CHAPTER 12

Theory of Constrained Optimization

The second part of this book is about minimizing functions subject to constraints on the variables. A general formulation for these problems is

$$\min_{x \in \mathbb{R}^n} f(x) \qquad \text{subject to} \quad \begin{cases} c_i(x) = 0, & i \in \mathcal{E}, \\ c_i(x) \geq 0, & i \in \mathcal{I}, \end{cases} \tag{12.1}$$

where f and the functions c_i are all smooth, real-valued functions on a subset of $\mathbb{R}^n$, and $\mathcal{I}$ and $\mathcal{E}$ are two finite sets of indices. As before, we call f the *objective* function, while c_i,

$i \in \mathcal{E}$ are the *equality constraints* and c_i, $i \in \mathcal{I}$ are the *inequality constraints*. We define the *feasible set* Ω to be the set of points x that satisfy the constraints; that is,

$$\Omega = \{x \mid c_i(x) = 0, \quad i \in \mathcal{E}; \quad c_i(x) \geq 0, \quad i \in \mathcal{I}\}, \tag{12.2}$$

so that we can rewrite (12.1) more compactly as

$$\min_{x \in \Omega} \; f(x). \tag{12.3}$$

In this chapter we derive mathematical characterizations of the solutions of (12.3). As in the unconstrained case, we discuss optimality conditions of two types. *Necessary* conditions are conditions that must be satisfied by any solution point (under certain assumptions). *Sufficient* conditions are those that, if satisfied at a certain point x^*, guarantee that x^* is in fact a solution.

For the unconstrained optimization problem of Chapter 2, the optimality conditions were as follows:

> Necessary conditions: Local unconstrained minimizers have $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ positive semidefinite.
>
> Sufficient conditions: Any point x^* at which $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ is positive definite is a strong local minimizer of f.

In this chapter, we derive analogous conditions to characterize the solutions of constrained optimization problems.

LOCAL AND GLOBAL SOLUTIONS

We have seen already that global solutions are difficult to find even when there are no constraints. The situation may be improved when we add constraints, since the feasible set might exclude many of the local minima and it may be comparatively easy to pick the global minimum from those that remain. However, constraints can also make things more difficult. As an example, consider the problem

$$\min \, (x_2 + 100)^2 + 0.01x_1^2, \qquad \text{subject to } x_2 - \cos x_1 \geq 0, \tag{12.4}$$

illustrated in Figure 12.1. Without the constraint, the problem has the unique solution $(0, -100)^T$. With the constraint, there are local solutions near the points

$$x^{(k)} = (k\pi, -1)^T, \qquad \text{for } k = \pm 1, \pm 3, \pm 5, \ldots.$$

Definitions of the different types of local solutions are simple extensions of the corresponding definitions for the unconstrained case, except that now we restrict consideration to the *feasible* points in the neighborhood of x^*. We have the following definition.

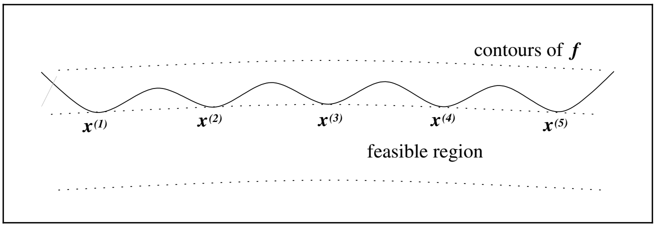

Figure 12.1 Constrained problem with many isolated local solutions.

A vector x^* is a *local solution* of the problem (12.3) if $x^* \in \Omega$ and there is a neighborhood $\mathcal{N}$ of x^* such that $f(x) \geq f(x^*)$ for $x \in \mathcal{N} \cap \Omega$.

Similarly, we can make the following definitions:

A vector x^* is a *strict local solution* (also called a *strong local solution*) if $x^* \in \Omega$ and there is a neighborhood $\mathcal{N}$ of x^* such that $f(x) > f(x^*)$ for all $x \in \mathcal{N} \cap \Omega$ with $x \neq x^*$.

A point x^* is an *isolated local solution* if $x^* \in \Omega$ and there is a neighborhood $\mathcal{N}$ of x^* such that x^* is the only local solution in $\mathcal{N} \cap \Omega$.

Note that isolated local solutions are strict, but that the reverse is not true (see Exercise 12.2).

SMOOTHNESS

Smoothness of objective functions and constraints is an important issue in characterizing solutions, just as in the unconstrained case. It ensures that the objective function and the constraints all behave in a reasonably predictable way and therefore allows algorithms to make good choices for search directions.

We saw in Chapter 2 that graphs of nonsmooth functions contain "kinks" or "jumps" where the smoothness breaks down. If we plot the feasible region for any given constrained optimization problem, we usually observe many kinks and sharp edges. Does this mean that the constraint functions that describe these regions are nonsmooth? The answer is often no, because the nonsmooth boundaries can often be described by a collection of smooth constraint functions. Figure 12.2 shows a diamond-shaped feasible region in $\mathbb{R}^2$ that could be described by the single nonsmooth constraint

$$\|x\|_1 = |x_1| + |x_2| \leq 1. \tag{12.5}$$

It can also be described by the following set of smooth (in fact, linear) constraints:

$$x_1 + x_2 \leq 1, \quad x_1 - x_2 \leq 1, \quad -x_1 + x_2 \leq 1, \quad -x_1 - x_2 \leq 1. \tag{12.6}$$

Each of the four constraints represents one edge of the feasible polytope. In general, the constraint functions are chosen so that each one represents a smooth piece of the boundary of Ω.

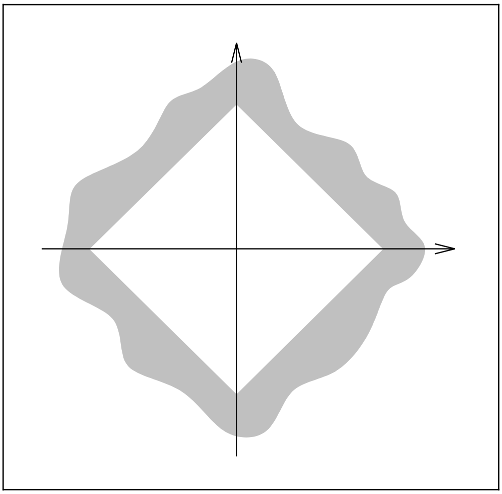

Figure 12.2
A feasible region with a nonsmooth boundary can be described by smooth constraints.

Nonsmooth, unconstrained optimization problems can sometimes be reformulated as smooth constrained problems. An example is the unconstrained minimization of a function

$$f(x) = \max(x^2, x), \tag{12.7}$$

which has kinks at $x = 0$ and $x = 1$, and the solution at $x^* = 0$. We obtain a smooth, constrained formulation of this problem by adding an artificial variable t and writing

$$\min t \qquad \text{s.t.} \qquad t \geq x, \quad t \geq x^2. \tag{12.8}$$

Reformulation techniques such as (12.6) and (12.8) are used often in cases where f is a maximum of a collection of functions or when f is a 1-norm or ∞-norm of a vector function.

In the examples above we expressed inequality constraints in a slightly different way from the form $c_i(x) \geq 0$ that appears in the definition (12.1). However, any collection of inequality constraints with $\geq$ and $\leq$ and nonzero right-hand-sides can be expressed in the form $c_i(x) \geq 0$ by simple rearrangement of the inequality.

12.1 EXAMPLES

To introduce the basic principles behind the characterization of solutions of constrained optimization problems, we work through three simple examples. The discussion here is informal; the ideas introduced will be made rigorous in the sections that follow.

We start by noting one important item of terminology that recurs throughout the rest of the book.

Definition 12.1.

The active set $\mathcal{A}(x)$ *at any feasible* x *consists of the equality constraint indices from* $\mathcal{E}$ *together with the indices of the inequality constraints* i *for which* $c_i(x) = 0$*; that is,*

$$\mathcal{A}(x) = \mathcal{E} \cup \{i \in \mathcal{I} \mid c_i(x) = 0\}.$$

At a feasible point x, the inequality constraint $i \in \mathcal{I}$ is said to be *active* if $c_i(x) = 0$ and *inactive* if the strict inequality $c_i(x) > 0$ is satisfied.

A SINGLE EQUALITY CONSTRAINT

❑ EXAMPLE 12.1

Our first example is a two-variable problem with a single equality constraint:

$$\min\ x_1 + x_2 \qquad \text{s.t.} \qquad x_1^2 + x_2^2 - 2 = 0 \tag{12.9}$$

(see Figure 12.3). In the language of (12.1), we have $f(x) = x_1 + x_2$, $\mathcal{I} = \emptyset$, $\mathcal{E} = \{1\}$, and $c_1(x) = x_1^2 + x_2^2 - 2$. We can see by inspection that the feasible set for this problem is the circle of radius $\sqrt{2}$ centered at the origin—just the boundary of this circle, not its interior. The solution x^* is obviously $(-1, -1)^T$. From any other point on the circle, it is easy to find a way to move that *stays feasible* (that is, remains on the circle) while *decreasing* f. For instance, from the point $x = (\sqrt{2}, 0)^T$ any move in the clockwise direction around the circle has the desired effect.

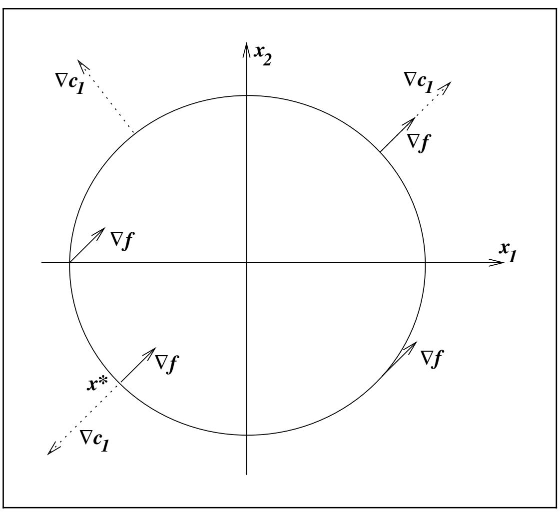

Figure 12.3
Problem (12.9), showing constraint and function gradients at various feasible points.

We also see from Figure 12.3 that at the solution x^*, the *constraint normal* $\nabla c_1(x^*)$ is parallel to $\nabla f(x^*)$. That is, there is a scalar λ_1^* (in this case $\lambda_1^* = -1/2$) such that

$$\nabla f(x^*) = \lambda_1^* \nabla c_1(x^*). \tag{12.10}$$

❐

We can derive (12.10) by examining first-order Taylor series approximations to the objective and constraint functions. To retain feasibility with respect to the function $c_1(x) = 0$, we require any small (but nonzero) step s to satisfy that $c_1(x+s) = 0$; that is,

$$0 = c_1(x+s) \approx c_1(x) + \nabla c_1(x)^T s = \nabla c_1(x)^T s. \tag{12.11}$$

Hence, the step s retains feasibility with respect to c_1, to first order, when it satisfies

$$\nabla c_1(x)^T s = 0. \tag{12.12}$$

Similarly, if we want s to produce a decrease in f, we would have so that

$$0 > f(x+s) - f(x) \approx \nabla f(x)^T s,$$

or, to first order,

$$\nabla f(x)^T s < 0. \tag{12.13}$$

Existence of a small step s that satisfies both (12.12) and (12.13) strongly suggests existence of a direction d (where the size of d is *not* small; we could have $d \approx s/\|s\|$ to ensure that the norm of d is close to 1) with the same properties, namely

$$\nabla c_1(x)^T d = 0 \quad \text{and} \quad \nabla f(x)^T d < 0. \tag{12.14}$$

If, on the other hand, there is *no* direction d with the properties (12.14), then is it likely that we cannot find a small step s with the properties (12.12) and (12.13). In this case, x^* would appear to be a local minimizer.

By drawing a picture, the reader can check that the only way that a d satisfying (12.14) does *not* exist is if $\nabla f(x)$ and $\nabla c_1(x)$ are parallel, that is, if the condition $\nabla f(x) = \lambda_1 \nabla c_1(x)$ holds at x, for some scalar λ_1. If in fact $\nabla f(x)$ and $\nabla c_1(x)$ are *not* parallel, we can set

$$\bar{d} = -\left(I - \frac{\nabla c_1(x) \nabla c_1(x)^T}{\|\nabla c_1(x)\|^2} \right) \nabla f(x); \qquad d = \frac{\bar{d}}{\|\bar{d}\|}. \tag{12.15}$$

It is easy to verify that this d satisfies (12.14).

