition.

b. Do each of the points found in part a satisfy the second-order necessary condition?

c. Do each of the points found in part a satisfy the second-order sufficient condition?

21.9 Consider the problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 \\ \text{subject to} & x_1 + \cdots + x_n = 1 \\ & x_1, \ldots, x_n \geq 0. \end{array}$$

a. Write down the KKT condition for the problem.

b. Define what it means for a feasible point $\boldsymbol{x}^*$ to be *regular* in this particular problem. Are there any feasible points in this problem that are not regular? If yes, find them. If not, explain why not.

21.10 Let $g : \mathbb{R}^n \to \mathbb{R}$ and $\boldsymbol{x}_0 \in \mathbb{R}^n$ be given, where $g(\boldsymbol{x}_0) > 0$. Consider the problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|\boldsymbol{x} - \boldsymbol{x}_0\|^2 \\ \text{subject to} & g(\boldsymbol{x}) \leq 0. \end{array}$$

Suppose that $\boldsymbol{x}^*$ is a solution to the problem and $g \in \mathcal{C}^1$. Use the KKT theorem to decide which of the following equations/inequalities hold:

i. $g(\boldsymbol{x}^*) < 0$.

ii. $g(\boldsymbol{x}^*) = 0$.

iii. $(\boldsymbol{x}^* - \boldsymbol{x}_0)^\top \nabla g(\boldsymbol{x}^*) < 0$.

iv. $(\boldsymbol{x}^* - \boldsymbol{x}_0)^\top \nabla g(\boldsymbol{x}^*) = 0$.

v. $(\boldsymbol{x}^* - \boldsymbol{x}_0)^\top \nabla g(\boldsymbol{x}^*) > 0$.

21.11 Consider a square room with corners located at $[0,0]^\top$, $[0,2]^\top$, $[2,0]^\top$, and $[2,2]^\top$ (in $\mathbb{R}^2$). We wish to find the point in the room that is closest to the point $[3,4]^\top$.

a. Guess which point in the room is the closest point in the room to the point $[3,4]^\top$.

b. Use the second-order sufficient conditions to prove that the point you have guessed is a strict local minimizer.

Hint: Minimizing the distance is the same as minimizing the square distance.

21.12 Consider the *quadratic programming problem*

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}, \end{array}$$

where $\boldsymbol{Q} = \boldsymbol{Q}^\top > 0$, $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, and $\boldsymbol{b} \geq \boldsymbol{0}$. Find all points satisfying the KKT condition.

21.13 Consider the linear programming problem

$$\begin{array}{ll} \text{minimize} & ax_1 + bx_2 \\ \text{subject to} & cx_1 + dx_2 = e \\ & x_1, x_2 \geq 0, \end{array}$$

where $a, b, c, d, e \in \mathbb{R}$ are all nonzero constants. Suppose that $\boldsymbol{x}^*$ is an optimal basic feasible solution to the problem.

a. Write down the Karush-Kuhn-Tucker condition involving $\boldsymbol{x}^*$ (specifying clearly the number of Lagrange and KKT multipliers).

b. Is $\boldsymbol{x}^*$ regular? Explain.

c. Find the tangent space $T(\boldsymbol{x}^*)$ (defined by the active constraints) for this problem.

d. Assume that the relative cost coefficients of all nonbasic variables are strictly positive. Does $\boldsymbol{x}^*$ satisfy the second-order sufficient condition? Explain.

21.14 Consider the problem

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \le \boldsymbol{0}, \end{aligned}$$

where $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, $m < n$, is of full rank. Use the KKT theorem to show that if there exists a solution, then the optimal objective function value is 0.

21.15 Consider a linear programming problem in standard form (see Chapter 15).

a. Write down the Karush-Kuhn-Tucker condition for the problem.

b. Use part a to show that if there exists an optimal feasible solution to the linear program, then there exists a feasible solution to the corresponding dual problem that achieves an objective function value that is the same as the optimal value of the primal (compare this with Theorem 17.1).

c. Use parts a and b to prove that if $\boldsymbol{x}^*$ is an optimal feasible solutions of the primal, then there exists a feasible solution $\boldsymbol{\lambda}^*$ to the dual such that $(\boldsymbol{c}^\top - \boldsymbol{\lambda}^{*\top}\boldsymbol{A})\boldsymbol{x}^* = 0$ (compare this with Theorem 17.3).

21.16 Consider the constraint set $S = \{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}, \boldsymbol{g}(\boldsymbol{x}) \le \boldsymbol{0}\}$. Let $\boldsymbol{x}^* \in S$ be a regular local minimizer of f over S and $J(\boldsymbol{x}^*)$ the index set of active inequality constraints. Show that $\boldsymbol{x}^*$ is also a regular local minimizer of f over the set $S' = \{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}, g_j(\boldsymbol{x}) = 0, j \in J(\boldsymbol{x}^*)\}$.

21.17 Solve the following optimization problem using the second-order sufficient conditions:

$$\begin{aligned} \text{minimize} \quad & x_1^2 + x_2^2 \\ \text{subject to} \quad & x_1^2 - x_2 - 4 \le 0 \\ & x_2 - x_1 - 2 \le 0. \end{aligned}$$

See Figure 22.1 for a graphical illustration of the problem.

21.18 Solve the following optimization problem using the second-order sufficient conditions:

$$\begin{array}{ll} \text{minimize} & x_1^2 + x_2^2 \\ \text{subject to} & x_1 - x_2^2 - 4 \geq 0 \\ & x_1 - 10 \leq 0. \end{array}$$

See Figure 22.2 for a graphical illustration of the problem.

21.19 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_1^2 + x_2^2 \\ \text{subject to} & 4 - x_1 - x_2^2 \leq 0 \\ & 3x_2 - x_1 \leq 0 \\ & -3x_2 - x_1 \leq 0. \end{array}$$

Figure 22.3 gives a graphical illustration of the problem. Deduce from the figure that the problem has two strict local minimizers, and use the second-order sufficient conditions to verify the graphical solutions.

21.20 Consider the following optimization problem with an inequality constraint:

$$\begin{array}{ll} \text{minimize} & 3x_1 \\ \text{subject to} & x_1 + x_2^2 \geq 2. \end{array}$$

a. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the KKT (first-order necessary) condition?

b. Does the point $\boldsymbol{x}^* = [2, 0]^\top$ satisfy the second-order necessary condition (for problems with inequality constraints)?

c. Is the point $\boldsymbol{x}^* = [2, 0]^\top$ a local minimizer?

(See Exercise 6.15 for a similar problem treated using set-constrained methods.)

21.21 Consider the problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|\boldsymbol{x}\|^2 \\ \text{subject to} & \boldsymbol{a}^\top \boldsymbol{x} = b \\ & \boldsymbol{x} \geq \boldsymbol{0}, \end{array}$$

where $\boldsymbol{a} \in \mathbb{R}^n$, $\boldsymbol{a} \geq \boldsymbol{0}$, and $b \in \mathbb{R}$, $b > 0$. Show that if a solution to the problem exists, then it is unique, and find an expression for it in terms of $\boldsymbol{a}$ and b.

21.22 Consider the problem

$$\begin{array}{ll} \text{minimize} & (x_1 - a)^2 + (x_2 - b)^2, \qquad x_1, x_2 \in \mathbb{R} \\ \text{subject to} & x_1^2 + x_2^2 \leq 1, \end{array}$$

where $a, b \in \mathbb{R}$ are given constants satisfying $a^2 + b^2 \geq 1$.

a. Let $\boldsymbol{x}^* = [x_1^*, x_2^*]^\top$ be a solution to the problem. Use the first-order necessary conditions for unconstrained optimization to show that $(x_1^*)^2 + (x_2^*)^2 = 1$.

b. Use the KKT theorem to show that the solution $\boldsymbol{x}^* = [x_1^*, x_2^*]^\top$ is unique and has the form $x_1^* = \alpha a$, $x_2^* = \alpha b$, where $\alpha \in \mathbb{R}$ is a positive constant.

c. Find an expression for α (from part b) in terms of a and b.

21.23 Consider the problem

$$\begin{array}{ll} \text{minimize} & x_1^2 + (x_2 + 1)^2, \qquad x_1, x_2 \in \mathbb{R} \\ \text{subject to} & x_2 \geq \exp(x_1) \end{array}$$

[$\exp(x) = e^x$ is the exponential of x]. Let $\boldsymbol{x}^* = [x_1^*, x_2^*]^\top$ be the solution to the problem.

a. Write down the KKT condition that must be satisfied by $\boldsymbol{x}^*$.

b. Prove that $x_2^* = \exp(x_1^*)$.

c. Prove that $-2 < x_1^* < 0$.

21.24 Consider the problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} + 8 \\ \text{subject to} & \frac{1}{2}\|\boldsymbol{x}\|^2 \leq 1, \end{array}$$

where $\boldsymbol{c} \in \mathbb{R}^n$, $\boldsymbol{c} \neq \boldsymbol{0}$. Suppose that $\boldsymbol{x}^* = \alpha \boldsymbol{e}$ is a solution to the problem, where $\alpha \in \mathbb{R}$ and $\boldsymbol{e} = [1, \ldots, 1]^\top$, and the corresponding objective value is 4.

a. Show that $\|\boldsymbol{x}^*\|^2 = 2$.

b. Find α and $\boldsymbol{c}$ (they may depend on n).

21.25 Consider the problem with equality constraint

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}. \end{array}$$

We can convert the above into the equivalent optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \dfrac{1}{2}\|\boldsymbol{h}(\boldsymbol{x})\|^2 \leq 0. \end{array}$$

Write down the KKT condition for the equivalent problem (with inequality constraint) and explain why the KKT theorem cannot be applied in this case.

CHAPTER 22

CONVEX OPTIMIZATION PROBLEMS

22.1 Introduction

The optimization problems posed at the beginning of this part are, in general, very difficult to solve. The source of these difficulties may be in the objective function or the constraints. Even if the objective function is simple and "well-behaved," the nature of the constraints may make the problem difficult to solve. We illustrate some of these difficulties in the following examples.

Example 22.1 Consider the optimization problem

$$\begin{array}{rl} \text{minimize} & x_1^2 + x_2^2 \\ \text{subject to} & x_2 - x_1 - 2 \leq 0 \\ & x_1^2 - x_2 - 4 \leq 0. \end{array}$$

The problem is depicted in Figure 22.1, where, as we can see, the constrained minimizer is the same as the unconstrained minimizer. At the minimizer, all the constraints are inactive. If we had only known this fact, we could have approached this problem as an unconstrained optimization problem using techniques from Part II. ∎

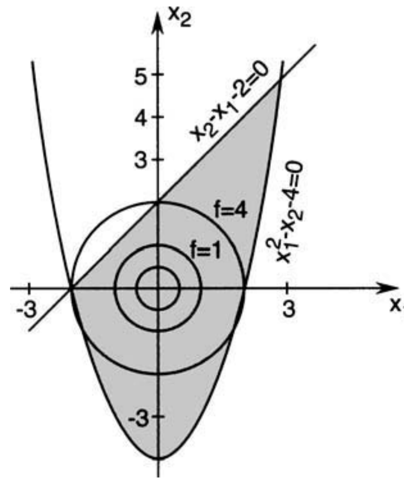

Figure 22.1 Situation where the constrained and the unconstrained minimizers are the same.

Example 22.2 Consider the optimization problem

$$\begin{aligned} \text{minimize} \quad & x_1^2 + x_2^2 \\ \text{subject to} \quad & x_1 - 10 \leq 0 \\ & x_1 - x_2^2 - 4 \geq 0. \end{aligned}$$

The problem is depicted in Figure 22.2. At the solution, only one constraint is active. If we had only known about this we could have handled this problem as a constrained optimization problem using the Lagrange multiplier method. ■

Example 22.3 Consider the optimization problem

$$\begin{aligned} \text{minimize} \quad & x_1^2 + x_2^2 \\ \text{subject to} \quad & 4 - x_1 - x_2^2 \leq 0 \\ & 3x_2 - x_1 \leq 0 \\ & -3x_2 - x_1 \leq 0. \end{aligned}$$

The problem is depicted in Figure 22.3. This example illustrates the situation where the constraints introduce local minimizers, even though the objective function itself has only one unconstrained global minimizer. ■

Some of the difficulties illustrated in the examples above can be eliminated if we restrict our problems to convex feasible regions. Admittedly, some important real-life problems do not fit into this framework. On the other hand,

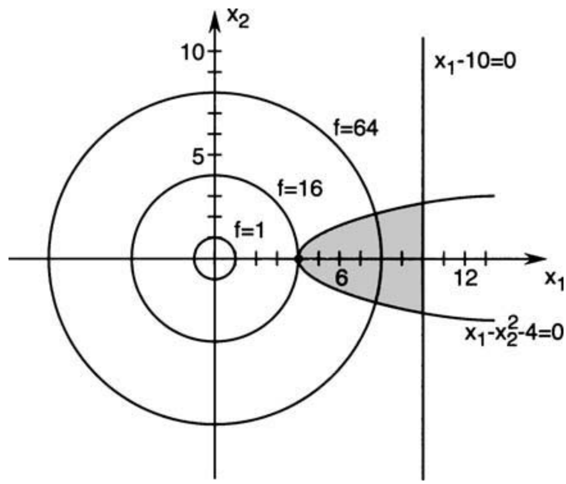

Figure 22.2 Situation where only one constraint is active.