By introducing the *Lagrangian function*

$$\mathcal{L}(x, \lambda_1) = f(x) - \lambda_1 c_1(x), \tag{12.16}$$

and noting that $\nabla_x \mathcal{L}(x, \lambda_1) = \nabla f(x) - \lambda_1 \nabla c_1(x)$, we can state the condition (12.10) equivalently as follows: At the solution x^*, there is a scalar λ_1^* such that

$$\nabla_x \mathcal{L}(x^*, \lambda_1^*) = 0. \tag{12.17}$$

This observation suggests that we can search for solutions of the equality-constrained problem (12.9) by seeking stationary points of the Lagrangian function. The scalar quantity λ_1 in (12.16) is called a *Lagrange multiplier* for the constraint $c_1(x) = 0$.

Though the condition (12.10) (equivalently, (12.17)) appears to be *necessary* for an optimal solution of the problem (12.9), it is clearly not *sufficient*. For instance, in Example 12.1, condition (12.10) is satisfied at the point $x = (1, 1)^T$ (with $\lambda_1 = \frac{1}{2}$), but this point is obviously not a solution—in fact, it *maximizes* the function f on the circle. Moreover, in the case of equality-constrained problems, we cannot turn the condition (12.10) into a sufficient condition simply by placing some restriction on the sign of λ_1. To see this, consider replacing the constraint $x_1^2 + x_2^2 - 2 = 0$ by its negative $2 - x_1^2 - x_2^2 = 0$ in Example 12.1. The solution of the problem is not affected, but the value of λ_1^* that satisfies the condition (12.10) changes from $\lambda_1^* = -\frac{1}{2}$ to $\lambda_1^* = \frac{1}{2}$.

A SINGLE INEQUALITY CONSTRAINT

❑ EXAMPLE 12.2

This is a slight modification of Example 12.1, in which the equality constraint is replaced by an inequality. Consider

$$\min\ x_1 + x_2 \qquad \text{s.t.} \qquad 2 - x_1^2 - x_2^2 \geq 0, \tag{12.18}$$

for which the feasible region consists of the circle of problem (12.9) and its interior (see Figure 12.4). Note that the constraint normal ∇c_1 points toward the interior of the feasible region at each point on the boundary of the circle. By inspection, we see that the solution is still $(-1, -1)^T$ and that the condition (12.10) holds for the value $\lambda_1^* = \frac{1}{2}$. However, this inequality-constrained problem differs from the equality-constrained problem (12.9) of Example 12.1 in that the sign of the Lagrange multiplier plays a significant role, as we now argue.

❐

As before, we conjecture that a given feasible point x is *not* optimal if we can find a small step s that both retains feasibility and decreases the objective function f to first order. The main difference between problems (12.9) and (12.18) comes in the handling of the feasibility condition. As in (12.13), the step s improves the objective function, to first order, if $\nabla f(x)^T s < 0$. Meanwhile, s retains feasibility if

$$0 \leq c_1(x+s) \approx c_1(x) + \nabla c_1(x)^T s,$$

so, to first order, feasibility is retained if

$$c_1(x) + \nabla c_1(x)^T s \geq 0. \tag{12.19}$$

In determining whether a step s exists that satisfies both (12.13) and (12.19), we consider the following two cases, which are illustrated in Figure 12.4.

Case I: Consider first the case in which x lies *strictly inside* the circle, so that the strict inequality $c_1(x) > 0$ holds. In this case, *any* step vector s satisfies the condition (12.19), provided only that its length is sufficiently small. In fact, whenever $\nabla f(x) \neq 0$, we can obtain a step s that satisfies both (12.13) and (12.19) by setting

$$s = -\alpha \nabla f(x),$$

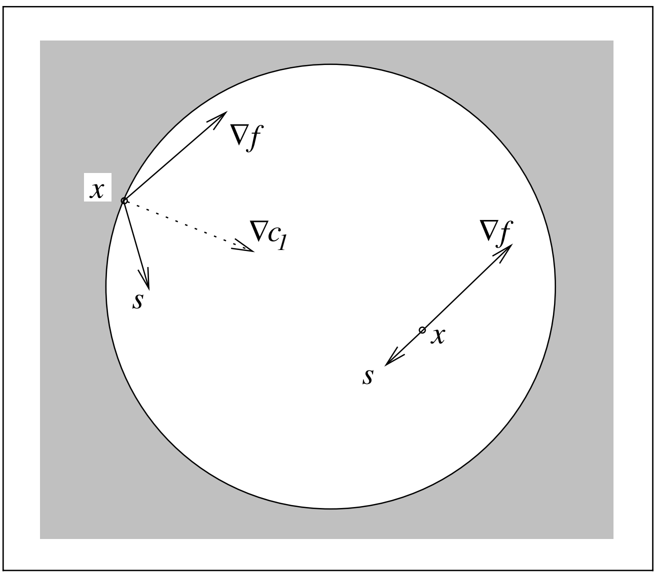

Figure 12.4 Improvement directions s from two feasible points x for the problem (12.18) at which the constraint is active and inactive, respectively.

for any positive scalar α sufficiently small. However, this definition does not give a step s with the required properties when

$$\nabla f(x) = 0, \tag{12.20}$$

Case II: Consider now the case in which x lies on the boundary of the circle, so that $c_1(x) = 0$. The conditions (12.13) and (12.19) therefore become

$$\nabla f(x)^T s < 0, \quad \nabla c_1(x)^T s \geq 0.$$

The first of these conditions defines an open half-space, while the second defines a closed half-space, as illustrated in Figure 12.5. It is clear from this figure that the intersection of these two regions is empty only when $\nabla f(x)$ and $\nabla c_1(x)$ point in the same direction, that is, when

$$\nabla f(x) = \lambda_1 \nabla c_1(x), \quad \text{for some } \lambda_1 \geq 0. \tag{12.21}$$

Note that the sign of the multiplier is significant here. If (12.10) were satisfied with a *negative* value of λ_1, then $\nabla f(x)$ and $\nabla c_1(x)$ would point in opposite directions, and we see from Figure 12.5 that the set of directions that satisfy both (12.13) and (12.19) would make up an entire open half-plane.

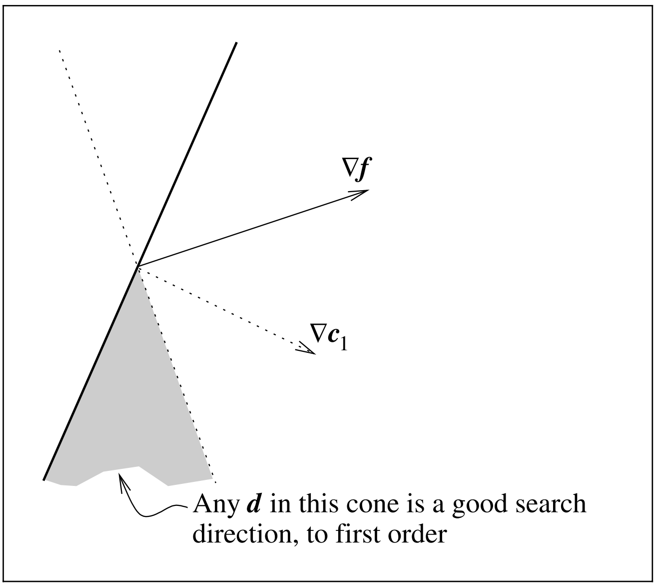

Figure 12.5 A direction d that satisfies both (12.13) and (12.19) lies in the intersection of a closed half-plane and an open half-plane.

The optimality conditions for both cases I and II can again be summarized neatly with reference to the Lagrangian function $\mathcal{L}$ defined in (12.16). When no first-order feasible descent direction exists at some point x^*, we have that

$$\nabla_x \mathcal{L}(x^*, \lambda_1^*) = 0, \qquad \text{for some } \lambda_1^* \geq 0, \tag{12.22}$$

where we also require that

$$\lambda_1^* c_1(x^*) = 0. \tag{12.23}$$

Condition (12.23) is known as a *complementarity condition*; it implies that the Lagrange multiplier λ_1 can be strictly positive *only when the corresponding constraint c_1 is active.* Conditions of this type play a central role in constrained optimization, as we see in the sections that follow. In case I, we have that $c_1(x^*) > 0$, so (12.23) requires that $\lambda_1^* = 0$. Hence, (12.22) reduces to $\nabla f(x^*) = 0$, as required by (12.20). In case II, (12.23) allows λ_1^* to take on a nonnegative value, so (12.22) becomes equivalent to (12.21).

TWO INEQUALITY CONSTRAINTS

❑ EXAMPLE 12.3

Suppose we add an extra constraint to the problem (12.18) to obtain

$$\min x_1 + x_2 \qquad \text{s.t.} \qquad 2 - x_1^2 - x_2^2 \geq 0, \quad x_2 \geq 0, \tag{12.24}$$

for which the feasible region is the half-disk illustrated in Figure 12.6. It is easy to see that the solution lies at $(-\sqrt{2}, 0)^T$, a point at which both constraints are active. By repeating the arguments for the previous examples, we would expect a direction d of first-order feasible descent to satisfy

$$\nabla c_i(x)^T d \geq 0, \qquad i \in \mathcal{I} = \{1, 2\}, \qquad \nabla f(x)^T d < 0. \tag{12.25}$$

However, it is clear from Figure 12.6 that no such direction can exist when $x = (-\sqrt{2}, 0)^T$. The conditions $\nabla c_i(x)^T d \geq 0$, $i = 1, 2$, are both satisfied only if d lies in the quadrant defined by $\nabla c_1(x)$ and $\nabla c_2(x)$, but it is clear by inspection that all vectors d in this quadrant satisfy $\nabla f(x)^T d \geq 0$.

Let us see how the Lagrangian and its derivatives behave for the problem (12.24) and the solution point $(-\sqrt{2}, 0)^T$. First, we include an additional term $\lambda_i c_i(x)$ in the Lagrangian for each additional constraint, so the definition of $\mathcal{L}$ becomes

$$\mathcal{L}(x, \lambda) = f(x) - \lambda_1 c_1(x) - \lambda_2 c_2(x),$$

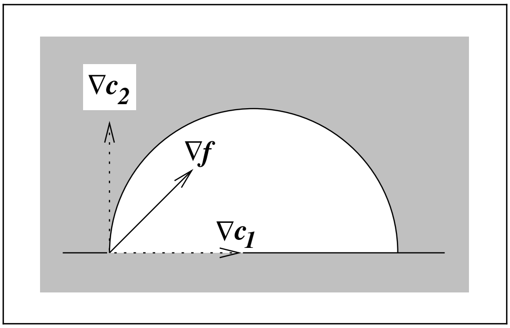

Figure 12.6
Problem (12.24), illustrating the gradients of the active constraints and objective at the solution.

where $\lambda = (\lambda_1, \lambda_2)^T$ is the vector of Lagrange multipliers. The extension of condition (12.22) to this case is

$$\nabla_x \mathcal{L}(x^*, \lambda^*) = 0, \qquad \text{for some } \lambda^* \geq 0, \tag{12.26}$$

where the inequality $\lambda^* \geq 0$ means that all components of λ^* are required to be nonnegative. By applying the complementarity condition (12.23) to both inequality constraints, we obtain

$$\lambda_1^* c_1(x^*) = 0, \qquad \lambda_2^* c_2(x^*) = 0. \tag{12.27}$$

When $x^* = (-\sqrt{2}, 0)^T$, we have

$$\nabla f(x^*) = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \nabla c_1(x^*) = \begin{bmatrix} 2\sqrt{2} \\ 0 \end{bmatrix}, \quad \nabla c_2(x^*) = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

so that it is easy to verify that $\nabla_x \mathcal{L}(x^*, \lambda^*) = 0$ when we select λ^* as follows:

$$\lambda^* = \begin{bmatrix} 1/(2\sqrt{2}) \\ 1 \end{bmatrix}.$$

Note that both components of λ^* are positive, so that (12.26) is satisfied.

We consider now some other feasible points that are *not* solutions of (12.24), and examine the properties of the Lagrangian and its gradient at these points.

For the point $x = (\sqrt{2}, 0)^T$, we again have that both constraints are active (see Figure 12.7). However, it s easy to identify vectors d that satisfies (12.25): $d = (-1, 0)^T$ is one such vector (there are many others). For this value of x it is easy to verify that the condition $\nabla_x \mathcal{L}(x, \lambda) = 0$ is satisfied only when $\lambda = (-1/(2\sqrt{2}), 1)^T$. Note that the first component λ_1 is negative, so that the conditions (12.26) are not satisfied at this point.

Finally, we consider the point $x = (1, 0)^T$, at which only the second constraint c_2 is active. Since any small step s away from this point will continue to satisfy $c_1(x + s) > 0$, we need to consider only the behavior of c_2 and f in determining whether s is indeed a feasible

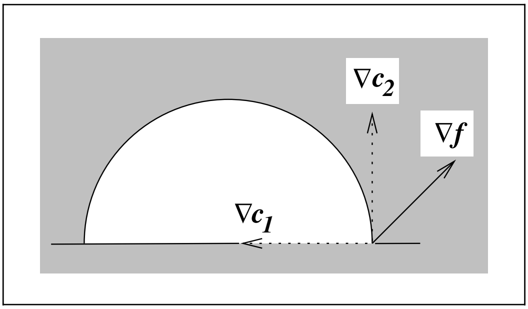

Figure 12.7
Problem (12.24), illustrating the gradients of the active constraints and objective at a nonoptimal point.

descent step. Using the same reasoning as in the earlier examples, we find that the direction of feasible descent d must satisfy

$$\nabla c_2(x)^T d \geq 0, \quad \nabla f(x)^T d < 0. \tag{12.28}$$

By noting that

$$\nabla f(x) = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \nabla c_2(x) = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

it is easy to verify that the vector $d = \left(-\frac{1}{2}, \frac{1}{4}\right)^T$ satisfies (12.28) and is therefore a descent direction.

To show that optimality conditions (12.26) and (12.27) fail, we note first from (12.27) that since $c_1(x) > 0$, we must have $\lambda_1 = 0$. Therefore, in trying to satisfy $\nabla_x \mathcal{L}(x, \lambda) = 0$, we are left to search for a value λ_2 such that $\nabla f(x) - \lambda_2 \nabla c_2(x) = 0$. No such λ_2 exists, and thus this point fails to satisfy the optimality conditions. ❐

12.2 TANGENT CONE AND CONSTRAINT QUALIFICATIONS

In this section we define the tangent cone $T_\Omega(x^*)$ to the closed convex set Ω at a point $x^* \in \Omega$, and also the set $\mathcal{F}(x^*)$ of first-order feasible directions at x^*. We also discuss *constraint qualifications.* In the previous section, we determined whether or not it was possible to take a feasible descent step away from a given feasible point x by examining the first derivatives of f and the constraint functions c_i. We used the first-order Taylor series expansion of these functions about x to form an approximate problem in which both objective and constraints are linear. This approach makes sense, however, only when the linearized approximation captures the essential geometric features of the feasible set near the point x in question. If, near x, the linearization is fundamentally different from the feasible set (for instance, it is an entire plane, while the feasible set is a single point) then we cannot expect the linear approximation to yield useful information about the original problem. Hence, we need to make assumptions about the nature of the constraints c_i that are active at x to ensure that the linearized approximation is similar to the feasible set, near x. Constraint qualifications are assumptions that ensure similarity of the constraint set Ω and its linearized approximation, in a neighborhood of x^*.

Given a feasible point x, we call $\{z_k\}$ a *feasible sequence approaching* x if $z_k \in \Omega$ for all k sufficiently large and $z_k \to x$.

Later, we characterize a local solution of (12.1) as a point x at which all feasible sequences approaching x have the property that $f(z_k) \geq f(x)$ for all k sufficiently large, and we will derive practical, verifiable conditions under which this property holds. We lay the groundwork in this section by characterizing the directions in which we can step away from x while remaining feasible.

A *tangent* is a limiting direction of a feasible sequence.

Definition 12.2.

The vector d is said to be a tangent *(or* tangent vector*) to Ω at a point x if there are a feasible sequence $\{z_k\}$ approaching x and a sequence of positive scalars $\{t_k\}$ with $t_k \to 0$ such that*

$$\lim_{k \to \infty} \frac{z_k - x}{t_k} = d. \tag{12.29}$$

The set of all tangents to Ω at x^* is called the tangent cone *and is denoted by $T_\Omega(x^*)$.*

It is easy to see that the tangent cone is indeed a cone, according to the definition (A.36). If d is a tangent vector with corresponding sequences $\{z_k\}$ and $\{t_k\}$, then by replacing each t_k by $\alpha^{-1} t_k$, for any $\alpha > 0$, we find that $\alpha d \in T_\Omega(x^*)$ also. We obtain that $0 \in T_\Omega(x)$ by setting $z_k \equiv x$ in the definition of feasible sequence.

We turn now to the linearized feasible direction set, which we define as follows.

Definition 12.3.

Given a feasible point x and the active constraint set $\mathcal{A}(x)$ of Definition 12.1, the set of linearized feasible directions $\mathcal{F}(x)$ *is*

$$\mathcal{F}(x) = \left\{ d \;\middle|\; \begin{array}{ll} d^T \nabla c_i(x) = 0, & \text{for all } i \in \mathcal{E}, \\ d^T \nabla c_i(x) \geq 0, & \text{for all } i \in \mathcal{A}(x) \cap \mathcal{I} \end{array} \right\}.$$

As with the tangent cone, it is easy to verify that $\mathcal{F}(x)$ is a cone, according to the definition (A.36).

It is important to note that the definition of tangent cone does not rely on the algebraic specification of the set Ω, only on its geometry. The linearized feasible direction set does, however, depend on the definition of the constraint functions c_i, $i \in \mathcal{E} \cup \mathcal{I}$.

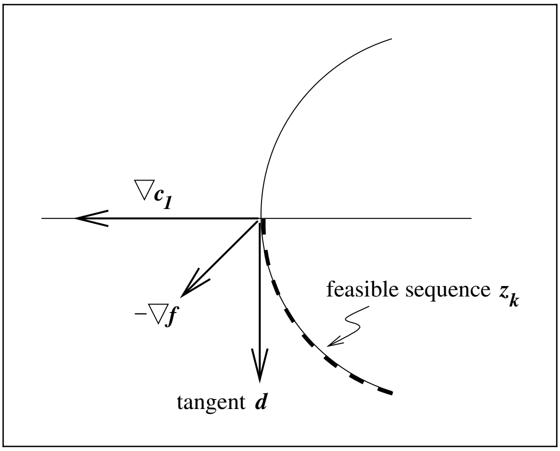

Figure 12.8
Constraint normal, objective gradient, and feasible sequence for problem (12.9).

We illustrate the tangent cone and the linearized feasible direction set by revisiting Examples 12.1 and 12.2.

❑ EXAMPLE 12.4 (EXAMPLE 12.1, REVISITED)

Figure 12.8 shows the problem (12.9), the equality-constrained problem in which the feasible set is a circle of radius $\sqrt{2}$, near the nonoptimal point $x = (-\sqrt{2}, 0)^T$. The figure also shows a feasible sequence approaching x. This sequence could be defined analytically by the formula

$$z_k = \begin{bmatrix} -\sqrt{2 - 1/k^2} \\ -1/k \end{bmatrix}. \tag{12.30}$$

By choosing $t_k = \|z_k - x\|$, we find that $d = (0, -1)^T$ is a tangent. Note that the objective function $f(x) = x_1 + x_2$ increases as we move along the sequence (12.30); in fact, we have $f(z_{k+1}) > f(z_k)$ for all $k = 2, 3, \ldots$. It follows that $f(z_k) < f(x)$ for $k = 2, 3, \ldots$, so x cannot be a solution of (12.9).

Another feasible sequence is one that approaches $x = (-\sqrt{2}, 0)^T$ from the opposite direction. Its elements are defined by

$$z_k = \begin{bmatrix} -\sqrt{2 - 1/k^2} \\ 1/k \end{bmatrix}.$$

It is easy to show that f *decreases* along this sequence and that the tangents corresponding to this sequence are $d = (0, \alpha)^T$. In summary, the tangent cone at $x = (-\sqrt{2}, 0)^T$ is $\{(0, d_2)^T \mid d_2 \in \mathbb{R}\}$.

For the definition (12.9) of this set, and Definition 12.3, we have that $d = (d_1, d_2)^T \in \mathcal{F}(x)$ if

$$0 = \nabla c_1(x)^T d = \begin{bmatrix} 2x_1 \\ 2x_2 \end{bmatrix}^T \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = -2\sqrt{2} d_1.$$

Therefore, we obtain $\mathcal{F}(x) = \{(0, d_2)^T \mid d_2 \in \mathbb{R}\}$. In this case, we have $T_\Omega(x) = \mathcal{F}(x)$.

Suppose that the feasible set is defined instead by the formula

$$\Omega = \{x \mid c_1(x) = 0\}, \quad \text{where } c_1(x) = (x_1^2 + x_2^2 - 2)^2 = 0. \tag{12.31}$$

(Note that Ω is the same, but its algebraic specification has changed.) The vector d belongs to the linearized feasible set if

$$0 = \nabla c_1(x)^T d = \begin{bmatrix} 4(x_1^2 + x_2^2 - 2)x_1 \\ 4(x_1^2 + x_2^2 - 2)x_2 \end{bmatrix}^T \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}^T \begin{bmatrix} d_1 \\ d_2 \end{bmatrix},$$

which is true for all $(d_1, d_2)^T$. Hence, we have $\mathcal{F}(x) = \mathbb{R}^2$, so for this algebraic specification of Ω, the tangent cone and linearized feasible sets differ. ❐

❑ EXAMPLE 12.5 (EXAMPLE 12.2, REVISITED)

We now reconsider problem (12.18) in Example 12.2. The solution $x = (-1, -1)^T$ is the same as in the equality-constrained case, but there is a much more extensive collection of feasible sequences that converge to any given feasible point (see Figure 12.9).

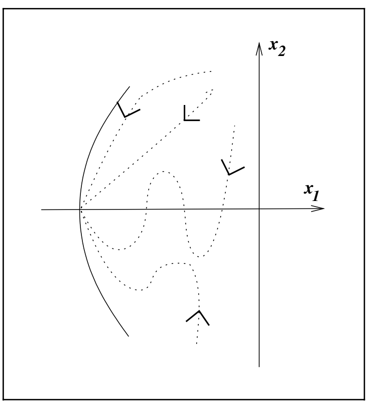

Figure 12.9
Feasible sequences converging to a particular feasible point for the region defined by $x_1^2 + x_2^2 \le 2$.

From the point $x = (-\sqrt{2}, 0)^T$, the various feasible sequences defined above for the equality-constrained problem are still feasible for (12.18). There are also infinitely many feasible sequences that converge to $x = (-\sqrt{2}, 0)^T$ along a straight line from the interior of the circle. These sequences have the form

$$z_k = (-\sqrt{2}, 0)^T + (1/k)w,$$

where w is any vector whose first component is positive ($w_1 > 0$). The point z_k is feasible provided that $\|z_k\| \le \sqrt{2}$, that is,

$$(-\sqrt{2} + w_1/k)^2 + (w_2/k)^2 \le 2,$$

which is true when $k \ge (w_1^2 + w_2^2)/(2\sqrt{2}w_1)$. In addition to these straight-line feasible sequences, we can also define an infinite variety of sequences that approach $(-\sqrt{2}, 0)^T$ along a curve from the interior of the circle. To summarize, the tangent cone to this set at $(-\sqrt{2}, 0)^T$ is $\{(w_1, w_2)^T \mid w_1 \ge 0\}$.

For the definition (12.18) of this feasible set, we have from Definition 12.3 that $d \in \mathcal{F}(x)$ if

$$0 \le \nabla c_1(x)^T d = \begin{bmatrix} -2x_1 \\ -2x_2 \end{bmatrix}^T \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = 2\sqrt{2}d_1.$$

Hence, we obtain $\mathcal{F}(x) = T_\Omega(x)$ for this particular algebraic specification of the feasible set.

❐

Constraint qualifications are conditions under which the linearized feasible set $\mathcal{F}(x)$ is similar to the tangent cone $T_\Omega(x)$. In fact, most constraint qualifications ensure that these two sets are identical. As mentioned earlier, these conditions ensure that the $\mathcal{F}(x)$, which is constructed by linearizing the algebraic description of the set Ω at x, captures the essential geometric features of the set Ω in the vicinity of x, as represented by $T_\Omega(x)$.