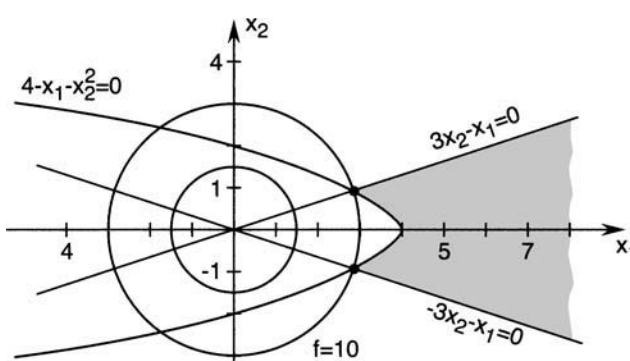

Figure 22.3 Situation where the constraints introduce local minimizers.

it is possible to give results of a *global* nature for this class of optimization problems. In the next section, we introduce the notion of a *convex function*, which plays an important role in our subsequent treatment of such problems.

22.2 Convex Functions

We begin with a definition of the graph of a real-valued function.

Definition 22.1 The *graph* of $f : \Omega \to \mathbb{R}$, $\Omega \subset \mathbb{R}^n$, is the set of points in $\Omega \times \mathbb{R} \subset \mathbb{R}^{n+1}$ given by

$$\left\{ \begin{bmatrix} \boldsymbol{x} \\ f(\boldsymbol{x}) \end{bmatrix} : \boldsymbol{x} \in \Omega \right\}.$$

■

We can visualize the graph of f as simply the set of points on a "plot" of $f(\boldsymbol{x})$ versus $\boldsymbol{x}$ (see Figure 22.4). We next define the epigraph of a real-valued function.

Definition 22.2 The *epigraph* of a function $f : \Omega \to \mathbb{R}$, $\Omega \subset \mathbb{R}^n$, denoted $\operatorname{epi}(f)$, is the set of points in $\Omega \times \mathbb{R}$ given by

$$\operatorname{epi}(f) = \left\{ \begin{bmatrix} \boldsymbol{x} \\ \beta \end{bmatrix} : \boldsymbol{x} \in \Omega,\ \beta \in \mathbb{R},\ \beta \geq f(\boldsymbol{x}) \right\}.$$

■

The epigraph $\operatorname{epi}(f)$ of a function f is simply the set of points in $\Omega \times \mathbb{R}$ on or above the graph of f (see Figure 22.4). We can also think of $\operatorname{epi}(f)$ as a subset of $\mathbb{R}^{n+1}$.

Recall that a set $\Omega \subset \mathbb{R}^n$ is convex if for every $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \Omega$ and $\alpha \in (0,1)$, $\alpha\boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in \Omega$ (see Section 4.3). We now introduce the notion of a convex function.

Definition 22.3 A function $f : \Omega \to \mathbb{R}$, $\Omega \subset \mathbb{R}^n$, is *convex* on Ω if its epigraph is a convex set. ■

Theorem 22.1 *If a function $f : \Omega \to \mathbb{R}$, $\Omega \subset \mathbb{R}^n$, is convex on Ω, then Ω is a convex set.* □

Proof. We prove this theorem by contraposition. Suppose that Ω is not a convex set. Then, there exist two points $\boldsymbol{y}_1$ and $\boldsymbol{y}_2$ such that for some $\alpha \in (0,1)$,

$$\boldsymbol{z} = \alpha\boldsymbol{y}_1 + (1-\alpha)\boldsymbol{y}_2 \notin \Omega.$$

Let

$$\beta_1 = f(\boldsymbol{y}_1), \quad \beta_2 = f(\boldsymbol{y}_2).$$

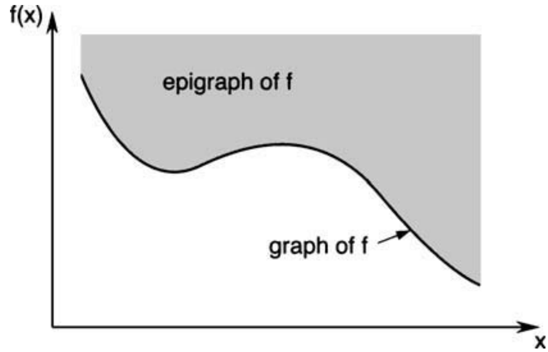

Figure 22.4 Graph and epigraph of a function $f : \mathbb{R} \to \mathbb{R}$.

Then, the pairs

$$\begin{bmatrix} \boldsymbol{y}_1 \\ \beta_1 \end{bmatrix}, \quad \begin{bmatrix} \boldsymbol{y}_2 \\ \beta_2 \end{bmatrix}$$

belong to the graph of f, and hence also the epigraph of f. Let

$$\boldsymbol{w} = \alpha \begin{bmatrix} \boldsymbol{y}_1 \\ \beta_1 \end{bmatrix} + (1-\alpha) \begin{bmatrix} \boldsymbol{y}_2 \\ \beta_2 \end{bmatrix}.$$

We have

$$\boldsymbol{w} = \begin{bmatrix} \boldsymbol{z} \\ \alpha\beta_1 + (1-\alpha)\beta_2 \end{bmatrix}.$$

But note that $\boldsymbol{w} \notin \text{epi}(f)$, because $\boldsymbol{z} \notin \Omega$. Therefore, epi$(f)$ is not convex, and hence f is not a convex function. ∎

The next theorem gives a very useful characterization of convex functions. This characterization is often used as a definition for a convex function.

Theorem 22.2 *A function $f : \Omega \to \mathbb{R}$ defined on a convex set $\Omega \subset \mathbb{R}^n$ is convex if and only if for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and all $\alpha \in (0,1)$, we have*

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

□

Proof. $\Leftarrow$: Assume that for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and $\alpha \in (0,1)$,

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

Let $[\boldsymbol{x}^\top, a]^\top$ and $[\boldsymbol{y}^\top, b]^\top$ be two points in epi(f), where $a, b \in \mathbb{R}$. From the definition of epi(f) it follows that

$$f(\boldsymbol{x}) \leq a, \quad f(\boldsymbol{y}) \leq b.$$

Therefore, using the first inequality above, we have

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

Because Ω is convex, $\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y} \in \Omega$. Hence,

$$\begin{bmatrix} \alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y} \\ \alpha a + (1-\alpha)b \end{bmatrix} \in \operatorname{epi}(f),$$

which implies that $\operatorname{epi}(f)$ is a convex set, and hence f is a convex function.

$\Rightarrow$: Assume that $f : \Omega \to \mathbb{R}$ is a convex function. Let $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and

$$f(\boldsymbol{x}) = a, \qquad f(\boldsymbol{y}) = b.$$

Thus,

$$\begin{bmatrix} \boldsymbol{x} \\ a \end{bmatrix}, \begin{bmatrix} \boldsymbol{y} \\ b \end{bmatrix} \in \operatorname{epi}(f).$$

Because f is a convex function, its epigraph is a convex subset of $\mathbb{R}^{n+1}$. Therefore, for all $\alpha \in (0,1)$, we have

$$\alpha \begin{bmatrix} \boldsymbol{x} \\ a \end{bmatrix} + (1-\alpha) \begin{bmatrix} \boldsymbol{y} \\ b \end{bmatrix} = \begin{bmatrix} \alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y} \\ \alpha a + (1-\alpha)b \end{bmatrix} \in \operatorname{epi}(f).$$

The above implies that for all $\alpha \in (0,1)$,

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

This completes the proof. ∎

A geometric interpretation of Theorem 22.2 is given in Figure 22.5. The theorem states that if $f : \Omega \to \mathbb{R}$ is a convex function over a convex set Ω, then for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, the points on the line segment in $\mathbb{R}^{n+1}$ connecting $[\boldsymbol{x}^\top, f(\boldsymbol{x})]^\top$ and $[\boldsymbol{y}^\top, f(\boldsymbol{y})]^\top$ must lie on or above the graph of f.

Using Theorem 22.2, it is straightforward to show that any nonnegative scaling of a convex function is convex, and that the sum of convex functions is convex.

Theorem 22.3 *Suppose that f, f_1, and f_2 are convex functions. Then, for any $a \geq 0$, the function af is convex. Moreover, $f_1 + f_2$ is convex.* □

Proof. Let $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and $\alpha \in (0,1)$. Fix $a \geq 0$. For convenience, write $\bar{f} = af$. We have

$$\begin{aligned} \bar{f}(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) &= af(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) \\ &\leq a\left(\alpha f(\boldsymbol{x}) + (1-\alpha)f(\boldsymbol{y})\right) \text{ because } f \text{ is convex and } a \geq 0 \\ &= \alpha(af(\boldsymbol{x})) + (1-\alpha)(af(\boldsymbol{y})) \\ &= \alpha \bar{f}(\boldsymbol{x}) + (1-\alpha)\bar{f}(\boldsymbol{y}), \end{aligned}$$

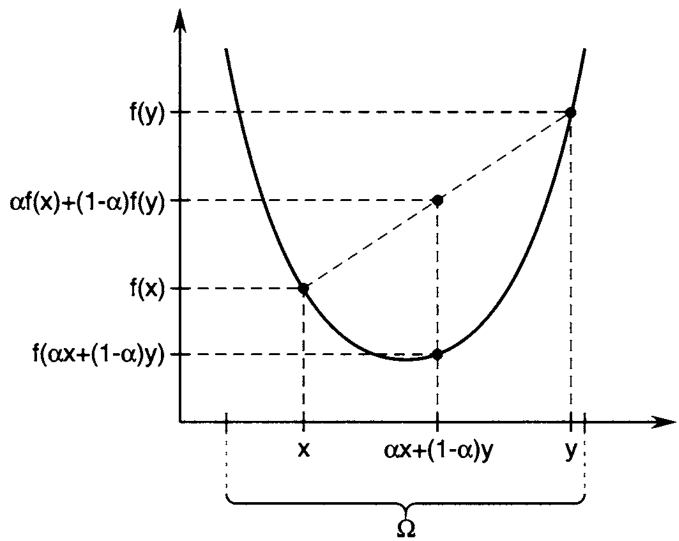

Figure 22.5 Geometric interpretation of Theorem 22.2.

which implies that $\bar{f}$ is convex.

Next, write $f_3 = f_1 + f_2$. We have

$$\begin{aligned} f_3(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) &= f_1(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) + f_2(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) \\ &\leq (\alpha f_1(\boldsymbol{x}) + (1-\alpha) f_1(\boldsymbol{y})) + (\alpha f_2(\boldsymbol{x}) + (1-\alpha) f_2(\boldsymbol{y})) \\ &\quad \text{by convexity of } f_1 \text{ and } f_2 \\ &= \alpha(f_1(\boldsymbol{x}) + f_2(\boldsymbol{x})) + (1-\alpha)(f_1(\boldsymbol{y}) + f_2(\boldsymbol{y})) \\ &= \alpha f_3(\boldsymbol{x}) + (1-\alpha) f_3(\boldsymbol{y}), \end{aligned}$$

which implies that f_3 is convex. ∎

Theorem 22.3 implies that for any given collection of convex functions $f_1, \ldots, f_\ell$ and nonnegative numbers $c_1, \ldots, c_\ell$, the function $c_1 f_2 + \cdots + c_\ell f_\ell$ is convex. Using a method of proof similar to that used in Theorem 22.3, it is similarly straightforward to show that the function $\max\{f_1, \ldots, f_\ell\}$ is convex (see Exercise 22.6).

We now define the notion of strict convexity.

Definition 22.4 A function $f : \Omega \to \mathbb{R}$ on a convex set $\Omega \subset \mathbb{R}^n$ is *strictly convex* if for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, $\boldsymbol{x} \neq \boldsymbol{y}$, and $\alpha \in (0,1)$, we have

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

∎

From this definition, we see that for a strictly convex function, all points on the open line segment connecting the points $[\boldsymbol{x}^\top, f(\boldsymbol{x})]^\top$ and $[\boldsymbol{y}^\top, f(\boldsymbol{y})]^\top$ lie (strictly) above the graph of f.

Definition 22.5 A function $f : \Omega \to \mathbb{R}$ on a convex set $\Omega \subset \mathbb{R}^n$ is (strictly) *concave* if $-f$ is (strictly) convex. ∎

Note that the graph of a strictly concave function always lies above the line segment connecting any two points on its graph.

To show that a function is not convex, we need only produce a pair of points $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and an $\alpha \in (0,1)$ such that the inequality in Theorem 22.2 is violated.

Example 22.4 Let $f(\boldsymbol{x}) = x_1 x_2$. Is f convex over $\Omega = \{\boldsymbol{x} : x_1 \geq 0,\ x_2 \geq 0\}$?

The answer is no. Take, for example, $\boldsymbol{x} = [1,2]^\top \in \Omega$ and $\boldsymbol{y} = [2,1]^\top \in \Omega$. Then,

$$\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y} = \begin{bmatrix} 2-\alpha \\ 1+\alpha \end{bmatrix}.$$

Hence,

$$f(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) = (2-\alpha)(1+\alpha) = 2 + \alpha - \alpha^2$$

and

$$\alpha f(\boldsymbol{x}) + (1-\alpha) f(\boldsymbol{y}) = 2.$$

If, for example, $\alpha = 1/2 \in (0,1)$, then

$$f\left(\frac{1}{2}\boldsymbol{x} + \frac{1}{2}\boldsymbol{y}\right) = \frac{9}{4} > \frac{1}{2} f(\boldsymbol{x}) + \frac{1}{2} f(\boldsymbol{y}),$$

which shows that f is not convex over Ω. ∎

Example 22.4 is an illustration of the following general result.