Revisiting Example 12.4, we see that both $T_\Omega(x)$ and $\mathcal{F}(x)$ consist of the vertical axis, which is qualitatively similar to the set $\Omega - \{x\}$ in the neighborhood of x. As a further example, consider the constraints

$$c_1(x) = 1 - x_1^2 - (x_2 - 1)^2 \ge 0, \qquad c_2(x) = -x_2 \ge 0, \tag{12.32}$$

for which the feasible set is the single point $\Omega = \{(0, 0)^T\}$ (see Figure 12.10). For this point $x = (0, 0)^T$, it is obvious that that tangent cone is $T_\Omega(x) = \{(0, 0)^T\}$, since all feasible sequences approaching x must have $z_k = x = (0, 0)^T$ for all k sufficiently large. Moreover, it is easy to show that linearized approximation to the feasible set $\mathcal{F}(x)$ is

$$\mathcal{F}(x^*) = \{(d_1, 0)^T \mid d_1 \in \mathbb{R}\},$$

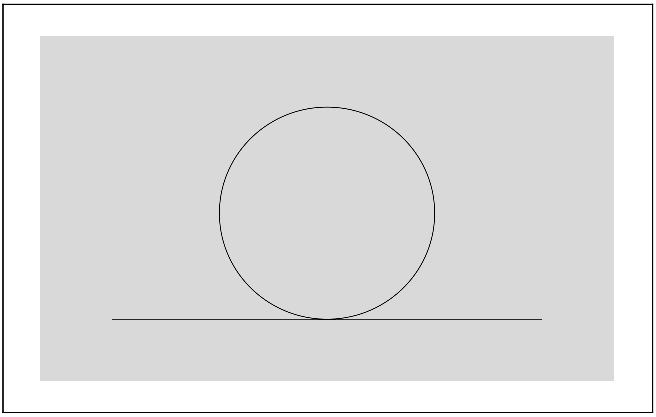

Figure 12.10 Problem (12.32), for which the feasible set is the single point of intersection between circle and line.

that is, the entire horizontal axis. In this case, the linearized feasible direction set does not capture the geometry of the feasible set, so constraint qualifications are not satisfied.

The constraint qualification most often used in the design of algorithms is the subject of the next definition.

Definition 12.4 (LICQ).

Given the point x *and the active set* $\mathcal{A}(x)$ *defined in Definition 12.1, we say that the* linear independence constraint qualification (LICQ) *holds if the set of active constraint gradients* $\{\nabla c_i(x), i \in \mathcal{A}(x)\}$ *is linearly independent.*

Note that this condition is *not* satisfied for the examples (12.32) and (12.31). In general, if LICQ holds, none of the active constraint gradients can be zero. We mention other constraint qualifications in Section 12.6.

12.3 FIRST-ORDER OPTIMALITY CONDITIONS

In this section, we state first-order necessary conditions for x^* to be a local minimizer and show how these conditions are satisfied on a small example. The proof of the result is presented in subsequent sections.

As a preliminary to stating the necessary conditions, we define the Lagrangian function for the general problem (12.1).

$$\mathcal{L}(x, \lambda) = f(x) - \sum_{i \in \mathcal{E} \cup \mathcal{I}} \lambda_i c_i(x). \tag{12.33}$$

(We had previously defined special cases of this function for the examples of Section 12.1.)

The necessary conditions defined in the following theorem are called *first-order conditions* because they are concerned with properties of the gradients (first-derivative vectors) of the objective and constraint functions. These conditions are the foundation for many of the algorithms described in the remaining chapters of the book.

Theorem 12.1 (First-Order Necessary Conditions).

Suppose that x^* is a local solution of (12.1), that the functions f and c_i in (12.1) are continuously differentiable, and that the LICQ holds at x^*. Then there is a Lagrange multiplier vector λ^*, with components λ_i^*, $i \in \mathcal{E} \cup \mathcal{I}$, such that the following conditions are satisfied at (x^*, λ^*)

$$\begin{aligned} \nabla_x \mathcal{L}(x^*, \lambda^*) &= 0, && && (12.34a)\\ c_i(x^*) &= 0, && \text{for all } i \in \mathcal{E}, && (12.34b)\\ c_i(x^*) &\geq 0, && \text{for all } i \in \mathcal{I}, && (12.34c)\\ \lambda_i^* &\geq 0, && \text{for all } i \in \mathcal{I}, && (12.34d)\\ \lambda_i^* c_i(x^*) &= 0, && \text{for all } i \in \mathcal{E} \cup \mathcal{I}. && (12.34e) \end{aligned}$$

The conditions (12.34) are often known as the *Karush–Kuhn–Tucker conditions*, or *KKT conditions* for short. The conditions (12.34e) are *complementarity conditions*; they imply that either constraint i is active or $\lambda_i^* = 0$, or possibly both. In particular, the Lagrange multipliers corresponding to inactive inequality constraints are zero, we can omit the terms for indices $i \notin \mathcal{A}(x^*)$ from (12.34a) and rewrite this condition as

$$0 = \nabla_x \mathcal{L}(x^*, \lambda^*) = \nabla f(x^*) - \sum_{i \in \mathcal{A}(x^*)} \lambda_i^* \nabla c_i(x^*). \qquad (12.35)$$

A special case of complementarity is important and deserves its own definition.

Definition 12.5 (Strict Complementarity).

Given a local solution x^* of (12.1) and a vector λ^* satisfying (12.34), we say that the strict complementarity condition *holds if* exactly one *of λ_i^* and $c_i(x^*)$ is zero for each index $i \in \mathcal{I}$. In other words, we have that $\lambda_i^* > 0$ for each $i \in \mathcal{I} \cap \mathcal{A}(x^*)$.*

Satisfaction of the strict complementarity property usually makes it easier for algorithms to determine the active set $\mathcal{A}(x^*)$ and converge rapidly to the solution x^*.

For a given problem (12.1) and solution point x^*, there may be many vectors λ^* for which the conditions (12.34) are satisfied. When the LICQ holds, however, the optimal λ^* is unique (see Exercise 12.17).

The proof of Theorem 12.1 is quite complex, but it is important to our understanding of constrained optimization, so we present it in the next section. First, we illustrate the KKT conditions with another example.

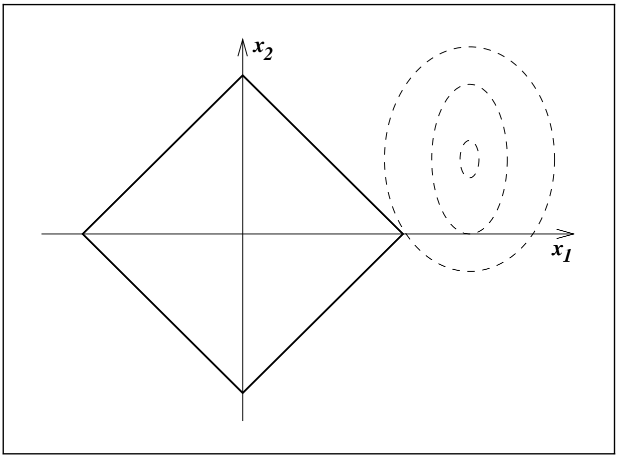

Figure 12.11 Inequality-constrained problem (12.36) with solution at $(1, 0)^T$.

❑ EXAMPLE 12.6

Consider the feasible region illustrated in Figure 12.2 and described by the four constraints (12.6). By restating the constraints in the standard form of (12.1) and including an objective function, the problem becomes

$$\min_x \left(x_1 - \frac{3}{2}\right)^2 + \left(x_2 - \frac{1}{2}\right)^4 \quad \text{s.t.} \quad \begin{bmatrix} 1 - x_1 - x_2 \\ 1 - x_1 + x_2 \\ 1 + x_1 - x_2 \\ 1 + x_1 + x_2 \end{bmatrix} \geq 0. \tag{12.36}$$

It is fairly clear from Figure 12.11 that the solution is $x^* = (1, 0)^T$. The first and second constraints in (12.36) are active at this point. Denoting them by c_1 and c_2 (and the inactive constraints by c_3 and c_4), we have

$$\nabla f(x^*) = \begin{bmatrix} -1 \\ 1 \\ -\frac{1}{2} \end{bmatrix}, \quad \nabla c_1(x^*) = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \quad \nabla c_2(x^*) = \begin{bmatrix} -1 \\ 1 \end{bmatrix}.$$

Therefore, the KKT conditions (12.34a)–(12.34e) are satisfied when we set

$$\lambda^* = \left(\tfrac{3}{4}, \tfrac{1}{4}, 0, 0\right)^T.$$

❐

12.4 FIRST-ORDER OPTIMALITY CONDITIONS: PROOF

We now develop a proof of Theorem 12.1. A number of key subsidiary results are required, so the development is quite long. However, a complete treatment is worthwhile, since these results are so fundamental to the field of optimization.

RELATING THE TANGENT CONE AND THE FIRST-ORDER FEASIBLE DIRECTION SET

The following key result uses a constraint qualification (LICQ) to relate the tangent cone of Definition 12.2 to the set $\mathcal{F}$ of first-order feasible directions of Definition 12.3. In the proof below and in later results, we use the notation $A(x^*)$ to represent the matrix whose rows are the active constraint gradients at the optimal point, that is,

$$A(x^*)^T = [\nabla c_i(x^*)]_{i \in \mathcal{A}(x^*)}, \tag{12.37}$$

where the active set $\mathcal{A}(x^*)$ is defined as in Definition 12.1.

Lemma 12.2.

Let x^* be a feasible point. The following two statements are true.

(i) $T_\Omega(x^*) \subset \mathcal{F}(x^*)$.

(ii) *If the LICQ condition is satisfied at x^*, then $\mathcal{F}(x^*) = T_\Omega(x^*)$.*

PROOF. Without loss of generality, let us assume that all the constraints $c_i(\cdot)$, $i = 1, 2, \ldots, m$, are active at x^*. (We can arrive at this convenient ordering by simply dropping all inactive constraints—which are irrelevant in some neighborhood of x^*—and renumbering the active constraints that remain.)

To prove (i), let $\{z_k\}$ and $\{t_k\}$ be the sequences for which (12.29) is satisfied, that is,

$$\lim_{k \to \infty} \frac{z_k - x^*}{t_k} = d.$$

(Note in particular that $t_k > 0$ for all k.) From this definition, we have that

$$z_k = x^* + t_k d + o(t_k). \tag{12.38}$$

By taking $i \in \mathcal{E}$ and using Taylor's theorem, we have that

$$\begin{aligned} 0 &= \frac{1}{t_k} c_i(z_k) \\ &= \frac{1}{t_k} \left[c_i(x^*) + t_k \nabla c_i(x^*)^T d + o(t_k) \right] \\ &= \nabla c_i(x^*)^T d + \frac{o(t_k)}{t_k}. \end{aligned}$$

By taking the limit as $k \to \infty$, the last term in this expression vanishes, and we have $\nabla c_i(x^*)^T d = 0$, as required. For the active inequality constraints $i \in \mathcal{A}(x^*) \cap \mathcal{I}$, we have similarly that

$$\begin{aligned} 0 &\le \frac{1}{t_k} c_i(z_k) \\ &= \frac{1}{t_k}\left[c_i(x^*) + t_k \nabla c_i(x^*)^T d + o(t_k)\right] \\ &= \nabla c_i(x^*)^T d + \frac{o(t_k)}{t_k}. \end{aligned}$$

Hence, by a similar limiting argument, we have that $\nabla c_i(x^*)^T d \ge 0$, as required.

For (ii), we use the implicit function theorem (see the Appendix or Lang [187, p. 131] for a statement of this result). First, since the LICQ holds, we have from Definition 12.4 that the $m \times n$ matrix $A(x^*)$ of active constraint gradients has full row rank m. Let Z be a matrix whose columns are a basis for the null space of $A(x^*)$; that is,

$$Z \in \mathbb{R}^{n \times (n-m)}, \qquad Z \text{ has full column rank}, \qquad A(x^*)Z = 0. \tag{12.39}$$

(See the related discussion in Chapter 16.) Choose $d \in \mathcal{F}(x^*)$ arbitrarily, and suppose that $\{t_k\}_{k=0}^{\infty}$ is any sequence of positive scalars such $\lim_{k\to\infty} t_k = 0$. Define the parametrized system of equations $R : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$ by

$$R(z, t) = \begin{bmatrix} c(z) - tA(x^*)d \\ Z^T(z - x^* - td) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \tag{12.40}$$

We claim that the solutions $z = z_k$ of this system for small $t = t_k > 0$ give a feasible sequence that approaches x^* and satisfies the definition (12.29).

At $t = 0$, $z = x^*$, and the Jacobian of R at this point is

$$\nabla_z R(x^*, 0) = \begin{bmatrix} A(x^*) \\ Z^T \end{bmatrix}, \tag{12.41}$$

which is nonsingular by construction of Z. Hence, according to the implicit function theorem, the system (12.40) has a unique solution z_k for all values of t_k sufficiently small. Moreover, we have from (12.40) and Definition 12.3 that

$$i \in \mathcal{E} \Rightarrow c_i(z_k) = t_k \nabla c_i(x^*)^T d = 0, \tag{12.42a}$$

$$i \in \mathcal{A}(x^*) \cap \mathcal{I} \Rightarrow c_i(z_k) = t_k \nabla c_i(x^*)^T d \ge 0, \tag{12.42b}$$

so that z_k is indeed feasible.

It remains to verify that (12.29) holds for this choice of $\{z_k\}$. Using the fact that $R(z_k, t_k) = 0$ for all k together with Taylor's theorem, we find that

$$
\begin{aligned}
0 = R(z_k, t_k) &= \begin{bmatrix} c(z_k) - t_k A(x^*)d \\ Z^T(z_k - x^* - t_k d) \end{bmatrix} \\
&= \begin{bmatrix} A(x^*)(z_k - x^*) + o(\|z_k - x^*\|) - t_k A(x^*)d \\ Z^T(z_k - x^* - t_k d) \end{bmatrix} \\
&= \begin{bmatrix} A(x^*) \\ Z^T \end{bmatrix} (z_k - x^* - t_k d) + o(\|z_k - x^*\|).
\end{aligned}
$$

By dividing this expression by t_k and using nonsingularity of the coefficient matrix in the first term, we obtain

$$
\frac{z_k - x^*}{t_k} = d + o\left(\frac{\|z_k - x^*\|}{t_k}\right),
$$

from which it follows that (12.29) is satisfied (for $x = x^*$). Hence, $d \in T_\Omega(x^*)$ for an arbitrary $d \in \mathcal{F}(x^*)$, so the proof of (ii) is complete. □

A FUNDAMENTAL NECESSARY CONDITION

As mentioned above, a local solution of (12.1) is a point x at which all feasible sequences have the property that $f(z_k) \geq f(x)$ for all k sufficiently large. The following result shows that if such a sequence exists, then its limiting directions must make a nonnegative inner product with the objective function gradient.

Theorem 12.3.

If x^* is a local solution of (12.1), then we have

$$
\nabla f(x^*)^T d \geq 0, \quad \text{for all } d \in T_\Omega(x^*). \tag{12.43}
$$

Proof. Suppose for contradiction that there is a tangent d for which $\nabla f(x^*)^T d < 0$. Let $\{z_k\}$ and $\{t_k\}$ be the sequences satisfying Definition 12.2 for this d. We have that

$$
\begin{aligned}
f(z_k) &= f(x^*) + (z_k - x^*)^T \nabla f(x^*) + o(\|z_k - x^*\|) \\
&= f(x^*) + t_k d^T \nabla f(x^*) + o(t_k),
\end{aligned}
$$

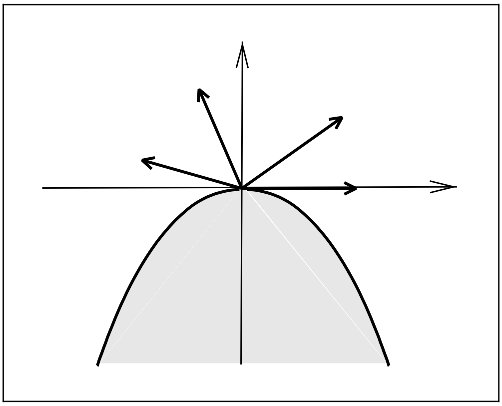

Figure 12.12
Problem (12.44), showing various limiting directions of feasible sequences at the point $(0, 0)^T$.

where the second line follows from (12.38). Since $d^T \nabla f(x^*) < 0$, the remainder term is eventually dominated by the first-order term, that is,

$$f(z_k) < f(x^*) + \tfrac{1}{2} t_k d^T \nabla f(x^*), \qquad \text{for all } k \text{ sufficiently large.}$$

Hence, given any open neighborhood of x^*, we can choose k sufficiently large that z_k lies within this neighborhood and has a lower value of the objective f. Therefore, x^* is not a local solution. □

The converse of this result is not necessarily true. That is, we may have $\nabla f(x^*)^T d \geq 0$ for all $d \in T_\Omega(x^*)$, yet x^* is not a local minimizer. An example is the following problem in two unknowns, illustrated in Figure 12.12

$$\min x_2 \qquad \text{subject to } \; x_2 \geq -x_1^2. \tag{12.44}$$

This problem is actually unbounded, but let us examine its behavior at $x^* = (0, 0)^T$. It is not difficult to show that all limiting directions d of feasible sequences must have $d_2 \geq 0$, so that $\nabla f(x^*)^T d = d_2 \geq 0$. However, x^* is clearly not a local minimizer; the point $(\alpha, -\alpha^2)^T$ for $\alpha > 0$ has a smaller function value than x^*, and can be brought arbitrarily close to x^* by setting α sufficiently small.

FARKAS' LEMMA

The most important step in proving Theorem 12.1 is a classical theorem of the alternative known as *Farkas' Lemma*. This lemma considers a cone K defined as follows:

$$K = \{By + Cw \mid y \geq 0\}, \tag{12.45}$$

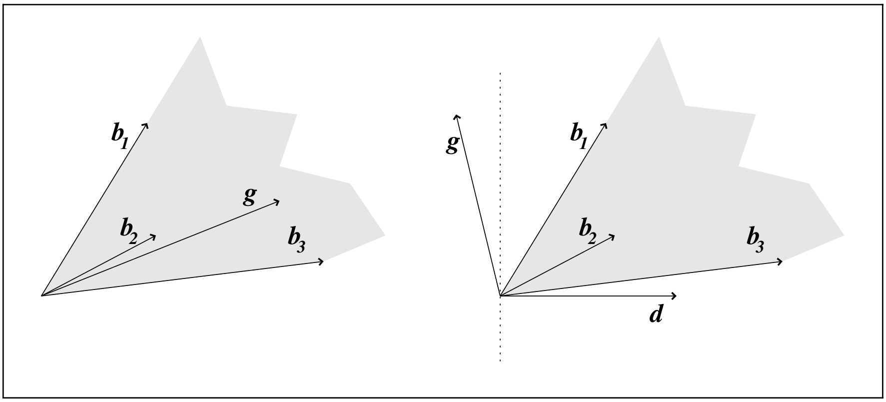

Figure 12.13 Farkas' Lemma: Either $g \in K$ (left) or there is a separating hyperplane (right).

where B and C are matrices of dimension $n \times m$ and $n \times p$, respectively, and y and w are vectors of appropriate dimensions. Given a vector $g \in \mathbb{R}^n$, Farkas' Lemma states that one (and only one) of two alternatives is true. Either $g \in K$, or else there is a vector $d \in \mathbb{R}^n$ such that

$$g^T d < 0, \quad B^T d \geq 0, \quad C^T d = 0. \tag{12.46}$$

The two cases are illustrated in Figure 12.13 for the case of B with three columns, C null, and $n = 2$. Note that in the second case, the vector d defines a *separating hyperplane*, which is a plane in $\mathbb{R}^n$ that separates the vector g from the cone K.

Lemma 12.4 (Farkas).

Let the cone K be defined as in (12.45). Given any vector $g \in \mathbb{R}^n$, we have either that $g \in K$ or that there exists $d \in \mathbb{R}^n$ satisfying (12.46), but not both.

Proof. We show first that the two alternatives cannot hold simultaneously. If $g \in K$, there exist vectors $y \geq 0$ and w such that $g = By + Cw$. If there also exists a d with the property (12.46), we have by taking inner products that

$$0 > d^T g = d^T By + d^T Cw = (B^T d)^T y + (C^T d)^T w \geq 0,$$

where the final inequality follows from $C^T d = 0$, $B^T d \geq 0$, and $y \geq 0$. Hence, we cannot have both alternatives holding at once.

We now show that *one* of the alternatives holds. To be precise, we show how to construct d with the properties (12.46) in the case that $g \notin K$. For this part of the proof, we need to use the property that K is a *closed* set—a fact that is intuitively obvious but not trivial to prove (see Lemma 12.15 in the Notes and References below). Let $\hat{s}$ be the vector in K that is closest to g in the sense of the Euclidean norm. Because K is closed, $\hat{s}$ is well defined and is given by the solution of the following optimization problem:

$$\min \|s - g\|_2^2 \quad \text{subject to } s \in K. \tag{12.47}$$

Since $\hat{s} \in K$, we have from the fact that K is a cone that $\alpha\hat{s} \in K$ for all scalars $\alpha \geq 0$. Since $\|\alpha\hat{s} - g\|_2^2$ is minimized by $\alpha = 1$, we have by simple calculus that

$$\begin{aligned} \frac{d}{d\alpha}\|\alpha\hat{s} - g\|_2^2\Big|_{\alpha=1} = 0 &\Rightarrow \left(-2\hat{s}^T g + 2t\hat{s}^T\hat{s}\right)\Big|_{\alpha=1} = 0 \\ &\Rightarrow \hat{s}^T(\hat{s} - g) = 0. \end{aligned} \tag{12.48}$$

Now, let s be any other vector in K. Since K is convex, we have by the minimizing property of $\hat{s}$ that

$$\|\hat{s} + \theta(s - \hat{s}) - g\|_2^2 \geq \|\hat{s} - g\|_2^2 \quad \text{for all } \theta \in [0, 1],$$

and hence

$$2\theta(s - \hat{s})^T(\hat{s} - g) + \theta^2\|s - \hat{s}\|_2^2 \geq 0.$$

By dividing this expression by θ and taking the limit as $\theta \downarrow 0$, we have $(s - \hat{s})^T(\hat{s} - g) \geq 0$. Therefore, because of (12.48),

$$s^T(\hat{s} - g) \geq 0, \qquad \text{for all } s \in K. \tag{12.49}$$

We claim now that the vector

$$d = \hat{s} - g$$

satisfies the conditions (12.46). Note that $d \neq 0$ because $g \notin K$. We have from (12.48) that

$$d^T g = d^T(\hat{s} - d) = (\hat{s} - g)^T\hat{s} - d^T d = -\|d\|_2^2 < 0,$$

so that d satisfies the first property in (12.46).

From (12.49), we have that $d^T s \geq 0$ for all $s \in K$, so that

$$d^T(By + Cw) \geq 0 \;\; \text{for all } y \geq 0 \text{ and all } w.$$

By fixing $y = 0$ we have that $(C^T d)^T w \geq 0$ for all w, which is true only if $C^T d = 0$. By fixing $w = 0$, we have that $(B^T d)^T y \geq 0$ for all $y \geq 0$, which is true only if $B^T d \geq 0$. Hence, d also satisfies the second and third properties in (12.46) and our proof is complete. □

By applying Lemma 12.4 to the cone N defined by

$$N = \left\{ \sum_{i \in \mathcal{A}(x^*)} \lambda_i \nabla c_i(x^*), \quad \lambda_i \geq 0 \text{ for } i \in \mathcal{A}(x^*) \cap \mathcal{I} \right\}, \tag{12.50}$$

and setting $g = \nabla f(x^*)$, we have that either

$$\nabla f(x^*) = \sum_{i \in \mathcal{A}(x^*)} \lambda_i \nabla c_i(x^*) = A(x^*)^T \lambda^*, \quad \lambda_i \geq 0 \text{ for } i \in \mathcal{A}(x^*) \cap \mathcal{I}, \tag{12.51}$$

or else there is a direction d such that $d^T \nabla f(x^*) < 0$ and $d \in \mathcal{F}(x^*)$.

PROOF OF THEOREM 12.1

Lemmas 12.2 and 12.4 can be combined to give the KKT conditions described in Theorem 12.1. We work through the final steps of the proof here. Suppose that $x^* \in \mathbb{R}^n$ is a feasible point at which the LICQ holds. The theorem claims that if x^* is a local solution for (12.1), then there is a vector $\lambda^* \in \mathbb{R}^m$ that satisfies the conditions (12.34).

We show first that there are multipliers λ_i, $i \in \mathcal{A}(x^*)$, such that (12.51) is satisfied. Theorem 12.3 tells us that $d^T \nabla f(x^*) \geq 0$ for all tangent vectors $d \in T_\Omega(x^*)$. From Lemma 12.2, since LICQ holds, we have that $T_\Omega(x^*) = \mathcal{F}(x^*)$. By putting these two statements together, we find that $d^T \nabla f(x^*) \geq 0$ for all $d \in \mathcal{F}(x^*)$. Hence, from Lemma 12.4, there is a vector λ for which (12.51) holds, as claimed.

We now define the vector λ^* by

$$\lambda_i^* = \begin{cases} \lambda_i, & i \in \mathcal{A}(x^*), \\ 0, & i \in \mathcal{I} \backslash \mathcal{A}(x^*), \end{cases} \tag{12.52}$$

and show that this choice of λ^*, together with our local solution x^*, satisfies the conditions (12.34). We check these conditions in turn.