Proposition 22.1 *A quadratic form $f : \Omega \to \mathbb{R}$, $\Omega \subset \mathbb{R}^n$, given by $f(\boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{Q} \boldsymbol{x}$, $\boldsymbol{Q} \in \mathbb{R}^{n\times n}$, $\boldsymbol{Q} = \boldsymbol{Q}^\top$, is convex on Ω if and only if for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, $(\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{Q} (\boldsymbol{x}-\boldsymbol{y}) \geq 0$.* □

Proof. The result follows from Theorem 22.2. Indeed, the function $f(\boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{Q} \boldsymbol{x}$ is convex if and only if for every $\alpha \in (0,1)$, and every $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, we have

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

or, equivalently,

$$\alpha f(\boldsymbol{x}) + (1-\alpha) f(\boldsymbol{y}) - f(\alpha \boldsymbol{x} + (1-\alpha)\boldsymbol{y}) \geq 0.$$

Substituting for f into the left-hand side of this equation yields

$$\begin{aligned}
&\alpha \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + (1-\alpha)\boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} - (\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y})^\top \boldsymbol{Q}(\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y}) \\
&\qquad = \alpha \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} + \boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} - \alpha \boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} - \alpha^2 \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \\
&\qquad\quad - (2\alpha - 2\alpha^2)\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{y} - (1 - 2\alpha + \alpha^2)\boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} \\
&\qquad = \alpha(1-\alpha)\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - 2\alpha(1-\alpha)\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{y} + \alpha(1-\alpha)\boldsymbol{y}^\top \boldsymbol{Q}\boldsymbol{y} \\
&\qquad = \alpha(1-\alpha)(\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{Q}(\boldsymbol{x}-\boldsymbol{y}).
\end{aligned}$$

Therefore, f is convex if and only if

$$\alpha(1-\alpha)(\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{Q}(\boldsymbol{x}-\boldsymbol{y}) \geq 0,$$

which proves the result. ∎

Example 22.5 In Example 22.4, $f(\boldsymbol{x}) = x_1 x_2$, which can be written as $f(\boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x}$, where

$$\boldsymbol{Q} = \frac{1}{2}\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Let $\Omega = \{\boldsymbol{x} : \boldsymbol{x} \geq \boldsymbol{0}\}$, and $\boldsymbol{x} = [2,2]^\top \in \Omega$, $\boldsymbol{y} = [1,3]^\top \in \Omega$. We have

$$\boldsymbol{y} - \boldsymbol{x} = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

and

$$(\boldsymbol{y}-\boldsymbol{x})^\top \boldsymbol{Q}(\boldsymbol{y}-\boldsymbol{x}) = \frac{1}{2}[-1,1]\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}\begin{bmatrix} -1 \\ 1 \end{bmatrix} = -1 < 0.$$

Hence, by Proposition 22.1, f is not convex on Ω. ∎

Differentiable convex functions can be characterized using the following theorem.

Theorem 22.4 *Let $f : \Omega \to \mathbb{R}$, $f \in \mathcal{C}^1$, be defined on an open convex set $\Omega \subset \mathbb{R}^n$. Then, f is convex on Ω if and only if for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$,*

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}).$$

□

Proof. $\Rightarrow$: Suppose that $f : \Omega \to \mathbb{R}$ is differentiable and convex. Then, by Theorem 22.2, for any $\boldsymbol{y}, \boldsymbol{x} \in \Omega$ and $\alpha \in (0,1)$ we have

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

Rearranging terms yields

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

Upon dividing both sides of this inequality by α, we get

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

If we now take the limit as $\alpha \to 0$ and apply the definition of the directional derivative of f at $\boldsymbol{x}$ in the direction $\boldsymbol{y}-\boldsymbol{x}$ (see Section 6.2), we get

$$Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}) \leq f(\boldsymbol{y})-f(\boldsymbol{x})$$

or

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x})+Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}).$$

$\Leftarrow$: Assume that Ω is convex, $f : \Omega \to \mathbb{R}$ is differentiable, and for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$,

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x})+Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}).$$

Let $\boldsymbol{u}, \boldsymbol{v} \in \Omega$ and $\alpha \in (0,1)$. Because Ω is convex,

$$\boldsymbol{w}=\alpha\boldsymbol{u}+(1-\alpha)\boldsymbol{v} \in \Omega.$$

We also have

$$f(\boldsymbol{u}) \geq f(\boldsymbol{w})+Df(\boldsymbol{w})(\boldsymbol{u}-\boldsymbol{w})$$

and

$$f(\boldsymbol{v}) \geq f(\boldsymbol{w})+Df(\boldsymbol{w})(\boldsymbol{v}-\boldsymbol{w}).$$

Multiplying the first of this inequalities by α and the second by $(1-\alpha)$ and then adding them together yields

$$\alpha f(\boldsymbol{u})+(1-\alpha)f(\boldsymbol{v}) \geq f(\boldsymbol{w})+Df(\boldsymbol{w})\left(\alpha\boldsymbol{u}+(1-\alpha)\boldsymbol{v}-\boldsymbol{w}\right).$$

But

$$\boldsymbol{w}=\alpha\boldsymbol{u}+(1-\alpha)\boldsymbol{v}.$$

Hence,

$$\alpha f(\boldsymbol{u})+(1-\alpha)f(\boldsymbol{v}) \geq f(\alpha\boldsymbol{u}+(1-\alpha)\boldsymbol{v}).$$

Hence, by Theorem 22.2, f is a convex function. ∎

In Theorem 22.4, the assumption that Ω be open is not necessary, as long as $f \in \mathcal{C}^1$ on some open set that contains Ω (e.g., $f \in \mathcal{C}^1$ on $\mathbb{R}^n$).

A geometric interpretation of Theorem 22.4 is given in Figure 22.6. To explain the interpretation, let $\boldsymbol{x}_0 \in \Omega$. The function $\ell(\boldsymbol{x})=f(\boldsymbol{x}_0)+Df(\boldsymbol{x}_0)(\boldsymbol{x}-\boldsymbol{x}_0)$ is the linear approximation to f at $\boldsymbol{x}_0$. The theorem says that the graph of f always lies above its linear approximation at any point. In other words,

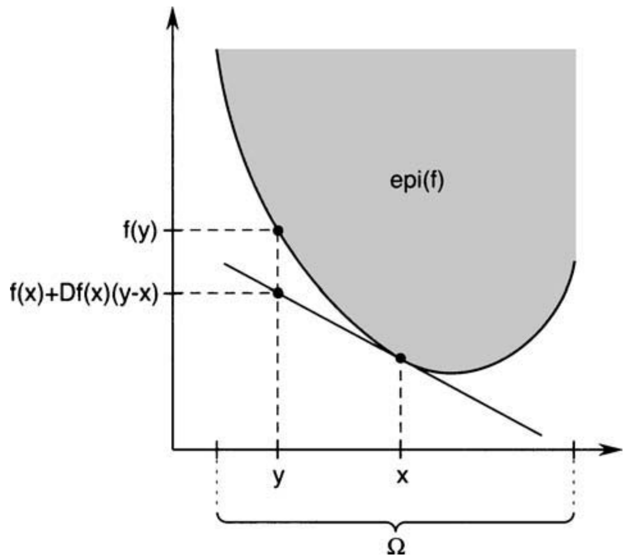

Figure 22.6 Geometric interpretation of Theorem 22.4.

the linear approximation to a convex function f at any point of its domain lies below epi(f).

This geometric idea leads to a generalization of the gradient to the case where f is not differentiable. Let $f : \Omega \to \mathbb{R}$ be defined on an open convex set $\Omega \subset \mathbb{R}^n$. A vector $\boldsymbol{g} \in \mathbb{R}^n$ is said to be a *subgradient* of f at a point $\boldsymbol{x} \in \Omega$ if for all $\boldsymbol{y} \in \Omega$,

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + \boldsymbol{g}^\top(\boldsymbol{y} - \boldsymbol{x}).$$

As in the case of the standard gradient, if $\boldsymbol{g}$ is a subgradient, then for a given $\boldsymbol{x}_0 \in \Omega$, the function $\ell(\boldsymbol{x}) = f(\boldsymbol{x}_0) + \boldsymbol{g}^\top(\boldsymbol{x} - \boldsymbol{x}_0)$ lies below epi(f).

For functions that are twice continuously differentiable, the following theorem gives another possible characterization of convexity.

Theorem 22.5 *Let $f : \Omega \to \mathbb{R}$, $f \in \mathcal{C}^2$, be defined on an open convex set $\Omega \subset \mathbb{R}^n$. Then, f is convex on Ω if and only if for each $\boldsymbol{x} \in \Omega$, the Hessian $\boldsymbol{F}(\boldsymbol{x})$ of f at $\boldsymbol{x}$ is a positive semidefinite matrix.* □

Proof. $\Leftarrow$: Let $\boldsymbol{x}, \boldsymbol{y} \in \Omega$. Because $f \in \mathcal{C}^2$, by Taylor's theorem there exists $\alpha \in (0, 1)$ such that

$$f(\boldsymbol{y}) = f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x}) + \frac{1}{2}(\boldsymbol{y} - \boldsymbol{x})^\top \boldsymbol{F}(\boldsymbol{x} + \alpha(\boldsymbol{y} - \boldsymbol{x}))(\boldsymbol{y} - \boldsymbol{x}).$$

Because $\boldsymbol{F}(\boldsymbol{x} + \alpha(\boldsymbol{y} - \boldsymbol{x}))$ is positive semidefinite,

$$(\boldsymbol{y} - \boldsymbol{x})^\top \boldsymbol{F}(\alpha \boldsymbol{y} + (1 - \alpha)\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x}) \geq 0.$$

Therefore, we have

$$f(\boldsymbol{y}) \geq f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}),$$

which implies that f is convex, by Theorem 22.4.

$\Rightarrow$: We use contraposition. Assume that there exists $\boldsymbol{x} \in \Omega$ such that $\boldsymbol{F}(\boldsymbol{x})$ is not positive semidefinite. Therefore, there exists $\boldsymbol{d} \in \mathbb{R}^n$ such that $\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x})\boldsymbol{d} < 0$. By assumption, Ω is open; thus, the point $\boldsymbol{x}$ is an interior point. By the continuity of the Hessian matrix, there exists a nonzero $s \in \mathbb{R}$ such that $\boldsymbol{x} + s\boldsymbol{d} \in \Omega$, and if we write $\boldsymbol{y} = \boldsymbol{x} + s\boldsymbol{d}$, then for all points $\boldsymbol{z}$ on the line segment joining $\boldsymbol{x}$ and $\boldsymbol{y}$, we have $\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{z})\boldsymbol{d} < 0$. By Taylor's theorem there exists $\alpha \in (0,1)$ such that

$$\begin{aligned} f(\boldsymbol{y}) &= f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}) + \frac{1}{2}(\boldsymbol{y}-\boldsymbol{x})^\top \boldsymbol{F}(\boldsymbol{x}+\alpha(\boldsymbol{y}-\boldsymbol{x}))(\boldsymbol{y}-\boldsymbol{x}) \\ &= f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}) + \frac{1}{2}s^2\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}+\alpha s\boldsymbol{d})\boldsymbol{d}. \end{aligned}$$

Because $\alpha \in (0,1)$, the point $\boldsymbol{x} + \alpha s\boldsymbol{d}$ is on the line segment joining $\boldsymbol{x}$ and $\boldsymbol{y}$, and therefore

$$\boldsymbol{d}^\top \boldsymbol{F}(\boldsymbol{x}+\alpha s\boldsymbol{d})\boldsymbol{d} < 0.$$

Because $s \neq 0$, we have $s^2 > 0$, and hence

$$f(\boldsymbol{y}) < f(\boldsymbol{x}) + Df(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}).$$

Therefore, by Theorem 22.4, f is not a convex function. ∎

Theorem 22.5 can be strengthened to include nonopen sets by modifying the condition to be $(\boldsymbol{y}-\boldsymbol{x})^\top \boldsymbol{F}(\boldsymbol{x})(\boldsymbol{y}-\boldsymbol{x}) \geq 0$ for all $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ (and assuming that $f \in \mathcal{C}^2$ on some open set that contains Ω; for example, $f \in \mathcal{C}^2$ on $\mathbb{R}^n$). A proof similar to that above can be used in this case.

Note that by definition of concavity, a function $f : \Omega \to \mathbb{R}$, $f \in \mathcal{C}^2$, is concave over the convex set $\Omega \subset \mathbb{R}^n$ if and only if for all $\boldsymbol{x} \in \Omega$, the Hessian $\boldsymbol{F}(\boldsymbol{x})$ of f is negative semidefinite.

Example 22.6 Determine whether the following functions are convex, concave, or neither:

1. $f : \mathbb{R} \to \mathbb{R}$, $f(x) = -8x^2$.
2. $f : \mathbb{R}^3 \to \mathbb{R}$, $f(\boldsymbol{x}) = 4x_1^2 + 3x_2^2 + 5x_3^2 + 6x_1x_2 + x_1x_3 - 3x_1 - 2x_2 + 15$.
3. $f : \mathbb{R}^2 \to \mathbb{R}$, $f(\boldsymbol{x}) = 2x_1x_2 - x_1^2 - x_2^2$.

Solution:

1. We use Theorem 22.5. We first compute the Hessian, which in this case is just the second derivative: $(d^2f/dx^2)(x) = -16 < 0$ for all $x \in \mathbb{R}$. Hence, f is concave over $\mathbb{R}$.