- The condition (12.34a) follows immediately from (12.51) and the definitions (12.33) of the Lagrangian function and (12.52) of λ^*.
- Since x^* is feasible, the conditions (12.34b) and (12.34c) are satisfied.
- We have from (12.51) that $\lambda_i^* \geq 0$ for $i \in \mathcal{A}(x^*) \cap \mathcal{I}$, while from (12.52), $\lambda_i^* = 0$ for $i \in \mathcal{I} \backslash \mathcal{A}(x^*)$. Hence, $\lambda_i^* \geq 0$ for $i \in \mathcal{I}$, so that (12.34d) holds.
- We have for $i \in \mathcal{A}(x^*) \cap \mathcal{I}$ that $c_i(x^*) = 0$, while for $i \in \mathcal{I} \backslash \mathcal{A}(x^*)$, we have $\lambda_i^* = 0$. Hence $\lambda_i^* c_i(x^*) = 0$ for $i \in \mathcal{I}$, so that (12.34e) is satisfied as well.

The proof is complete.

12.5 SECOND-ORDER CONDITIONS

So far, we have described first-order conditions—the KKT conditions—which tell us how the first derivatives of f and the active constraints c_i are related to each other at a solution x^*. When these conditions are satisfied, a move along any vector w from $\mathcal{F}(x^*)$ either increases the first-order approximation to the objective function (that is, $w^T \nabla f(x^*) > 0$), or else keeps this value the same (that is, $w^T \nabla f(x^*) = 0$).

What role do the *second* derivatives of f and the constraints c_i play in optimality conditions? We see in this section that second derivatives play a "tiebreaking" role. For the directions $w \in \mathcal{F}(x^*)$ for which $w^T \nabla f(x^*) = 0$, we cannot determine from first derivative information alone whether a move along this direction will increase or decrease the objective function f. Second-order conditions examine the second derivative terms in the Taylor series expansions of f and c_i, to see whether this extra information resolves the issue of increase or decrease in f. Essentially, the second-order conditions concern the curvature of the Lagrangian function in the "undecided" directions—the directions $w \in \mathcal{F}(x^*)$ for which $w^T \nabla f(x^*) = 0$.

Since we are discussing second derivatives, stronger smoothness assumptions are needed here than in the previous sections. For the purpose of this section, f and c_i, $i \in \mathcal{E} \cup \mathcal{I}$, are all assumed to be twice continuously differentiable.

Given $\mathcal{F}(x^*)$ from Definition 12.3 and some Lagrange multiplier vector λ^* satisfying the KKT conditions (12.34), we define the *critical cone* $\mathcal{C}(x^*, \lambda^*)$ as follows:

$$\mathcal{C}(x^*, \lambda^*) = \{w \in \mathcal{F}(x^*) \mid \nabla c_i(x^*)^T w = 0, \text{ all } i \in \mathcal{A}(x^*) \cap \mathcal{I} \text{ with } \lambda_i^* > 0\}.$$

Equivalently,

$$w \in \mathcal{C}(x^*, \lambda^*) \;\Leftrightarrow\; \begin{cases} \nabla c_i(x^*)^T w = 0, & \text{for all } i \in \mathcal{E}, \\ \nabla c_i(x^*)^T w = 0, & \text{for all } i \in \mathcal{A}(x^*) \cap \mathcal{I} \text{ with } \lambda_i^* > 0, \\ \nabla c_i(x^*)^T w \geq 0, & \text{for all } i \in \mathcal{A}(x^*) \cap \mathcal{I} \text{ with } \lambda_i^* = 0. \end{cases} \qquad (12.53)$$

The critical cone contains those directions w that would tend to "adhere" to the active inequality constraints even when we were to make small changes to the objective (those indices $i \in \mathcal{I}$ for which the Lagrange multiplier component λ_i^* is positive), as well as to the equality constraints. From the definition (12.53) and the fact that $\lambda_i^* = 0$ for all inactive components $i \in \mathcal{I} \backslash \mathcal{A}(x^*)$, it follows immediately that

$$w \in \mathcal{C}(x^*, \lambda^*) \;\Rightarrow\; \lambda_i^* \nabla c_i(x^*)^T w = 0 \;\text{ for all } i \in \mathcal{E} \cup \mathcal{I}. \qquad (12.54)$$

Hence, from the first KKT condition (12.34a) and the definition (12.33) of the Lagrangian function, we have that

$$w \in \mathcal{C}(x^*, \lambda^*) \;\Rightarrow\; w^T \nabla f(x^*) = \sum_{i \in \mathcal{E} \cup \mathcal{I}} \lambda_i^* w^T \nabla c_i(x^*) = 0. \qquad (12.55)$$

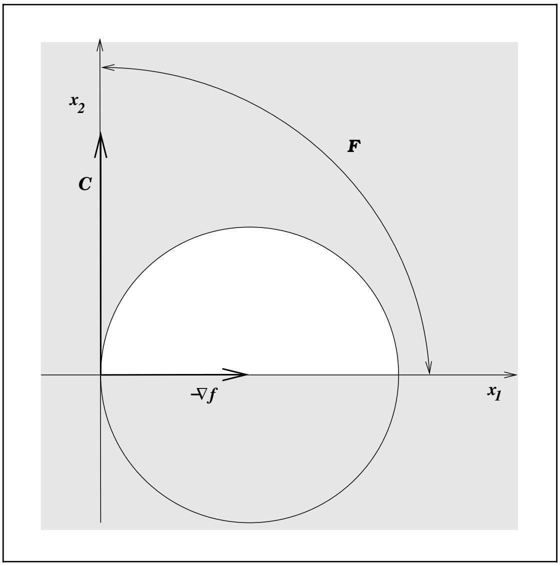

Figure 12.14
Problem (12.56), showing $\mathcal{F}(x^*)$ and $\mathcal{C}(x^*, \lambda^*)$.

Hence the critical cone $\mathcal{C}(x^*, \lambda^*)$ contains directions from $\mathcal{F}(x^*)$ for which it is not clear from first derivative information alone whether f will increase or decrease.

❑ EXAMPLE 12.7

Consider the problem

$$\min\ x_1 \qquad \text{subject to}\ \ x_2 \geq 0,\ \ 1 - (x_1 - 1)^2 - x_2^2 \geq 0, \tag{12.56}$$

illustrated in Figure 12.14. It is not difficult to see that the solution is $x^* = (0, 0)^T$, with active set $\mathcal{A}(x^*) = \{1, 2\}$ and a unique optimal Lagrange multiplier $\lambda^* = (0, 0.5)^T$. Since the gradients of the active constraints at x^* are $(0, 1)^T$ and $(2, 0)^T$, respectively, the LICQ holds, so the optimal multiplier is unique. The linearized feasible set is then

$$\mathcal{F}(x^*) = \{d \mid d \geq 0\},$$

while the critical cone is

$$\mathcal{C}(x^*, \lambda^*) = \{(0, w_2)^T \mid w_2 \geq 0\}.$$

❐

The first theorem defines a *necessary* condition involving the second derivatives: If x^* is a local solution, then the Hessian of the Lagrangian has nonnegative curvature along critical directions (that is, the directions in $\mathcal{C}(x^*, \lambda^*)$).

Theorem 12.5 (Second-Order Necessary Conditions).

Suppose that x^* is a local solution of (12.1) and that the LICQ condition is satisfied. Let λ^* be the Lagrange multiplier vector for which the KKT conditions (12.34) are satisfied. Then

$$w^T \nabla^2_{xx} \mathcal{L}(x^*, \lambda^*) w \geq 0, \quad \text{for all } w \in \mathcal{C}(x^*, \lambda^*). \tag{12.57}$$

PROOF. Since x^* is a local solution, all feasible sequences $\{z_k\}$ approaching x^* must have $f(z_k) \geq f(x^*)$ for all k sufficiently large. Our approach in this proof is to construct a feasible sequence whose limiting direction is w and show that the property $f(z_k) \geq f(x^*)$ implies that (12.57) holds.

Since $w \in \mathcal{C}(x^*, \lambda^*) \subset \mathcal{F}(x^*)$, we can use the technique in the proof of Lemma 12.2 to choose a sequence $\{t_k\}$ of positive scalars and to construct a feasible sequence $\{z_k\}$ approaching x^* such that

$$\lim_{k \to \infty} \frac{z_k - x^*}{t_k} = w, \tag{12.58}$$

which we can write also as (12.58) that

$$z_k - x^* = t_k w + o(t_k). \tag{12.59}$$

Because of the construction technique for $\{z_k\}$, we have from formula (12.42) that

$$c_i(z_k) = t_k \nabla c_i(x^*)^T w, \quad \text{for all } i \in \mathcal{A}(x^*) \tag{12.60}$$

From (12.33), (12.60), and (12.54), we have

$$\begin{aligned} \mathcal{L}(z_k, \lambda^*) &= f(z_k) - \sum_{i \in \mathcal{E} \cup \mathcal{I}} \lambda_i^* c_i(z_k) \\ &= f(z_k) - t_k \sum_{i \in \mathcal{A}(x^*)} \lambda_i^* \nabla c_i(x^*)^T w \\ &= f(z_k), \end{aligned} \tag{12.61}$$

On the other hand, we can perform a Taylor series expansion to obtain an estimate of $\mathcal{L}(z_k, \lambda^*)$ near x^*. By using Taylor's theorem expression (2.6) and continuity of the Hessians $\nabla^2 f$ and $\nabla^2 c_i$, $i \in \mathcal{E} \cup \mathcal{I}$, we obtain

$$\begin{aligned} \mathcal{L}(z_k, \lambda^*) = \mathcal{L}(x^*, \lambda^*) &+ (z_k - x^*)^T \nabla_x \mathcal{L}(x^*, \lambda^*) \\ &+ \tfrac{1}{2}(z_k - x^*)^T \nabla^2_{xx} \mathcal{L}(x^*, \lambda^*)(z_k - x^*) + o(\|z_k - x^*\|^2). \end{aligned} \tag{12.62}$$

By the complementarity conditions (12.34e), we have $\mathcal{L}(x^*, \lambda^*) = f(x^*)$. From (12.34a), the second term on the right-hand side is zero. Hence, using (12.59), we can rewrite (12.62) as

$$\mathcal{L}(z_k, \lambda^*) = f(x^*) + \tfrac{1}{2} t_k^2 w^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) + o(t_k^2). \tag{12.63}$$

By substituting into (12.63), we obtain

$$f(z_k) = f(x^*) + \tfrac{1}{2} t_k^2 w^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) w + o(t_k^2). \tag{12.64}$$

If $w^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) w < 0$, then (12.64) would imply that $f(z_k) < f(x^*)$ for all k sufficiently large, contradicting the fact that x^* is a local solution. Hence, the condition (12.57) must hold, as claimed. □

Sufficient conditions are conditions on f and c_i, $i \in \mathcal{E} \cup \mathcal{I}$, that ensure that x^* is a local solution of the problem (12.1). (They take the opposite tack to necessary conditions, which assume that x^* is a local solution and deduce properties of f and c_i, for the active indices i.) The second-order sufficient condition stated in the next theorem looks very much like the necessary condition just discussed, but it differs in that the constraint qualification is not required, and the inequality in (12.57) is replaced by a strict inequality.

Theorem 12.6 (Second-Order Sufficient Conditions).

Suppose that for some feasible point $x^* \in \mathbb{R}^n$ there is a Lagrange multiplier vector λ^* such that the KKT conditions (12.34) are satisfied. Suppose also that

$$w^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) w > 0, \qquad \text{for all } w \in \mathcal{C}(x^*, \lambda^*),\ w \neq 0. \tag{12.65}$$

Then x^* is a strict local solution for (12.1).

PROOF. First, note that the set $\bar{\mathcal{C}} = \{d \in \mathcal{C}(x^*, \lambda^*) \mid \|d\| = 1\}$ is a compact subset of $\mathcal{C}(x^*, \lambda^*)$, so by (12.65), the minimizer of $d^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) d$ over this set is a strictly positive number, say σ. Since $\mathcal{C}(x^*, \lambda^*)$ is a cone, we have that $(w/\|w\|) \in \bar{\mathcal{C}}$ if and only if $w \in \mathcal{C}(x^*, \lambda^*)$, $w \neq 0$. Therefore, condition (12.65) by

$$w^T \nabla_{xx}^2 \mathcal{L}(x^*, \lambda^*) w \geq \sigma \|w\|^2, \qquad \text{for all } w \in \mathcal{C}(x^*, \lambda^*), \tag{12.66}$$

for $\sigma > 0$ defined as above. (Note that this inequality holds trivially for $w = 0$.)

We prove the result by showing that every feasible sequence $\{z_k\}$ approaching x^* has $f(z_k) \geq f(x^*) + (\sigma/4)\|z_k - x^*\|^2$, for all k sufficiently large. Suppose for contradiction that this is *not* the case, and that there is a sequence $\{z_k\}$ approaching x^* with

$$f(z_k) < f(x^*) + (\sigma/4)\|z_k - x^*\|^2, \qquad \text{for all } k \text{ sufficiently large.} \tag{12.67}$$

By taking a subsequence if necessary, we can identify a limiting direction d such that

$$\lim_{k\to\infty} \frac{z_k - x^*}{\|z_k - x^*\|} = d. \tag{12.68}$$

We have from Lemma 12.2(i) and Definition 12.3 that $d \in \mathcal{F}(x^*)$. From (12.33) and the facts that $\lambda_i^* \geq 0$ and $c_i(z_k) \geq 0$ for $i \in \mathcal{I}$ and $c_i(z_k) = 0$ for $i \in \mathcal{E}$, we have that

$$\mathcal{L}(z_k, \lambda^*) = f(z_k) - \sum_{i\in\mathcal{A}(x^*)} \lambda_i^* c_i(z_k) \leq f(z_k), \tag{12.69}$$

while the Taylor series approximation (12.63) from the proof of Theorem 12.5 continues to hold.

If d were *not* in $\mathcal{C}(x^*, \lambda^*)$, we could identify some index $j \in \mathcal{A}(x^*) \cap \mathcal{I}$ such that the strict positivity condition

$$\lambda_j^* \nabla c_j(x^*)^T d > 0 \tag{12.70}$$

is satisfied, while for the remaining indices $i \in \mathcal{A}(x^*)$, we have

$$\lambda_i^* \nabla c_i(x^*)^T d \geq 0.$$

From Taylor's theorem and (12.68), we have for this particular value of j that

$$\begin{aligned}
\lambda_j^* c_j(z_k) &= \lambda_j^* c_j(x^*) + \lambda_j^* \nabla c_j(x^*)^T (z_k - x^*) + o(\|z_k - x^*\|) \\
&= \|z_k - x^*\| \lambda_j^* \nabla c_j(x^*)^T d + o(\|z_k - x^*\|).
\end{aligned}$$

Hence, from (12.69), we have that

$$\begin{aligned}
\mathcal{L}(z_k, \lambda^*) &= f(z_k) - \sum_{i\in\mathcal{A}(x^*)} \lambda_i^* c_i(z_k) \\
&\leq f(z_k) - \lambda_j^* c_j(z_k) \\
&\leq f(z_k) - \|z_k - x^*\| \lambda_j^* \nabla c_j(x^*)^T d + o(\|z_k - x^*\|).
\end{aligned} \tag{12.71}$$

From the Taylor series estimate (12.63), we have meanwhile that

$$\mathcal{L}(z_k, \lambda^*) = f(x^*) + O(\|z_k - x^*\|^2),$$

and by combining with (12.71), we obtain

$$f(z_k) \geq f(x^*) + \|z_k - x^*\| \lambda_j^* \nabla c_j(x^*)^T d + o(\|z_k - x^*\|).$$

Because of (12.70), this inequality is incompatible with (12.67). We conclude that $d \in \mathcal{C}(x^*, \lambda^*)$, and hence $d^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*)d \geq \sigma$.

By combining the Taylor series estimate (12.63) with (12.69) and using (12.68), we obtain

$$\begin{aligned} f(z_k) &\geq f(x^*) + \tfrac{1}{2}(z_k - x^*)^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*)(z_k - x^*) + o(\|z_k - x^*\|^2) \\ &= f(x^*) + \tfrac{1}{2} d^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*) d\|z_k - x^*\|^2 + o(\|z_k - x^*\|^2) \\ &\geq f(x^*) + (\sigma/2)\|z_k - x^*\|^2 + o(\|z_k - x^*\|^2). \end{aligned}$$

This inequality yields the contradiction to (12.67). We conclude that every feasible sequence $\{z_k\}$ approaching x^* must satisfy $f(z_k) \geq f(x^*) + (\sigma/4)\|z_k - x^*\|^2$, for all k sufficiently large, so x^* is a strict local solution. □

❑ EXAMPLE 12.8 (EXAMPLE 12.2, ONE MORE TIME)

We now return to Example 12.2 to check the second-order conditions for problem (12.18). In this problem we have $f(x) = x_1 + x_2$, $c_1(x) = 2 - x_1^2 - x_2^2$, $\mathcal{E} = \emptyset$, and $\mathcal{I} = \{1\}$. The Lagrangian is

$$\mathcal{L}(x, \lambda) = (x_1 + x_2) - \lambda_1(2 - x_1^2 - x_2^2),$$

and it is easy to show that the KKT conditions (12.34) are satisfied by $x^* = (-1, -1)^T$, with $\lambda_1^* = \frac{1}{2}$. The Lagrangian Hessian at this point is

$$\nabla^2_{xx}\mathcal{L}(x^*, \lambda^*) = \begin{bmatrix} 2\lambda_1^* & 0 \\ 0 & 2\lambda_1^* \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

This matrix is positive definite, so it certainly satisfies the conditions of Theorem 12.6. We conclude that $x^* = (-1, -1)^T$ is a strict local solution for (12.18). (In fact, it is the global solution of this problem, since, as we note later, this problem is a convex programming problem.)

❐

❑ EXAMPLE 12.9

For a more complex example, consider the problem

$$\min\ -0.1(x_1 - 4)^2 + x_2^2 \qquad \text{s.t.} \qquad x_1^2 + x_2^2 - 1 \geq 0, \tag{12.72}$$

in which we seek to minimize a nonconvex function over the *exterior* of the unit circle. Obviously, the objective function is not bounded below on the feasible region, since we can take the feasible sequence

$$\begin{bmatrix} 10 \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 20 \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 30 \\ 0 \end{bmatrix}, \quad \begin{bmatrix} 40 \\ 0 \end{bmatrix},$$

and note that $f(x)$ approaches $-\infty$ along this sequence. Therefore, no global solution exists, but it may still be possible to identify a strict local solution on the boundary of the constraint. We search for such a solution by using the KKT conditions (12.34) and the second-order conditions of Theorem 12.6.

By defining the Lagrangian for (12.72) in the usual way, it is easy to verify that

$$\nabla_x \mathcal{L}(x, \lambda) = \begin{bmatrix} -0.2(x_1 - 4) - 2\lambda_1 x_1 \\ 2x_2 - 2\lambda_1 x_2 \end{bmatrix}, \tag{12.73a}$$

$$\nabla^2_{xx} \mathcal{L}(x, \lambda) = \begin{bmatrix} -0.2 - 2\lambda_1 & 0 \\ 0 & 2 - 2\lambda_1 \end{bmatrix}. \tag{12.73b}$$

The point $x^* = (1, 0)^T$ satisfies the KKT conditions with $\lambda_1^* = 0.3$ and the active set $\mathcal{A}(x^*) = \{1\}$. To check that the second-order sufficient conditions are satisfied at this point, we note that

$$\nabla c_1(x^*) = \begin{bmatrix} 2 \\ 0 \end{bmatrix},$$

so that the set $\mathcal{C}$ defined in (12.53) is simply

$$\mathcal{C}(x^*, \lambda^*) = \{(0, w_2)^T \mid w_2 \in \mathbb{R}\}.$$

Now, by substituting x^* and λ^* into (12.73b), we have for any $w \in \mathcal{C}(x^*, \lambda^*)$ with $w \neq 0$ that $w_2 \neq 0$ and thus

$$w^T \nabla^2_{xx} \mathcal{L}(x^*, \lambda^*) w = \begin{bmatrix} 0 \\ w_2 \end{bmatrix}^T \begin{bmatrix} -0.4 & 0 \\ 0 & 1.4 \end{bmatrix} \begin{bmatrix} 0 \\ w_2 \end{bmatrix} = 1.4 w_2^2 > 0.$$

Hence, the second-order sufficient conditions are satisfied, and we conclude from Theorem 12.6 that $(1, 0)^T$ is a strict local solution for (12.72).

❐

SECOND-ORDER CONDITIONS AND PROJECTED HESSIANS

The second-order conditions are sometimes stated in a form that is slightly weaker but easier to verify than (12.57) and (12.65). This form uses a two-sided projection of the Lagrangian Hessian $\nabla^2_{xx}\mathcal{L}(x^*, \lambda^*)$ onto subspaces that are related to $\mathcal{C}(x^*, \lambda^*)$.

The simplest case is obtained when the multiplier λ^* that satisfies the KKT conditions (12.34) is unique (as happens, for example, when the LICQ condition holds) *and* strict complementarity holds. In this case, the definition (12.53) of $\mathcal{C}(x^*, \lambda^*)$ reduces to

$$\mathcal{C}(x^*, \lambda^*) = \text{Null}\left[\nabla c_i(x^*)^T\right]_{i\in\mathcal{A}(x^*)} = \text{Null}\, A(x^*),$$

where $A(x^*)$ is defined as in (12.37). In other words, $\mathcal{C}(x^*, \lambda^*)$ is the null space of the matrix whose rows are the active constraint gradients at x^*. As in (12.39), we can define the matrix Z with full column rank whose columns span the space $\mathcal{C}(x^*, \lambda^*)$;that is,

$$\mathcal{C}(x^*, \lambda^*) = \{Zu \mid u \in \mathbb{R}^{|\mathcal{A}(x^*)|}\}.$$

Hence, the condition (12.57) in Theorem 12.5 can be restated as

$$u^T Z^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*) Z u \geq 0 \quad \text{for all } u,$$

or, more succinctly,

$$Z^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*) Z \quad \text{is positive semidefinite.}$$

Similarly, the condition (12.65) in Theorem 12.6 can be restated as

$$Z^T \nabla^2_{xx}\mathcal{L}(x^*, \lambda^*) Z \quad \text{is positive definite.}$$

As we show next, Z can be computed numerically, so that the positive (semi)definiteness conditions can actually be checked by forming these matrices and finding their eigenvalues.

One way to compute the matrix Z is to apply a QR factorization to the matrix of active constraint gradients whose null space we seek. In the simplest case above (in which the multiplier λ^* is unique and strictly complementary holds), we define $A(x^*)$ as in (12.37) and write the QR factorization of its transpose as

$$A(x^*)^T = Q\begin{bmatrix} R \\ 0 \end{bmatrix} = \begin{bmatrix} Q_1 & Q_2 \end{bmatrix}\begin{bmatrix} R \\ 0 \end{bmatrix} = Q_1 R, \tag{12.74}$$

where R is a square upper triangular matrix and Q is $n \times n$ orthogonal. If R is nonsingular, we can set $Z = Q_2$. If R is singular (indicating that the active constraint gradients are linearly dependent), a slight enhancement of this procedure that makes use of column pivoting during the QR procedure can be used to identify Z.

12.6 OTHER CONSTRAINT QUALIFICATIONS

We now reconsider constraint qualifications, the conditions discussed in Sections 12.2 and 12.4 that ensure that the linearized approximation to the feasible set Ω captures the essential shape of Ω in a neighborhood of x^*.

One situation in which the linearized feasible direction set $\mathcal{F}(x^*)$ is obviously an adequate representation of the actual feasible set occurs when all the active constraints are already linear; that is,

$$c_i(x) = a_i^T x + b_i, \tag{12.75}$$

for some $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$. It is not difficult to prove a version of Lemma 12.2 for this situation.

Lemma 12.7.

Suppose that at some $x^* \in \Omega$, all active constraints $c_i(\cdot)$, $i \in \mathcal{A}(x^*)$, are linear functions. Then $\mathcal{F}(x^*) = T_\Omega(x^*)$.