2. The Hessian matrix of f is

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} 8 & 6 & 1 \\ 6 & 6 & 0 \\ 1 & 0 & 10 \end{bmatrix}.$$

The leading principal minors of $\boldsymbol{F}(\boldsymbol{x})$ are

$$\begin{aligned} \Delta_1 &= 8 > 0, \\ \Delta_2 &= \det \begin{bmatrix} 8 & 6 \\ 6 & 6 \end{bmatrix} = 12 > 0, \\ \Delta_3 &= \det \boldsymbol{F}(\boldsymbol{x}) = 114 > 0. \end{aligned}$$

Hence, $\boldsymbol{F}(\boldsymbol{x})$ is positive definite for all $\boldsymbol{x} \in \mathbb{R}^3$. Therefore, f is a convex function over $\mathbb{R}^3$.

3. The Hessian of f is

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} -2 & 2 \\ 2 & -2 \end{bmatrix},$$

which is negative semidefinite for all $\boldsymbol{x} \in \mathbb{R}^2$. Hence, f is concave on $\mathbb{R}^2$.

■

22.3 Convex Optimization Problems

In this section we consider optimization problems where the objective function is a convex function and the constraint set is a convex set. We refer to such problems as *convex optimization problems* or *convex programming problems*. Optimization problems that can be classified as convex programming problems include linear programs and optimization problems with quadratic objective function and linear constraints. Convex programming problems are interesting for several reasons. Specifically, as we shall see, local minimizers are global for such problems. Furthermore, first-order necessary conditions become sufficient conditions for minimization.

Our first theorem below states that in convex programming problems, local minimizers are also global.

Theorem 22.6 *Let $f : \Omega \to \mathbb{R}$ be a convex function defined on a convex set $\Omega \subset \mathbb{R}^n$. Then, a point is a global minimizer of f over Ω if and only if it is a local minimizer of f.* □

Proof. $\Rightarrow$: This is obvious.

$\Leftarrow$: We prove this by contraposition. Suppose that $\boldsymbol{x}^*$ is not a global minimizer of f over Ω. Then, for some $\boldsymbol{y} \in \Omega$, we have $f(\boldsymbol{y}) < f(\boldsymbol{x}^*)$. By assumption, the function f is convex, and hence for all $\alpha \in (0,1)$,

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

Because $f(\boldsymbol{y}) < f(\boldsymbol{x}^*)$, we have

$$\alpha f(\boldsymbol{y}) + (1-\alpha)f(\boldsymbol{x}^*) = \alpha(f(\boldsymbol{y}) - f(\boldsymbol{x}^*)) + f(\boldsymbol{x}^*) < f(\boldsymbol{x}^*).$$

Thus, for all $\alpha \in (0,1)$,

$$f(\alpha\boldsymbol{y} + (1-\alpha)\boldsymbol{x}^*) < f(\boldsymbol{x}^*).$$

Hence, there exist points that are arbitrarily close to $\boldsymbol{x}^*$ and have lower objective function value. For example, the sequence $\{\boldsymbol{y}_n\}$ of points given by

$$\boldsymbol{y}_n = \frac{1}{n}\boldsymbol{y} + \left(1 - \frac{1}{n}\right)\boldsymbol{x}^*$$

converges to $\boldsymbol{x}^*$, and $f(\boldsymbol{y}_n) < f(\boldsymbol{x}^*)$. Hence, $\boldsymbol{x}^*$ is not a local minimizer, which completes the proof. ∎

We now show that the set of global optimizers is convex. For this, we need the following lemma.

Lemma 22.1 *Let $g : \Omega \to \mathbb{R}$ be a convex function defined on a convex set $\Omega \subset \mathbb{R}^n$. Then, for each $c \in \mathbb{R}$, the set*

$$\Gamma_c = \{\boldsymbol{x} \in \Omega : g(\boldsymbol{x}) \leq c\}$$

is a convex set. □

Proof. Let $\boldsymbol{x}, \boldsymbol{y} \in \Gamma_c$. Then, $g(\boldsymbol{x}) \leq c$ and $g(\boldsymbol{y}) \leq c$. Because g is convex, for all $\alpha \in (0,1)$,

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

Hence, $\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y} \in \Gamma_c$, which implies that Γ_c is convex. ∎

Corollary 22.1 *Let $f : \Omega \to \mathbb{R}$ be a convex function defined on a convex set $\Omega \subset \mathbb{R}^n$. Then, the set of all global minimizers of f over Ω is a convex set.* □

Proof. The result follows immediately from Lemma 22.1 by setting

$$c = \min_{\boldsymbol{x} \in \Omega} f(\boldsymbol{x}).$$

∎

We now show that if the objective function is continuously differentiable and convex, then the first-order necessary condition (see Theorem 6.1) for a point to be a minimizer is also sufficient. We use the following lemma.

Lemma 22.2 *Let $f : \Omega \to \mathbb{R}$ be a convex function defined on the convex set $\Omega \subset \mathbb{R}^n$ and $f \in \mathcal{C}^1$ on an open convex set containing Ω. Suppose that the point $\boldsymbol{x}^* \in \Omega$ is such that for all $\boldsymbol{x} \in \Omega$, $\boldsymbol{x} \neq \boldsymbol{x}^*$, we have*

$$Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) \geq 0.$$

Then, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. □

Proof. Because the function f is convex, by Theorem 22.4, for all $\boldsymbol{x} \in \Omega$, we have

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*) + Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*).$$

Hence, the condition $Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) \geq 0$ implies that $f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*)$. ∎

Observe that for any $\boldsymbol{x} \in \Omega$, the vector $\boldsymbol{x} - \boldsymbol{x}^*$ can be interpreted as a feasible direction at $\boldsymbol{x}^*$ (see Definition 6.2). Using Lemma 22.2, we have the following theorem (cf. Theorem 6.1).

Theorem 22.7 *Let $f : \Omega \to \mathbb{R}$ be a convex function defined on the convex set $\Omega \subset \mathbb{R}^n$, and $f \in \mathcal{C}^1$ on an open convex set containing Ω. Suppose that the point $\boldsymbol{x}^* \in \Omega$ is such that for any feasible direction $\boldsymbol{d}$ at $\boldsymbol{x}^*$, we have*

$$\boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq 0.$$

Then, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. □

Proof. Let $\boldsymbol{x} \in \Omega$, $\boldsymbol{x} \neq \boldsymbol{x}^*$. By convexity of Ω,

$$\boldsymbol{x}^* + \alpha(\boldsymbol{x} - \boldsymbol{x}^*) = \alpha\boldsymbol{x} + (1 - \alpha)\boldsymbol{x}^* \in \Omega$$

for all $\alpha \in (0, 1)$. Hence, the vector $\boldsymbol{d} = \boldsymbol{x} - \boldsymbol{x}^*$ is a feasible direction at $\boldsymbol{x}^*$ (see Definition 6.2). By assumption,

$$Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) = \boldsymbol{d}^\top \nabla f(\boldsymbol{x}^*) \geq 0.$$

Hence, by Lemma 22.2, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. ∎

From Theorem 22.7, we easily deduce the following corollary (compare this with Corollary 6.1).

Corollary 22.2 *Let $f : \Omega \to \mathbb{R}$, $f \in \mathcal{C}^1$, be a convex function defined on the convex set $\Omega \subset \mathbb{R}^n$. Suppose that the point $\boldsymbol{x}^* \in \Omega$ is such that*

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

Then, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. □

We now consider the constrained optimization problem

$$\begin{array}{rl} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}. \end{array}$$

We assume that the feasible set is convex. An example where this is the case is when

$$\boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}.$$

The following theorem states that provided the feasible set is convex, the Lagrange condition is sufficient for a point to be a minimizer.

Theorem 22.8 *Let $f : \mathbb{R}^n \to \mathbb{R}$, $f \in \mathcal{C}^1$, be a convex function on the set of feasible points*

$$\Omega = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\},$$

where $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $\boldsymbol{h} \in \mathcal{C}^1$, and Ω is convex. Suppose that there exist $\boldsymbol{x}^* \in \Omega$ and $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ such that

$$Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) = \boldsymbol{0}^\top.$$

Then, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. □

Proof. By Theorem 22.4, for all $\boldsymbol{x} \in \Omega$, we have

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*) + Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*).$$

Substituting $Df(\boldsymbol{x}^*) = -\boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)$ into the inequality above yields

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*) - \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*).$$

Because Ω is convex, $(1-\alpha)\boldsymbol{x}^* + \alpha\boldsymbol{x} \in \Omega$ for all $\alpha \in (0,1)$. Thus,

$$\boldsymbol{h}(\boldsymbol{x}^* + \alpha(\boldsymbol{x} - \boldsymbol{x}^*)) = \boldsymbol{h}((1-\alpha)\boldsymbol{x}^* + \alpha\boldsymbol{x}) = \boldsymbol{0}$$

for all $\alpha \in (0,1)$. Premultiplying by $\boldsymbol{\lambda}^{*\top}$, subtracting $\boldsymbol{\lambda}^{*\top}\boldsymbol{h}(\boldsymbol{x}^*) = 0$, and dividing by α, we get

$$\frac{\boldsymbol{\lambda}^{*\top}\boldsymbol{h}(\boldsymbol{x}^* + \alpha(\boldsymbol{x} - \boldsymbol{x}^*)) - \boldsymbol{\lambda}^{*\top}\boldsymbol{h}(\boldsymbol{x}^*)}{\alpha} = 0$$

for all $\alpha \in (0,1)$. If we now take the limit as $\alpha \to 0$ and apply the definition of the directional derivative of $\boldsymbol{\lambda}^{*\top}\boldsymbol{h}$ at $\boldsymbol{x}^*$ in the direction $\boldsymbol{x} - \boldsymbol{x}^*$ (see Section 6.2), we get

$$\boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) = 0.$$

Hence,

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*),$$

which implies that $\boldsymbol{x}^*$ is a global minimizer of f over Ω. ∎

Consider the general constrained optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0} \\ & \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}. \end{array}$$

As before, we assume that the feasible set is convex. This is the case if, for example, the two sets $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\}$ and $\{\boldsymbol{x} : \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}\}$ are convex, because the feasible set is the intersection of these two sets (see also Theorem 4.1). We have already seen an example where the set $\{\boldsymbol{x} : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}\}$ is convex. On the other hand, an example where the set $\{\boldsymbol{x} : \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}\}$ is convex is when the components of $\boldsymbol{g} = [g_1, \ldots, g_p]^\top$ are all convex functions. Indeed, the set $\{\boldsymbol{x} : \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}\}$ is the intersection of the sets $\{\boldsymbol{x} : g_i(\boldsymbol{x}) \leq 0\}$, $i = 1, \ldots, p$. Because each of these sets is convex (by Lemma 22.1), their intersection is also convex.

We now prove that the Karush-Kuhn-Tucker (KKT) condition is sufficient for a point to be a minimizer to the problem above.

Theorem 22.9 *Let $f : \mathbb{R}^n \to \mathbb{R}$, $f \in \mathcal{C}^1$, be a convex function on the set of feasible points*

$$\Omega = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{h}(\boldsymbol{x}) = \boldsymbol{0}, \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{0}\},$$

where $\boldsymbol{h} : \mathbb{R}^n \to \mathbb{R}^m$, $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^p$, $\boldsymbol{h}, \boldsymbol{g} \in \mathcal{C}^1$, and Ω is convex. Suppose that there exist $\boldsymbol{x}^* \in \Omega$, $\boldsymbol{\lambda}^* \in \mathbb{R}^m$, and $\boldsymbol{\mu}^* \in \mathbb{R}^p$, such that

1. $\boldsymbol{\mu}^* \geq \boldsymbol{0}$.
2. $Df(\boldsymbol{x}^*) + \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*) + \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*) = \boldsymbol{0}^\top$.
3. $\boldsymbol{\mu}^{*\top} \boldsymbol{g}(\boldsymbol{x}^*) = 0$.

Then, $\boldsymbol{x}^*$ is a global minimizer of f over Ω. □

Proof. Suppose that $\boldsymbol{x} \in \Omega$. By convexity of f and Theorem 22.4,

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*) + Df(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*).$$

Using condition 2, we get

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}^*) - \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) - \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*).$$

As in the proof of Theorem 22.8, we can show that $\boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) = 0$. We now claim that $\boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) \leq 0$. To see this, note that because Ω is convex, $(1-\alpha)\boldsymbol{x}^* + \alpha\boldsymbol{x} \in \Omega$ for all $\alpha \in (0,1)$. Thus,

$$\boldsymbol{g}(\boldsymbol{x}^* + \alpha(\boldsymbol{x} - \boldsymbol{x}^*)) = \boldsymbol{g}((1-\alpha)\boldsymbol{x}^* + \alpha\boldsymbol{x}) \leq \boldsymbol{0}$$

for all $\alpha \in (0,1)$. Premultiplying by $\boldsymbol{\mu}^{*\top} \geq \mathbf{0}^\top$ (by condition 1), subtracting $\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*) = 0$ (by condition 3), and dividing by α, we get

$$\frac{\boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^* + \alpha(\boldsymbol{x} - \boldsymbol{x}^*)) - \boldsymbol{\mu}^{*\top}\boldsymbol{g}(\boldsymbol{x}^*)}{\alpha} \leq 0.$$

We now take the limit as $\alpha \to 0$ to obtain $\boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) \leq 0$.

From the above, we have

$$\begin{aligned} f(\boldsymbol{x}) &\geq f(\boldsymbol{x}^*) - \boldsymbol{\lambda}^{*\top} D\boldsymbol{h}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) - \boldsymbol{\mu}^{*\top} D\boldsymbol{g}(\boldsymbol{x}^*)(\boldsymbol{x} - \boldsymbol{x}^*) \\ &\geq f(\boldsymbol{x}^*) \end{aligned}$$

for all $\boldsymbol{x} \in \Omega$, which completes the proof. ∎

Example 22.7 A bank account starts out with 0 dollars. At the beginning of each month, we deposit some money into the bank account. Denote by x_k the amount deposited in the kth month, $k = 1, 2, \ldots$. Suppose that the monthly interest rate is $r > 0$ and the interest is paid into the account at the end of each month (and compounded). We wish to maximize the total amount of money accumulated at the end of n months, such that the total money deposited during the n months does not exceed D dollars (where $D > 0$).

To solve this problem we first show that the problem can be posed as a linear program, and is therefore a convex optimization problem. Let y_k be the total amount in the bank at the end of the kth month. Then, $y_k = (1+r)(y_{k-1} + x_k)$, $k \geq 1$, with $y_0 = 0$. Therefore, we want to maximize y_n subject to the constraint that $x_k \geq 0$, $k = 1, \ldots, n$, and $x_1 + \cdots + x_n \leq D$. It is easy to deduce that

$$y_n = (1+r)^n x_1 + (1+r)^{n-1} x_2 + \cdots + (1+r) x_n.$$

Let $\boldsymbol{c}^\top = [(1+r)^n, (1+r)^{n-1}, \ldots, (1+r)]$, $\boldsymbol{e}^\top = [1, \ldots, 1]$, and $\boldsymbol{x} = [x_1, \ldots, x_n]^\top$. Then, we can write the problem as

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{e}^\top \boldsymbol{x} \leq D \\ & \boldsymbol{x} \geq \mathbf{0}, \end{aligned}$$

which is a linear program.

It is intuitively clear that the optimal strategy is to deposit D dollars in the first month. To show that this strategy is indeed optimal, we use Theorem 22.9. Let $\boldsymbol{x}^* = [D, 0, \ldots, 0]^\top \in \mathbb{R}^n$. Because the problem is a convex programming problem, it suffices to show that $\boldsymbol{x}^*$ satisfies the KKT condition (see Theorem 22.9). The KKT condition for this problem is

$$\begin{aligned}
-\boldsymbol{c}^\top + \mu^{(1)}\boldsymbol{e}^\top - \boldsymbol{\mu}^{(2)\top} &= 0, \\
\mu^{(1)}(\boldsymbol{e}^\top \boldsymbol{x}^* - D) &= 0, \\
\boldsymbol{\mu}^{(2)\top}\boldsymbol{x}^* &= 0, \\
\boldsymbol{e}^\top \boldsymbol{x}^* - D &\leq 0, \\
-\boldsymbol{x}^* &\leq \boldsymbol{0}, \\
\mu^{(1)} &\geq 0, \\
\boldsymbol{\mu}^{(2)} &\geq \boldsymbol{0}, \\
\boldsymbol{e}^\top \boldsymbol{x} &\leq D, \\
\boldsymbol{x} &\geq \boldsymbol{0},
\end{aligned}$$

where $\mu^{(1)} \in \mathbb{R}$ and $\boldsymbol{\mu}^{(2)} \in \mathbb{R}^n$. Let $\mu^{(1)} = (1+r)^n$ and $\boldsymbol{\mu}^{(2)} = (1+r)^n\boldsymbol{e} - \boldsymbol{c}$. Then, it is clear that the KKT condition is satisfied. Therefore, $\boldsymbol{x}^*$ is a global minimizer. ∎

An entire book devoted to the vast topic of convexity and optimization is [7]. For extensions of the theory of convex optimization, we refer the reader to [136, Chapter 10]. The study of convex programming problems also serves as a prerequisite to *nondifferentiable optimization* (see, e.g., [38]).

22.4 Semidefinite Programming

Semidefinite programming is a subfield of convex optimization concerned with minimizing a linear objective function subject to a linear matrix inequality. The linear matrix inequality constraint defines a convex feasible set over which the linear objective function is to be minimized. Semidefinite programming can be viewed as an extension of linear programming, where the componentwise inequalities on vectors are replaced by matrix inequalities (see Exercise 22.20). For further reading on the subject of semidefinite programming, we recommend an excellent survey paper by Vandenberghe and Boyd [128].

Linear Matrix Inequalities and Their Properties

Consider $n+1$ real symmetric matrices

$$\boldsymbol{F}_i = \boldsymbol{F}_i^\top \in \mathbb{R}^{m\times m}, \quad i = 0, 1, \ldots, n$$

and a vector

$$\boldsymbol{x} = [x_1, \ldots, x_n]^\top \in \mathbb{R}^n.$$

Then,

$$
\begin{aligned}
\boldsymbol{F}(\boldsymbol{x}) &= \boldsymbol{F}_0 + x_1\boldsymbol{F}_1 + \cdots + x_n\boldsymbol{F}_n \\
&= \boldsymbol{F}_0 + \sum_{i=1}^{n} x_i\boldsymbol{F}_i
\end{aligned}
$$

is an affine function of $\boldsymbol{x}$, because $\boldsymbol{F}(\boldsymbol{x})$ is composed of a linear term $\sum_{i=1}^{n} x_i\boldsymbol{F}_i$ and a constant term $\boldsymbol{F}_0$.

Consider now an inequality constraint of the form

$$\boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{F}_0 + x_1\boldsymbol{F}_1 + \cdots + x_n\boldsymbol{F}_n \geq 0.$$

The inequality constraint above is to be interpreted as the set of vectors $\boldsymbol{x}$ such that

$$\boldsymbol{z}^\top \boldsymbol{F}(\boldsymbol{x})\boldsymbol{z} \geq 0 \ \text{ for all } \ \boldsymbol{z} \in \mathbb{R}^m;$$

that is, $\boldsymbol{F}(\boldsymbol{x})$ is positive semidefinite [or, in the usual notation, $\boldsymbol{F}(\boldsymbol{x}) \geq 0$]. Recall that the terms $\boldsymbol{F}_i$ represent constant matrices, $\boldsymbol{x}$ is unknown, and $\boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{F}(\boldsymbol{x})^\top$ is an affine function $\boldsymbol{x}$. The expression $\boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{F}_0 + x_1\boldsymbol{F}_1 + \cdots + x_n\boldsymbol{F}_n \geq 0$ is referred to in the literature as a *linear matrix inequality* (LMI), although the term *affine matrix inequality* would seem to be more appropriate. It is easy to verify that the set $\{\boldsymbol{x} : \boldsymbol{F}(\boldsymbol{x}) \geq 0\}$ is convex (see Exercise 22.20).

We can speak similarly of LMIs of the form $\boldsymbol{F}(\boldsymbol{x}) > 0$, where the requirement is for $\boldsymbol{F}(\boldsymbol{x})$ to be positive definite (rather than just positive semidefinite). It is again easy to see that the set $\{\boldsymbol{x} : \boldsymbol{F}(\boldsymbol{x}) > 0\}$ is convex.

A system of LMIs

$$\boldsymbol{F}_1(\boldsymbol{x}) \geq 0, \ \boldsymbol{F}_2(\boldsymbol{x}) \geq 0, \ \ldots, \ \boldsymbol{F}_k(\boldsymbol{x}) \geq 0$$

can be represented as one single LMI:

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} \boldsymbol{F}_1(\boldsymbol{x}) & & & \\ & \boldsymbol{F}_2(\boldsymbol{x}) & & \\ & & \ddots & \\ & & & \boldsymbol{F}_k(\boldsymbol{x}) \end{bmatrix} \geq 0.$$

As an example, a linear inequality involving an $m \times n$ real constant matrix $\boldsymbol{A}$ of the form

$$\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}$$

can be represented as m LMIs:

$$b_i - \boldsymbol{a}_i^\top \boldsymbol{x} \geq 0, \quad i = 1, 2, \ldots, m,$$

where $\boldsymbol{a}_i^\top$ is the ith row of the matrix $\boldsymbol{A}$. We can view each scalar inequality as an LMI. We then represent m LMIs as one LMI:

$$\boldsymbol{F}(\boldsymbol{x}) = \begin{bmatrix} b_1 - \boldsymbol{a}_1^\top \boldsymbol{x} & & & \\ & b_2 - \boldsymbol{a}_2^\top \boldsymbol{x} & & \\ & & \ddots & \\ & & & b_m - \boldsymbol{a}_m^\top \boldsymbol{x} \end{bmatrix} \geq 0.$$

With the foregoing facts as background, we can now give an example of semidefinite programming:

$$\begin{array}{rl} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{F}(\boldsymbol{x}) \geq 0. \end{array}$$

The matrix property that we discuss next is useful when converting certain LMIs or nonlinear matrix inequalities into equivalent LMIs. We start with a simple observation. Let $\boldsymbol{P}$ be a nonsingular $n \times n$ matrix and let $\boldsymbol{x} = \boldsymbol{M}\boldsymbol{z}$, where $\boldsymbol{M} \in \mathbb{R}^{n \times n}$ such that $\det \boldsymbol{M} \neq 0$. Then, we have

$$\boldsymbol{x}^\top \boldsymbol{P} \boldsymbol{x} \geq 0 \text{ if and only if } \boldsymbol{z}^\top \boldsymbol{M}^\top \boldsymbol{P} \boldsymbol{M} \boldsymbol{z} \geq 0;$$

that is,

$$\boldsymbol{P} \geq 0 \text{ if and only if } \boldsymbol{M}^\top \boldsymbol{P} \boldsymbol{M} \geq 0.$$

Similarly,

$$\boldsymbol{P} > 0 \text{ if and only if } \boldsymbol{M}^\top \boldsymbol{P} \boldsymbol{M} > 0.$$

Suppose that we have a square matrix

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}^\top & \boldsymbol{D} \end{bmatrix}.$$

Then, by the observation above,

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}^\top & \boldsymbol{D} \end{bmatrix} \geq 0 \text{ if and only if } \begin{bmatrix} \boldsymbol{O} & \boldsymbol{I} \\ \boldsymbol{I} & \boldsymbol{O} \end{bmatrix} \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}^\top & \boldsymbol{D} \end{bmatrix} \begin{bmatrix} \boldsymbol{O} & \boldsymbol{I} \\ \boldsymbol{I} & \boldsymbol{O} \end{bmatrix} \geq 0,$$

where $\boldsymbol{I}$ is an identity matrix of appropriate dimension. In other words,

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}^\top & \boldsymbol{D} \end{bmatrix} \geq 0 \text{ if and only if } \begin{bmatrix} \boldsymbol{D} & \boldsymbol{B}^\top \\ \boldsymbol{B} & \boldsymbol{A} \end{bmatrix} \geq 0.$$

We now introduce the notion of the Schur complement, useful in studying LMIs. Consider a square matrix of the form

$$\begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{bmatrix},$$

where A_{11} and A_{22} are square submatrices. Suppose that the matrix A_{11} is invertible. Then, we have

$$\begin{bmatrix} I & O \\ -A_{21}A_{11}^{-1} & I \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} I & -A_{11}^{-1}A_{12} \\ O & I \end{bmatrix} = \begin{bmatrix} A_{11} & O \\ O & A_{22} - A_{21}A_{11}^{-1}A_{12} \end{bmatrix}.$$

Let

$$\Delta_{11} = A_{22} - A_{21}A_{11}^{-1}A_{12},$$

which is called the *Schur complement* of A_{11}. For the case when $A_{12} = A_{21}^{\top}$, we have

$$\begin{bmatrix} I & O \\ -A_{21}A_{11}^{-1} & I \end{bmatrix} \begin{bmatrix} A_{11} & A_{21}^{\top} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} I & -A_{11}^{-1}A_{21}^{\top} \\ O & I \end{bmatrix} = \begin{bmatrix} A_{11} & O \\ O & \Delta_{11} \end{bmatrix},$$

where

$$\Delta_{11} = A_{22} - A_{21}A_{11}^{-1}A_{21}^{\top}.$$

Hence,

$$\begin{bmatrix} A_{11} & A_{21}^{\top} \\ A_{21} & A_{22} \end{bmatrix} > 0 \text{ if and only if } \begin{bmatrix} A_{11} & O \\ O & \Delta_{11} \end{bmatrix} > 0;$$

that is,

$$\begin{bmatrix} A_{11} & A_{21}^{\top} \\ A_{21} & A_{22} \end{bmatrix} > 0 \text{ if and only if } A_{11} > 0 \text{ and } \Delta_{11} > 0.$$

Given

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix},$$

we can similarly define the Schur complement of A_{22}, assuming that A_{22} is invertible. We have

$$\begin{bmatrix} I & -A_{12}A_{22}^{-1} \\ O & I \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} I & O \\ -A_{22}^{-1}A_{21} & I \end{bmatrix} = \begin{bmatrix} \Delta_{22} & O \\ O & A_{22} \end{bmatrix},$$

where $\Delta_{22} = A_{11} - A_{12}A_{22}^{-1}A_{21}$ is the Schur complement of A_{22}. So, for the case where $A_{12} = A_{21}^{\top}$,

$$\begin{bmatrix} A_{11} & A_{21}^{\top} \\ A_{21} & A_{22} \end{bmatrix} > 0 \text{ if and only if } A_{22} > 0 \text{ and } \Delta_{22} > 0.$$

Many problems of optimization, control design, and signal processing can be formulated in terms of LMIs. To determine whether or not there exists a point x such that $F(x) > 0$ is called a *feasibility problem*. We say that the LMI is nonfeasible if no such solution exists.

Example 22.8 We now present a simple example illustrating the LMI feasibility problem. Let $\boldsymbol{A} \in \mathbb{R}^{m \times m}$ be a given real constant square matrix. Suppose that we wish to determine if $\boldsymbol{A}$ has all its eigenvalues in the open left half-complex plane. It is well known that this condition is true if and only if there exists a real symmetric positive definite matrix $\boldsymbol{P}$ such that

$$\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} < 0,$$

or, equivalently,

$$-\boldsymbol{A}^\top \boldsymbol{P} - \boldsymbol{P}\boldsymbol{A} > 0$$

(also called the *Lyapunov inequality*; see [16]). Thus, the location of all eigenvalues of $\boldsymbol{A}$ being in the open left half-complex plane is equivalent to feasibility of the following matrix inequality:

$$\begin{bmatrix} \boldsymbol{P} & \boldsymbol{O} \\ \boldsymbol{O} & -\boldsymbol{A}^\top \boldsymbol{P} - \boldsymbol{P}\boldsymbol{A} \end{bmatrix} > 0;$$

that is, the existence of $\boldsymbol{P} = \boldsymbol{P}^\top > 0$ such that $\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} < 0$.

We now show that finding $\boldsymbol{P} = \boldsymbol{P}^\top > 0$ such that $\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} < 0$ is indeed an LMI. For this, let

$$\boldsymbol{P} = \begin{bmatrix} x_1 & x_2 & \cdots & x_m \\ x_2 & x_{m+1} & \cdots & x_{2m-1} \\ \vdots & & & \vdots \\ x_m & x_{2m-1} & \cdots & x_n \end{bmatrix},$$

where

$$n = \frac{m(m+1)}{2}.$$

We next define the following matrices:

$$\boldsymbol{P}_1 = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & & & & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix},$$

$$\boldsymbol{P}_2 = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & & & & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix},$$

$$\vdots$$

$$\boldsymbol{P}_n = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & & & & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

Note that $\boldsymbol{P}_i$ has only nonzero elements corresponding to $\boldsymbol{x}_i$ in $\boldsymbol{P}$. Let

$$\boldsymbol{F}_i = -\boldsymbol{A}^\top \boldsymbol{P}_i - \boldsymbol{P}_i \boldsymbol{A}, \quad i = 1, 2, \ldots, n.$$

We can then write

$$\begin{aligned} \boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} &= x_1 \left(\boldsymbol{A}^\top \boldsymbol{P}_1 + \boldsymbol{P}_1 \boldsymbol{A}\right) + x_2 \left(\boldsymbol{A}^\top \boldsymbol{P}_2 + \boldsymbol{P}_2 \boldsymbol{A}\right) + \cdots \\ &\quad + x_n \left(\boldsymbol{A}^\top \boldsymbol{P}_n + \boldsymbol{P}_n \boldsymbol{A}\right) \\ &= -x_1 \boldsymbol{F}_1 - x_2 \boldsymbol{F}_2 - \cdots - x_n \boldsymbol{F}_n \\ &< 0. \end{aligned}$$

Let

$$\boldsymbol{F}(\boldsymbol{x}) = x_1 \boldsymbol{F}_1 + x_2 \boldsymbol{F}_2 + \cdots + x_n \boldsymbol{F}_n.$$

Then,

$$\boldsymbol{P} = \boldsymbol{P}^\top > 0 \quad \text{and} \quad \boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} < 0$$

if and only if

$$\boldsymbol{F}(\boldsymbol{x}) > 0.$$

Note that this LMI involves a strict inequality. Most numerical solvers do not handle strict inequalities. Such solvers simply treat a strict inequality ($>$) as a non-strict inequality ($\geq$). ∎

LMI Solvers

The inequality $\boldsymbol{F}(\boldsymbol{x}) = \boldsymbol{F}_0 + x_1\boldsymbol{F}_1 + \cdots + x_n\boldsymbol{F}_n \geq 0$ is called the *canonical representation* of an LMI. Numerical LMI solvers do not deal directly with LMIs in canonical form because of various inefficiencies. Instead, LMI solvers use a structured representation of LMIs.

We can use MATLAB's LMI toolbox to solve LMIs efficiently. This toolbox has three types of LMI solvers, which we discuss next.

Finding a Feasible Solution Under LMI Constraints

First, we discuss MATLAB's LMI solver for solving the feasibility problem defined by a given system of LMI constraints. Using this solver, we can solve any system of LMIs of the form

$$\boldsymbol{N}^\top \boldsymbol{\mathcal{L}}(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)\boldsymbol{N} \leq \boldsymbol{M}^\top \boldsymbol{\mathcal{R}}(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)\boldsymbol{M},$$

where $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k$ are matrix variables, $\boldsymbol{N}$ is the left outer factor, $\boldsymbol{M}$ is the right outer factor, $\boldsymbol{\mathcal{L}}(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ is the left inner factor, and $\boldsymbol{\mathcal{R}}(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_k)$ is the right inner factor. The matrices $\boldsymbol{\mathcal{L}}(\cdot)$ and $\boldsymbol{\mathcal{R}}(\cdot)$ are, in general, symmetric block matrices. We note that the term left-hand side refers to what is on the "smaller" side of the inequality $0 \leq \boldsymbol{X}$. Thus in $\boldsymbol{X} \geq 0$, the matrix $\boldsymbol{X}$ is still on the right-hand side because it is on the "larger " side of the inequality.

We now provide a description of an approach that can be used to solve the given LMI system feasibility problem. To initialize the LMI system description, we type `setlmis([])`. Then we declare matrix variables using the command `lmivar`. The command `lmiterm` allows us to specify LMIs that constitute the LMI system under consideration. Next, we need to obtain an internal representation using the command `getlmis`. We next compute a feasible solution to the LMI system using the command `feasp`. After that, we extract matrix variable values with the command `dec2mat`. In summary, a general structure of a MATLAB program for finding a feasible solution to the set of LMIs could have the form

```
setlmis([])
lmivar
lmiterm
 .
 .
 .
lmiterm
getlmis
feasp
dec2mat
```

We now analyze these commands in some detail so that the reader can write simple MATLAB programs for solving LMIs after completing this section.

First, to create a new matrix-valued variable, say, `X`, in the given LMI system, we use the command

```
X = lmivar(type,structure)
```

The input `type` specifies the structure of the variable `X`. There may be three structures of matrix variables. When `type=1`, we have a symmetric block diagonal matrix variable. The input `type=2` refers to a full rectangular matrix variable. Finally, `type=3` refers to other cases. The second input `structure` gives additional information on the structure of the matrix variable `X`. For example, the matrix variable `X` could have the form

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{D}_1 & \boldsymbol{O} & \cdots & \boldsymbol{O} \\ \boldsymbol{O} & \boldsymbol{D}_2 & \cdots & \boldsymbol{O} \\ \vdots & & \ddots & \vdots \\ \boldsymbol{O} & \boldsymbol{O} & \cdots & \boldsymbol{D}_r \end{bmatrix},$$

where each $\boldsymbol{D}_i$ is a square symmetric matrix. For the example above we would use `type=1`. The matrix variable above has r blocks. The input `structure` is then an $r \times 2$ matrix whose ith row describes the ith block, where the first component of each row gives the corresponding block size, while the second element of each row specifies the block type. For example,

```
X = lmivar(1,[3 1])
```

specifies a full symmetric 3×3 matrix variable. On the other hand,

```
X = lmivar(2,[2 3])
```

specifies a rectangular 2×3 matrix variable. Finally, a matrix variable $\boldsymbol{S}$ of the form

$$\boldsymbol{S} = \left[\begin{array}{cc|cc} s_1 & 0 & 0 & 0 \\ 0 & s_1 & 0 & 0 \\ \hline 0 & 0 & s_2 & s_3 \\ 0 & 0 & s_3 & s_4 \end{array}\right]$$

can be declared as follows:

```
S = lmivar(1,[2 0;2 1])
```

Note above that the second component of the first row of the second input has the value of zero; that is, `structure(1,2)=0`. This describes a scalar block matrix of the form

$$\boldsymbol{D}_1 = s_1 \boldsymbol{I}_2.$$

Note that the second block is a 2×2 symmetric full block.

We next take a closer look at a command whose purpose is to specify the terms of the LMI system of interest. This command has the form

```
lmiterm(termid,A,B,flag)
```

We briefly describe each of the four inputs of this command. The first input, `termid`, is a row with four elements that specify the terms of each LMI of the LMI system. We have `termid(1)=n` to specify the left-hand side of the nth LMI. We use `termid(1)=-n` to specify the right-hand side of the nth LMI. The middle two elements of the input `termid` specify the block location. Thus `termid(2,3)=[i j]` refers to the term that belongs to the (i, j) block of the LMI specified by the first component. Finally, `termid(4)=0` for the constant term, `termid(4)=X` for the variable term in the form $\boldsymbol{AXB}$, while `termid(4)=-X` for the variable term in the form $\boldsymbol{AX}^\top\boldsymbol{B}$. The second and third inputs of the command `lmiterm` give the values of the left and right outer factors; that is, `A` and `B` give the values of the constant outer factors in the variable terms $\boldsymbol{AXB}$ and $\boldsymbol{AX}^\top\boldsymbol{B}$. Finally, the fourth input to `lmiterm` serves as a compact way to specify the expression

$$\boldsymbol{AXB} + (\boldsymbol{AXB})^\top.$$

Thus, `flag='s'` can be used to denote a symmetrized expression. We now illustrate the command above on the following LMI:

$$\boldsymbol{PA} + (\boldsymbol{PA})^\top \leq 0.$$

We have one LMI with two terms. We could use the following description of this single LMI:

```
lmiterm([1 1 1 P],1,A)
lmiterm([1 1 1 -P],A',1)
```

On the other hand, we can describe this LMI compactly using the `flag` as follows:

```
lmiterm([1 1 1 P],1,A,'s')
```

Now, to solve the feasibility problem we could have typed

```
[tmin,xfeas] = feas(lmis)
```

In general, for a given LMI feasibility problem of the form

$$\begin{aligned} &\text{find} && \boldsymbol{x} \\ &\text{such that} && \boldsymbol{L}(\boldsymbol{x}) \leq \boldsymbol{R}(\boldsymbol{x}), \end{aligned}$$

the command `feasp` solves the auxiliary convex problem

$$\begin{aligned} &\text{minimize} && t \\ &\text{subject to} && \boldsymbol{L}(\boldsymbol{x}) \leq \boldsymbol{R}(\boldsymbol{x}) + t\boldsymbol{I}. \end{aligned}$$

The system of LMIs is feasible if the minimal t is negative. We add that the current value of t is displayed by `feasp` at each iteration.

Finally, we convert the output of the LMI solver into matrix variables using the command

```
P = dec2mat(lmis,xfeas,P).
```

Example 22.9 Let

$$\boldsymbol{A}_1 = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \quad \text{and} \quad \boldsymbol{A}_2 = \begin{bmatrix} -2 & 0 \\ 1 & -1 \end{bmatrix}.$$

We use the commands of the LMI Control Toolbox discussed above to write a program that finds $\boldsymbol{P}$ such that $\boldsymbol{P} \geq 0.5\boldsymbol{I}_2$ and

$$\begin{aligned} \boldsymbol{A}_1^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A}_1 &\leq 0, \\ \boldsymbol{A}_2^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A}_2 &\leq 0. \end{aligned}$$

The program is as follows:

```
A_1 = [-1 0;0 -1];
A_2 = [-2 0;1 -1];
setlmis([])
P = lmivar(1,[2,1])
lmiterm([1 1 1 P],A_1',1,'s')
lmiterm([2 1 1 P],A_2',1,'s')
lmiterm([3 1 1 0],.5)
lmiterm([-3 1 1 P],1,1)
lmis=getlmis;
[tmin,xfeas] = feasp(lmis);
P = dec2mat(lmis,xfeas,P)
```

∎

Minimizing a Linear Objective Under LMI Constraints

The next solver we discuss solves the convex optimization problem

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}(\boldsymbol{x}) \leq \boldsymbol{B}(\boldsymbol{x}). \end{aligned}$$

The notation $\boldsymbol{A}(\boldsymbol{x}) \leq \boldsymbol{B}(\boldsymbol{x})$ is shorthand notation for a general structured LMI system.

This solver is invoked using the function `mincx`. Thus, to solve a `mincx` problem, in addition to specifying the LMI constraints as in the `feasp` problem, we also declare the linear objective function. Then we invoke the function mincx. We illustrate and contrast the feasp and mincx solvers in the following example.

Example 22.10 Consider the optimization problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}, \end{array}$$

where

$$\boldsymbol{c}^\top = \begin{bmatrix} 4 & 5 \end{bmatrix},$$
$$\boldsymbol{A} = \begin{bmatrix} 1 & 1 \\ 1 & 3 \\ 2 & 1 \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} 8 \\ 18 \\ 14 \end{bmatrix}.$$

We first solve the feasibility problem; that is, we find an $\boldsymbol{x}$ such that $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}$, using the feasp solver. After that, we solve the minimization problem above using the mincx solver. A simple MATLAB code accomplishing these tasks is shown below.

```
% Enter problem data
A = [1 1;1 3;2 1];
b = [8 18 14]';
c = [-4 -5]';
setlmis([]);
X = lmivar(2,[2 1]);
lmiterm([1 1 1 X],A(1,:),1);
lmiterm([1 1 1 0],-b(1));
lmiterm([1 2 2 X],A(2,:),1);
lmiterm([1 2 2 0],-b(2));
lmiterm([1 3 3 X],A(3,:),1);
lmiterm([1 3 3 0],-b(3));
lmis = getlmis;
%------------------------------------
disp('-------------feasp result---------------')
[tmin,xfeas] = feasp(lmis);
x_feasp = dec2mat(lmis,xfeas,X)
disp('-------------mincx result---------------')
[objective,x_mincx] = mincx(lmis,c,[0.0001 1000 0 0 1])
```

The feasp function produces

$$\boldsymbol{x}_{\text{feasp}} = \begin{bmatrix} -64.3996 \\ -25.1712 \end{bmatrix}.$$

The `mincx` function produces

$$\boldsymbol{x}_{\text{mincx}} = \begin{bmatrix} 3.0000 \\ 5.0000 \end{bmatrix}.$$

∎

In the next example, we discuss the function `defcx`, which we can use to construct the vector $\boldsymbol{c}$ used by the LMI solver `mincx`.

Example 22.11 Suppose that we wish to solve the optimization problem

$$\begin{aligned} \text{minimize} \quad & \operatorname{trace}(\boldsymbol{P}) \\ \text{subject to} \quad & \boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} \leq 0 \\ & \boldsymbol{P} \geq 0 \end{aligned}$$

where $\operatorname{trace}(\boldsymbol{P})$ is the sum of the diagonal elements of $\boldsymbol{P}$. We can use the function `mincx` to solve this problem. However, to use `mincx`, we need a vector $\boldsymbol{c}$ such that

$$\boldsymbol{c}^\top \boldsymbol{x} = \operatorname{trace}(\boldsymbol{P}).$$

After specifying the LMIs and obtaining their internal representation using, for example, the command `lmisys=getlmis`, we can obtain the desired $\boldsymbol{c}$ with the following MATLAB code,

```
q = decnbr(lmisys);
c = zeros(q,1);
for j = 1:q
    Pj = defcx(lmisys,j,P);
    c(j) = trace(Pj);
end
```

Having obtained the vector $\boldsymbol{c}$, we can use the function `mincx` to solve the optimization problem. ∎

Minimizing a Generalized Eigenvalue Under LMI Constraints

This problem can be stated as

$$\begin{aligned} \text{minimize} \quad & \lambda \\ \text{subject to} \quad & \boldsymbol{C}(\boldsymbol{x}) \leq \boldsymbol{D}(\boldsymbol{x}) \\ & 0 \leq \boldsymbol{B}(\boldsymbol{x}) \\ & \boldsymbol{A}(\boldsymbol{x}) \leq \lambda \boldsymbol{B}(\boldsymbol{x}). \end{aligned}$$

Here, we need to distinguish between standard LMI constraints of the form $\boldsymbol{C}(\boldsymbol{x}) \leq \boldsymbol{D}(\boldsymbol{x})$ and *linear-fractional LMIs* of the form $\boldsymbol{A}(\boldsymbol{x}) \leq \lambda \boldsymbol{B}(\boldsymbol{x})$, which are concerned with the generalized eigenvalue λ. The generalized eigenvalue minimization problem under LMI constraints can be solved using the solver **gevp**. The basic structure of the **gevp** solver has the form

```
[lopt,xopt] = gep{lmisys,nflc}
```

which returns `lopt`, the global minimum of the generalized eigenvalue, and `xopt`, the optimal decision vector variable. The argument `lmisys` is the system of LMIs, $\boldsymbol{C}(\boldsymbol{x}) \leq \boldsymbol{D}(\boldsymbol{x})$, $\boldsymbol{C}(\boldsymbol{x}) \leq \boldsymbol{D}(\boldsymbol{x})$, and $\boldsymbol{A}(\boldsymbol{x}) \leq \lambda \boldsymbol{B}(\boldsymbol{x})$ for $\lambda = 1$. As in the previous solvers, the corresponding optimal values of the matrix variables are obtained using `dec2mat`. The number of linear-fractional constraints is specified with `nflc`. There are other inputs to **gevp** but they are optional. For more information on this type of the LMI solver, we refer the reader to the LMI Lab in MATLAB's Robust Control Toolbox user's guide.

Example 22.12 Consider the problem of finding the smallest α such that

$$\begin{gathered} \boldsymbol{P} > 0 \\ \boldsymbol{A}^{\top}\boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} \leq -\alpha \boldsymbol{P}, \end{gathered}$$

where

$$\boldsymbol{A} = \begin{bmatrix} -1.1853 & 0.9134 & 0.2785 \\ 0.9058 & -1.3676 & 0.5469 \\ 0.1270 & 0.0975 & -3.0000 \end{bmatrix}.$$

This problem is related to finding the decay rate of the stable linear differential equation $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x}$. Finding α that solves the optimization problem above can be accomplished using the following LMIs:

```
A = [-1.1853   0.9134    0.2785
    0.9058  -1.3676    0.5469
    0.1270   0.0975   -3.0000];
setlmis([]);
P = lmivar(1,[3 1])
lmiterm([-1 1 1 P],1,1)    % P
lmiterm([1 1 1 0],.01)     % P >= 0.01*I
lmiterm([2 1 1 P],1,A,'s') % linear fractional constraint---LHS
lmiterm([-2 1 1 P],1,1)    % linear fractional constraint---RHS
lmis = getlmis;
[gamma,P_opt] = gevp(lmis,1);
P = dec2mat(lmis,P_opt,P)
alpha = -gamma
```

The result is

$$\alpha = 0.6561 \quad \text{and} \quad \boldsymbol{P} = \begin{bmatrix} 0.6996 & -0.7466 & -0.0296 \\ -0.7466 & 0.8537 & -0.2488 \\ -0.0296 & -0.2488 & 3.2307 \end{bmatrix}.$$

Notice that we used $\boldsymbol{P} \geq 0.01\boldsymbol{I}$ in place of $\boldsymbol{P} > 0$. ∎

More examples of linear matrix inequalities in system and control theory can be found in the book by Boyd et al. [16].

A quick introduction to MATLAB's LMI toolbox is the tutorial that can be accessed with the command `lmidem` within MATLAB. In addition to the MATLAB's LMI toolbox, there is another toolbox for solving LMIs called LMITOOL, a built-in software package in Scilab toolbox, developed at INRIA in France. Scilab offers free software for numerical optimization. There is a version of LMITOOL for MATLAB that can be obtained from the website of the Scilab Consortium.

Yet Another LMI Package, YALMIP, for solving LMIs was developed in Switzerland in the Automatic Control Laboratory at ETH. YALMIP is an "intuitive and flexible modelling language for solving optimization problems in MATLAB."

LMIs are tools of modern optimization. The following quote on numerical linear algebra from Gill, Murray, and Wright [52, p. 2] applies as well to the contents of this chapter: "At the heart of modern optimization methods are techniques associated with *linear algebra.* Numerical linear algebra applies not simply in optimization, but in all fields of scientific computation, including approximation, ordinary differential equations, and partial differential equations. *The importance of numerical linear algebra to modern scientific computing cannot be overstated.* Without fast and reliable linear algebraic building blocks, it is impossible to develop effective optimization methods; without some knowledge of the fundamental issues in linear algebra, it is impossible to understand what happens during the transition from equations in a textbook to actual computation."

EXERCISES

22.1 Find the range of values of the parameter α for which the function

$$f(x_1, x_2, x_3) = 2x_1x_3 - x_1^2 - x_2^2 - 5x_3^2 - 2\alpha x_1x_2 - 4x_2x_3$$

is concave.

22.2 Consider the function

$$f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} - \boldsymbol{x}^\top \boldsymbol{b},$$

where $\boldsymbol{Q} = \boldsymbol{Q}^\top > 0$ and $\boldsymbol{x}, \boldsymbol{b} \in \mathbb{R}^n$. Define the function $\phi : \mathbb{R} \to \mathbb{R}$ by $\phi(\alpha) = f(\boldsymbol{x} + \alpha\boldsymbol{d})$, where $\boldsymbol{x}, \boldsymbol{d} \in \mathbb{R}^n$ are fixed vectors and $\boldsymbol{d} \neq \boldsymbol{0}$. Show that $\phi(\alpha)$ is a strictly convex quadratic function of α.

22.3 Show that $f(\boldsymbol{x}) = x_1 x_2$ is a convex function on $\Omega = \{[a, ma]^\top : a \in \mathbb{R}\}$, where m is any given nonnegative constant.

22.4 Suppose that the set $\Omega = \{\boldsymbol{x} : h(\boldsymbol{x}) = c\}$ is convex, where $h : \mathbb{R}^n \to \mathbb{R}$ and $c \in \mathbb{R}$. Show that h is convex and concave over Ω.

22.5 Find all subgradients of

$$f(x) = |x|, \quad x \in \mathbb{R},$$

at $x = 0$ and at $x = 1$.

22.6 Let $\Omega \subset \mathbb{R}^n$ be a convex set, and $f_i : \Omega \to \mathbb{R}$, $i = 1, \dots, \ell$ be convex functions. Show that $\max\{f_1, \dots, f_\ell\}$ is a convex function.
Note: The notation $\max\{f_1, \dots, f_\ell\}$ denotes a function from Ω to $\mathbb{R}$ such that for each $\boldsymbol{x} \in \Omega$, its value is the largest among the numbers $f_i(\boldsymbol{x})$, $i = 1, \dots, \ell$.

22.7 Let $\Omega \subset \mathbb{R}^n$ be an open convex set. Show that a symmetric matrix $\boldsymbol{Q} \in \mathbb{R}^n$ is positive semidefinite if and only if for each $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, $(\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{Q}(\boldsymbol{x}-\boldsymbol{y}) \geq 0$. Show that a similar result for positive definiteness holds if we replace the "$\geq$" by "$>$" in the inequality above.

22.8 Consider the problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|^2 \\ \text{subject to} & x_1 + \cdots + x_n = 1 \\ & x_1, \dots, x_n \geq 0 \end{array}$$

(see also Exercise 21.9). Is the problem a convex optimization problem? If yes, give a complete proof. If no, explain why not, giving complete explanations.

22.9 Consider the optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega, \end{array}$$

where $f(\boldsymbol{x}) = x_1 x_2^2$, where $\boldsymbol{x} = [x_1, x_2]^\top$, and $\Omega = \{\boldsymbol{x} \in \mathbb{R}^2 : x_1 = x_2,\ x_1 \geq 0\}$. (See also Exercise 21.8.) Show that the problem is a convex optimization problem.

22.10 Consider the convex optimization problem

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{x} \in \Omega. \end{array}$$

Suppose that the points $\boldsymbol{y} \in \Omega$ and $\boldsymbol{z} \in \Omega$ are local minimizers. Determine the largest set of points $G \subset \Omega$ for which you can be sure that every point in G is a global minimizer.

22.11 Suppose that we have a convex optimization problem on $\mathbb{R}^3$.

a. Consider the following three feasible points: $[1,0,0]^\top$, $[0,1,0]^\top$, $[0,0,1]^\top$. Suppose that all three have objective function value 1. What can you say about the objective function value of the point $(1/3)[1,1,1]^\top$? Explain fully.

b. Suppose we know that the three points in part a are global minimizers. What can you say about the point $(1/3)[1,1,1]^\top$? Explain fully.

22.12 Consider the optimization problem

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\boldsymbol{x}^\top \boldsymbol{Q}\boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}, \end{array}$$

where $\boldsymbol{Q} \in \mathbb{R}^{n\times n}$, $\boldsymbol{Q} = \boldsymbol{Q}^\top > 0$, $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, and rank $\boldsymbol{A} = m$.

a. Find all points satisfying the Lagrange condition for the problem (in terms of $\boldsymbol{Q}$, $\boldsymbol{A}$, and $\boldsymbol{b}$).

b. Are the points (or point) global minimizers for this problem?

22.13 Let $f : \mathbb{R}^n \to \mathbb{R}$, $f \in \mathcal{C}^1$, be a convex function on the set of feasible points

$$\Omega = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{a}_i^\top \boldsymbol{x} + b_i \geq 0, \quad i = 1, \ldots, p\},$$

where $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_p \in \mathbb{R}^n$, and $b_1, \ldots, b_p \in \mathbb{R}$. Suppose that there exist $\boldsymbol{x}^* \in S$, and $\boldsymbol{\mu}^* \in \mathbb{R}^p$, $\boldsymbol{\mu}^* \leq \boldsymbol{0}$, such that

$$Df(\boldsymbol{x}^*) + \sum_{j \in J(\boldsymbol{x}^*)} \mu_j^* \boldsymbol{a}_j^\top = \boldsymbol{0}^\top,$$

where $J(\boldsymbol{x}^*) = \{i : \boldsymbol{a}_i^\top \boldsymbol{x}^* + b_i = 0\}$. Show that $\boldsymbol{x}^*$ is a global minimizer of f over Ω.

22.14 Consider the problem: minimize $\|\boldsymbol{x}\|^2$ ($\boldsymbol{x} \in \mathbb{R}^n$) subject to $\boldsymbol{a}^\top \boldsymbol{x} \geq b$, where $\boldsymbol{a} \in \mathbb{R}^n$ is a nonzero vector and $b \in \mathbb{R}$, $b > 0$. Suppose that $\boldsymbol{x}^*$ is a solution to the problem.

a. Show that the constraint set is convex.

b. Use the KKT theorem to show that $\boldsymbol{a}^\top \boldsymbol{x}^* = b$.

c. Show that $\boldsymbol{x}^*$ is unique, and find an expression for $\boldsymbol{x}^*$ in terms of $\boldsymbol{a}$ and b.

22.15 Consider the problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{R}^n \\ \text{subject to} & \boldsymbol{x} \geq \mathbf{0}. \end{array}$$

For this problem we have the following theorem (see also Exercise 17.16).
Theorem: *A solution to this problem exists if and only if* $\mathbf{c} \geq \mathbf{0}$. *Moreover, if a solution exists,* $\mathbf{0}$ *is a solution.*

a. Show that the problem is a convex programming problem.

b. Use the first-order necessary condition (for set constraints) to prove the theorem.

c. Use the KKT condition to prove the above theorem.

22.16 Consider a linear programming problem in standard form.

a. Derive the KKT condition for the problem.

b. Explain precisely why the KKT condition is sufficient for optimality in this case.

c. Write down the dual to the standard form primal problem (see Chapter 17).

d. Suppose that $\boldsymbol{x}^*$ and $\boldsymbol{\lambda}^*$ are feasible solutions to the primal and dual, respectively. Use the KKT condition to prove that if the complementary slackness condition $(\boldsymbol{c}^\top - \boldsymbol{\lambda}^{*\top}\boldsymbol{A})\boldsymbol{x}^* = 0$ holds, then $\boldsymbol{x}^*$ is an optimal solution to the primal problem. Compare the above with Exercise 21.15.

22.17 Consider two real-valued discrete-time signals, $\boldsymbol{s}^{(1)}$ and $\boldsymbol{s}^{(2)}$, defined over the time interval $[1, n]$. Let $s_i^{(1)}$ and $s_i^{(2)}$ be the values at time i of the signals $\boldsymbol{s}^{(1)}$ and $\boldsymbol{s}^{(2)}$, respectively. Assume that the energies of the two signals are 1 [i.e., $(s_1^{(1)})^2 + \cdots + (s_n^{(1)})^2 = 1$ and $(s_1^{(2)})^2 + \cdots + (s_n^{(2)})^2 = 1$].

Let S_a be the set of all signals that are linear combinations of $\boldsymbol{s}^{(1)}$ and $\boldsymbol{s}^{(2)}$ with the property that for each signal in S_a, the value of the signal over all time is no smaller than $a \in \mathbb{R}$. For each $\boldsymbol{s} \in S_a$, if $\boldsymbol{s} = x_1\boldsymbol{s}^{(1)} + x_2\boldsymbol{s}^{(2)}$, we call x_1 and x_2 the *coefficients* of $\boldsymbol{s}$.

We wish to find a signal in S_a such that the sum of the squares of its coefficients is minimized.

a. Formulate the problem as an optimization problem.

b. Derive the Karush-Kuhn-Tucker conditions for the problem.

c. Suppose that you have found a point satisfying the Karush-Kuhn-Tucker conditions. Does this point satisfy the second-order sufficient condition?

d. Is this problem a convex optimization problem?

22.18 Let a *probability vector* be any vector $\boldsymbol{p} \in \mathbb{R}^n$ satisfying $p_i > 0$, $i = 1, \ldots, n$, and $p_1 + \cdots + p_n = 1$.

Let $\boldsymbol{p} \in \mathbb{R}^n$ and $\boldsymbol{q} \in \mathbb{R}^n$ be two probability vectors. Define

$$D(\boldsymbol{p}, \boldsymbol{q}) = p_1 \log\left(\frac{p_1}{q_1}\right) + \cdots + p_n \log\left(\frac{p_n}{q_n}\right),$$

where "log" is the natural logarithm function.

a. Let Ω be the set of all probability vectors (with fixed n). Show that Ω is convex.

b. Show that for each fixed $\boldsymbol{p}$, the function f defined by $f(\boldsymbol{q}) = D(\boldsymbol{p}, \boldsymbol{q})$ is convex over Ω.

c. Show the following: $D(\boldsymbol{p}, \boldsymbol{q}) \geq 0$ for any probability vectors $\boldsymbol{p}$ and $\boldsymbol{q}$. Moreover, $D(\boldsymbol{p}, \boldsymbol{q}) = 0$ if and only if $\boldsymbol{p} = \boldsymbol{q}$.

d. Describe an application of the result of part c.

22.19 Let $\Omega \subset \mathbb{R}^n$ be a nonempty closed convex set and $\boldsymbol{z} \in \mathbb{R}^n$ be a given point such that $\boldsymbol{z} \notin \Omega$. Consider the optimization problem

$$\begin{array}{ll} \text{minimize} & \|\boldsymbol{x} - \boldsymbol{z}\| \\ \text{subject to} & \boldsymbol{x} \in \Omega. \end{array}$$

Does this problem have an optimal solution? If so, is it unique? Whatever your assertion, prove it.

Hint: (i) If $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are optimal solutions, what can you say about $\boldsymbol{x}_3 = (\boldsymbol{x}_1 + \boldsymbol{x}_2)/2$? (ii) The triangle inequality states that $\|\boldsymbol{x} + \boldsymbol{y}\| \leq \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$, with equality holding if and only if $\boldsymbol{x} = \alpha \boldsymbol{y}$ for some $\alpha \geq 0$ (or $\boldsymbol{x} = 0$ or $\boldsymbol{y} = 0$).

22.20 This exercise is about *semidefinite programming.*

a. Show that if $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ and $\boldsymbol{B} \in \mathbb{R}^{n\times n}$ are symmetric and $\boldsymbol{A} \geq 0$, $\boldsymbol{B} \geq 0$, then for any $\alpha \in (0, 1)$, we have $\alpha\boldsymbol{A} + (1 - \alpha)\boldsymbol{B} \geq 0$. As usual, the notation "≥ 0" denotes positive semidefiniteness.

b. Consider the following semidefinite programming problem, that is, an optimization problem with linear objective function and linear matrix inequality constraints:

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{F}_0 + \sum_{j=1}^{n} x_j \boldsymbol{F}_j \geq 0, \end{array}$$

where $\boldsymbol{x} = [x_1, \ldots, x_n]^\top \in \mathbb{R}^n$ is the decision variable, $\boldsymbol{c} \in \mathbb{R}^n$, and $\boldsymbol{F}_0, \boldsymbol{F}_1, \ldots, \boldsymbol{F}_n \in \mathbb{R}^{m \times m}$ are symmetric.

Show that this problem is a convex optimization problem.

c. Consider the linear programming problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}, \end{array}$$

where $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, $\boldsymbol{b} \in \mathbb{R}^m$, and the inequality $\boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}$ has the usual elementwise interpretation. Show that this linear programming problem can be converted to the problem in part b.
Hint: First consider diagonal $\boldsymbol{F}_j$.

22.21 Suppose that you have a cake and you need to divide it among n different children. Suppose that the ith child receives a fraction x_i of the cake. We will call the vector $\boldsymbol{x} = [x_1, \ldots, x_n]^\top$ an *allocation*. We require that every child receives at least some share of the cake, and that the entire cake is completely used up in the allocation. We also impose the additional condition that the first child ($i = 1$) is allocated a share that is at least twice that of any other child. We say that the allocation is feasible if it meets all these requirements.

A feasible allocation $\boldsymbol{x}$ is said to be *proportionally fair* if for any other allocation $\boldsymbol{y}$,

$$\sum_{i=1}^{n} \frac{y_i - x_i}{x_i} \leq 0.$$

a. Let Ω be the set of all feasible allocations. Show that Ω is convex.

b. Show that a feasible allocation is proportionally fair if and only if it solves the following optimization problem:

$$\begin{array}{ll} \text{maximize} & \sum_{i=1}^{n} \log(x_i) \\ \text{subject to} & \boldsymbol{x} \in \Omega. \end{array}$$

22.22 Let $U_i : \mathbb{R} \to \mathbb{R}$, $U_i \in \mathcal{C}^1$, $i = 1, \dots, n$, be a set of concave increasing functions. Consider the optimization problem

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{i=1}^{n} U_i(x_i) \\ \text{subject to} & \displaystyle\sum_{i=1}^{n} x_i \le C, \end{array}$$

where $C > 0$ is a given constant.

a. Show that the optimization problem above is a convex optimization problem.

b. Show that $\boldsymbol{x}^* = [x_1^*, \dots, x_n^*]^\top$ is an optimal solution to the optimization problem if and only if there exists a scalar $\mu^* \ge 0$ such that $x_i^* = \arg\max_x (U_i(x) - \mu^* x)$. [The quantity $U_i(x)$ has the interpretation of the "utility" of x, whereas μ^* has the interpretation of a "price" per unit of x.]

c. Show that $\sum_{i=1}^{n} x_i^* = C$.

22.23 Give an example of a function $f : \mathbb{R}^2 \to \mathbb{R}$, a set $\Omega = \{\boldsymbol{x} : g(\boldsymbol{x}) \le 0\}$, and a regular point $\boldsymbol{x}^* \in \Omega$, such that the following all hold simultaneously:

1. $\boldsymbol{x}^*$ satisfies the FONC for set constraint Ω (Theorem 6.1).
2. $\boldsymbol{x}^*$ satisfies the KKT condition for inequality constraint $g(\boldsymbol{x}) \le 0$ (Theorem 21.1).
3. $\boldsymbol{x}^*$ satisfies the SONC for set constraint Ω (Theorem 6.2).
4. $\boldsymbol{x}^*$ does not satisfy the SONC for inequality constraint $g(\boldsymbol{x}) \le 0$ (Theorem 21.2).

Be sure to show carefully that your choice of f, $\Omega = \{\boldsymbol{x} : g(\boldsymbol{x}) \le 0\}$, and $\boldsymbol{x}^*$ satisfies all the conditions above simultaneously.

22.24 This question is on duality theory for *nonlinear* programming problems, analogous to the theory for linear programming (Chapter 17). (A version for quadratic programming is considered in Exercise 17.24.)

Consider the following optimization problem:

$$\begin{array}{ll} \text{minimize} & f(\boldsymbol{x}) \\ \text{subject to} & \boldsymbol{g}(\boldsymbol{x}) \le \boldsymbol{0}, \end{array}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is convex, each component of $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^m$ is convex, and $f, \boldsymbol{g} \in \mathcal{C}^1$. Let us call this problem the *primal* problem.

Define the *dual* of the problem above as

$$\begin{array}{ll} \text{maximize} & q(\boldsymbol{\mu}) \\ \text{subject to} & \boldsymbol{\mu} \geq \mathbf{0}, \end{array}$$

where q is defined by

$$q(\boldsymbol{\mu}) = \min_{\boldsymbol{x} \in \mathbb{R}^n} l(\boldsymbol{x}, \boldsymbol{\mu}),$$

with $l(\boldsymbol{x}, \boldsymbol{\mu}) = f(\boldsymbol{x}) + \boldsymbol{\mu}^\top \boldsymbol{g}(\boldsymbol{x})$ the Lagrangian at $\boldsymbol{x}, \boldsymbol{\mu}$.

Prove the following results:

a. If $\boldsymbol{x}_0$ and $\boldsymbol{\mu}_0$ are feasible points in the primal and dual, respectively, then $f(\boldsymbol{x}_0) \geq q(\boldsymbol{\mu}_0)$. This is the *weak duality lemma* for nonlinear programming, analogous to Lemma 17.1.

b. If $\boldsymbol{x}_0$ and $\boldsymbol{\mu}_0$ are feasible points in the primal and dual, and $f(\boldsymbol{x}_0) = q(\boldsymbol{\mu}_0)$, then $\boldsymbol{x}_0$ and $\boldsymbol{\mu}_0$ are optimal solutions to the primal and dual, respectively.

c. If the primal has an optimal (feasible) solution, then so does the dual, and their objective function values are equal. (You may assume regularity.) This is the *duality theorem* for nonlinear programming, analogous to Theorem 17.2.

22.25 Consider the matrix

$$\boldsymbol{M} = \begin{bmatrix} 1 & \gamma & -1 \\ \gamma & 1 & 2 \\ -1 & 2 & 5 \end{bmatrix},$$

where γ is a parameter.

a. Find the Schur complement of $\boldsymbol{M}(1,1)$;

b. Find the Schur complement of $\boldsymbol{M}(2\!:\!3, 2\!:\!3)$ (the bottom-right 2×2 submatrix of $\boldsymbol{M}$, using MATLAB notation).

22.26 Represent the Lyapunov inequality

$$\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} < 0,$$

where

$$\boldsymbol{A} = \begin{bmatrix} 0 & 1 \\ -1 & -2 \end{bmatrix},$$

as a canonical LMI.

22.27 Let $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{R}$ be given matrices such that $\boldsymbol{R} = \boldsymbol{R}^\top > 0$. Suppose that we wish to find a symmetric positive definite matrix $\boldsymbol{P}$ satisfying the following quadratic inequality:

$$\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} + \boldsymbol{P}\boldsymbol{B}\boldsymbol{R}^{-1}\boldsymbol{B}^\top \boldsymbol{P} < 0.$$

Represent this inequality in the form of LMIs. (This inequality should not be confused with the *algebraic Riccati inequality*, which has a negative sign in front of the third term.)

22.28 Let

$$\boldsymbol{A} = \begin{bmatrix} -0.9501 & -0.4860 & -0.4565 \\ -0.2311 & -0.8913 & -0.0185 \\ -0.6068 & -0.7621 & -0.8214 \end{bmatrix}.$$

Write a MATLAB program that finds a matrix $\boldsymbol{P}$ satisfying $0.1\boldsymbol{I}_3 \leq \boldsymbol{P} \leq \boldsymbol{I}_3$ and

$$\boldsymbol{A}^\top \boldsymbol{P} + \boldsymbol{P}\boldsymbol{A} \leq 0.$$