PROOF. We have from Lemma 12.2 (i) that $T_\Omega(x^*) \subset \mathcal{F}(x^*)$. To prove that $\mathcal{F}(x^*) \subset T_\Omega(x^*)$, we choose an arbitrary $w \in \mathcal{F}(x^*)$ and show that $w \in T_\Omega(x^*)$. By Definition 12.3 and the form (12.75) of the constraints, we have

$$\mathcal{F}(x^*) = \left\{ d \;\middle|\; \begin{array}{ll} a_i^T d = 0, & \text{for all } i \in \mathcal{E}, \\ a_i^T d \geq 0, & \text{for all } i \in \mathcal{A}(x) \cap \mathcal{I} \end{array} \right\}.$$

First, note that there is a positive scalar $\bar{t}$ such that the inactive constraint remain inactive at $x^* + tw$, for all $t \in [0, \bar{t}]$, that is,

$$c_i(x^* + tw) > 0, \quad \text{for all } i \in \mathcal{I} \backslash \mathcal{A}(x^*) \text{ and all } t \in [0, \bar{t}].$$

Now define the sequence z_k by

$$z_k = x^* + (\bar{t}/k)w, \quad k = 1, 2, \ldots.$$

Since $a_i^T w \geq 0$ for all $i \in \mathcal{I} \cap \mathcal{A}(x^*)$, we have

$$c_i(z_k) = c_i(z_k) - c_i(x^*) = a_i^T(z_k - x^*) = \frac{\bar{t}}{k} a_i^T w \geq 0, \quad \text{for all } i \in \mathcal{I} \cap \mathcal{A}(x^*),$$

so that z_k is feasible with respect to the active inequality constraints c_i, $i \in \mathcal{I} \cap \mathcal{A}(x^*)$. By the choice of $\bar{t}$, we find that z_k is also feasible with respect to the inactive inequality constraints

$i \in \mathcal{I} \backslash \mathcal{A}(x^*)$, and it is easy to show that $c_i(z_k) = 0$ for the equality constraints $i \in \mathcal{E}$. Hence, z_k is feasible for each $k = 1, 2, \ldots$. In addition, we have that

$$\frac{z_k - x^*}{(\bar{t}/k)} = \frac{(\bar{t}/k)w}{(\bar{t}/k)} = w,$$

so that indeed w is the limiting direction of $\{z_k\}$. Hence, $w \in T_\Omega(x^*)$, and the proof is complete. □

We conclude from this result that the condition that all active constraints be linear is another possible constraint qualification. It is neither weaker nor stronger than the LICQ condition, that is, there are situations in which one condition is satisfied but not the other (see Exercise 12.12).

Another useful generalization of the LICQ is the Mangasarian–Fromovitz constraint qualification (MFCQ).

Definition 12.6 (MFCQ).

We say that the Mangasarian–Fromovitz constraint qualification (MFCQ) *holds if there exists a vector* $w \in \mathbb{R}^n$ *such that*

$$\begin{aligned} \nabla c_i(x^*)^T w &> 0, \quad \text{for all } i \in \mathcal{A}(x^*) \cap \mathcal{I}, \\ \nabla c_i(x^*)^T w &= 0, \quad \text{for all } i \in \mathcal{E}, \end{aligned}$$

and the set of equality constraint gradients $\{\nabla c_i(x^*), i \in \mathcal{E}\}$ *is linearly independent.*

Note the *strict* inequality involving the active inequality constraints.

The MFCQ is a weaker condition than LICQ. If LICQ is satisfied, then the system of equalities defined by

$$\begin{aligned} \nabla c_i(x^*)^T w &= 1, \quad \text{for all } i \in \mathcal{A}(x^*) \cap \mathcal{I}, \\ \nabla c_i(x^*)^T w &= 0, \quad \text{for all } i \in \mathcal{E}, \end{aligned}$$

has a solution w, by full rank of the active constraint gradients. Hence, we can choose the w of Definition 12.6 to be precisely this vector. On the other hand, it is easy to construct examples in which the MFCQ is satisfied but the LICQ is not; see Exercise 12.13.

It is possible to prove a version of the first-order necessary condition result (Theorem 12.1) in which MFCQ replaces LICQ in the assumptions. MFCQ gives rise to the nice property that it is equivalent to boundedness of the set of Lagrange multiplier vectors λ^* for which the KKT conditions (12.34) are satisfied. (In the case of LICQ, this set consists of a unique vector λ^*, and so is trivially bounded.)

Note that constraint qualifications are *sufficient* conditions for the linear approximation to be adequate, not necessary conditions. For instance, consider the set defined by $x_2 \geq -x_1^2$ and $x_2 \leq x_1^2$ and the feasible point $x^* = (0, 0)^T$. None of the constraint qualifications we have discussed are satisfied, but the linear approximation $\mathcal{F}(x^*) = \{(w_1, 0)^T \mid w_1 \in \mathbb{R}\}$ accurately reflects the geometry of the feasible set near x^*.

12.7 A GEOMETRIC VIEWPOINT

Finally, we mention an alternative first-order optimality condition that depends only on the geometry of the feasible set Ω and not on its particular algebraic description in terms of the constraint functions c_i, $i \in \mathcal{E} \cup \mathcal{I}$. In geometric terms, our problem (12.1) can be stated as

$$\min f(x) \quad \text{subject to } x \in \Omega, \tag{12.76}$$

where Ω is the feasible set.

To prove a "geometric" first-order condition, we need to define the normal cone to the set Ω at a feasible point x.

Definition 12.7.

The normal cone *to the set Ω at the point $x \in \Omega$ is defined as*

$$N_\Omega(x) = \{v \mid v^T w \le 0 \text{ for all } w \in T_\Omega(x)\}, \tag{12.77}$$

where $T_\Omega(x)$ is the tangent cone of Definition 12.2. Each vector $v \in N_\Omega(x)$ is said to be a normal vector.

Geometrically, each normal vector v makes an angle of at least $\pi/2$ with every tangent vector.

The first-order necessary condition for (12.76) is delightfully simple.

Theorem 12.8.

Suppose that x^* is a local minimizer of f in Ω. Then

$$-\nabla f(x^*) \in N_\Omega(x^*). \tag{12.78}$$

PROOF. Given any $d \in T_\Omega(x^*)$, we have for the sequences $\{t_k\}$ and $\{z_k\}$ in Definition 12.2 that

$$z_k \in \Omega, \qquad z_k = x^* + t_k d + o(t_k), \qquad \text{for all } k. \tag{12.79}$$

Since x^* is a local solution, we must have

$$f(z_k) \ge f(x^*)$$

for all k sufficiently large. Hence, since f is continuously differentiable, we have from Taylor's theorem (2.4) that

$$f(z_k) - f(x^*) = t_k \nabla f(x^*)^T d + o(t_k) \geq 0.$$

By dividing by t_k and taking limits as $k \to \infty$, we have

$$\nabla f(x^*)^T d \geq 0.$$

Recall that d was an arbitrary member of $T_\Omega(x^*)$, so we have $-\nabla f(x^*)^T d \leq 0$ for all $d \in T_\Omega(x^*)$. We conclude from Definition 12.7 that $-\nabla f(x^*) \in N_\Omega(x^*)$. □

This result suggests a close relationship between $N_\Omega(x^*)$ and the conic combination of active constraint gradients given by (12.50). When the linear independence constraint qualification holds, identical (to within a change of sign).

Lemma 12.9.

Suppose that the LICQ assumption (Definition 12.4) holds at x^*. Then t the normal cone $N_\Omega(x^*)$ is simply $-N$, where N is the set defined in (12.50).

PROOF. The proof follows from Farkas' Lemma (Lemma 12.4) and Definition 12.7 of $N_\Omega(x^*)$. From Lemma 12.4, we have that

$$g \in N \quad \Rightarrow \quad g^T d \geq 0 \;\; \text{for all } d \in \mathcal{F}(x^*).$$

Since we have $\mathcal{F}(x^*) = T_\Omega(x^*)$ from Lemma 12.2, it follows by switching the sign of this expression that

$$g \in -N \quad \Rightarrow \quad g^T d \leq 0 \;\; \text{for all } d \in T_\Omega(x^*).$$

We conclude from Definition 12.7 that $N_\Omega(x^*) = -N$, as claimed. □

12.8 LAGRANGE MULTIPLIERS AND SENSITIVITY

The importance of Lagrange multipliers in optimality theory should be clear, but what of their intuitive significance? We show in this section that each Lagrange multiplier λ_i^* tells us something about the *sensitivity* of the optimal objective value $f(x^*)$ to the presence of the constraint c_i. To put it another way, λ_i^* indicates how hard f is "pushing" or "pulling" the solution x^* against the particular constraint c_i.

We illustrate this point with some informal analysis. When we choose an inactive constraint $i \notin \mathcal{A}(x^*)$ such that $c_i(x^*) > 0$, the solution x^* and function value $f(x^*)$ are indifferent to whether this constraint is present or not. If we perturb c_i by a tiny amount, it will still be inactive and x^* will still be a local solution of the optimization problem. Since $\lambda_i^* = 0$ from (12.34e), the Lagrange multiplier indicates accurately that constraint i is not significant.

Suppose instead that constraint i is active, and let us perturb the right-hand-side of this constraint a little, requiring, say, that $c_i(x) \geq -\epsilon\|\nabla c_i(x^*)\|$ instead of $c_i(x) \geq 0$. Suppose that ϵ is sufficiently small that the perturbed solution $x^*(\epsilon)$ still has the same set of active constraints, and that the Lagrange multipliers are not much affected by the perturbation. (These conditions can be made more rigorous with the help of strict complementarity and second-order conditions.) We then find that

$$\begin{aligned} -\epsilon\|\nabla c_i(x^*)\| &= c_i(x^*(\epsilon)) - c_i(x^*) \approx (x^*(\epsilon) - x^*)^T \nabla c_i(x^*), \\ 0 &= c_j(x^*(\epsilon)) - c_j(x^*) \approx (x^*(\epsilon) - x^*)^T \nabla c_j(x^*), \\ &\qquad \text{for all } j \in \mathcal{A}(x^*) \text{ with } j \neq i. \end{aligned}$$

The value of $f(x^*(\epsilon))$, meanwhile, can be estimated with the help of (12.34a). We have

$$\begin{aligned} f(x^*(\epsilon)) - f(x^*) &\approx (x^*(\epsilon) - x^*)^T \nabla f(x^*) \\ &= \sum_{j \in \mathcal{A}(x^*)} \lambda_j^*(x^*(\epsilon) - x^*)^T \nabla c_j(x^*) \\ &\approx -\epsilon\|\nabla c_i(x^*)\|\lambda_i^*. \end{aligned}$$

By taking limits, we see that the family of solutions $x^*(\epsilon)$ satisfies

$$\frac{df(x^*(\epsilon))}{d\epsilon} = -\lambda_i^*\|\nabla c_i(x^*)\|. \tag{12.80}$$

A sensitivity analysis of this problem would conclude that if $\lambda_i^*\|\nabla c_i(x^*)\|$ is large, then the optimal value is sensitive to the placement of the ith constraint, while if this quantity is small, the dependence is not too strong. If λ_i^* is exactly zero for some active constraint, small perturbations to c_i in some directions will hardly affect the optimal objective value at all; the change is zero, to first order.

This discussion motivates the definition below, which classifies constraints according to whether or not their corresponding Lagrange multiplier is zero.

Definition 12.8.

Let x^* be a solution of the problem (12.1), and suppose that the KKT conditions (12.34) are satisfied. We say that an inequality constraint c_i is strongly active *or* binding *if $i \in \mathcal{A}(x^*)$ and $\lambda_i^* > 0$ for some Lagrange multiplier λ^* satisfying (12.34). We say that c_i is* weakly active *if $i \in \mathcal{A}(x^*)$ and $\lambda_i^* = 0$ for all λ^* satisfying (12.34).*

Note that the analysis above is independent of scaling of the individual constraints. For instance, we might change the formulation of the problem by replacing some active constraint c_i by $10c_i$. The new problem will actually be equivalent (that is, it has the same feasible set and same solution), but the optimal multiplier λ_i^* corresponding to c_i will be replaced by $\lambda_i^*/10$. However, since $\|\nabla c_i(x^*)\|$ is replaced by $10\|\nabla c_i(x^*)\|$, the product $\lambda_i^*\|\nabla c_i(x^*)\|$ does not change. If, on the other hand, we replace the objective function f by $10f$, the multipliers λ_i^* in (12.34) all will need to be replaced by $10\lambda_i^*$. Hence in (12.80) we see that the sensitivity of f to perturbations has increased by a factor of 10, which is exactly what we would expect.

12.9 DUALITY

In this section we present some elements of the duality theory for nonlinear programming. This theory is used to motivate and develop some important algorithms, including the augmented Lagrangian algorithms of Chapter 17. In its full generality, duality theory ranges beyond nonlinear programming to provide important insight into the fields of convex nonsmooth optimization and even discrete optimization. Its specialization to linear programming proved central to the development of that area; see Chapter 13. (We note that the discussion of linear programming duality in Section 13.1 can be read without consulting this section first.)

Duality theory shows how we can construct an alternative problem from the functions and data that define the original optimization problem. This alternative "dual" problem is related to the original problem (which is sometimes referred to in this context as the "primal" for purposes of contrast) in fascinating ways. In some cases, the dual problem is easier to solve computationally than the original problem. In other cases, the dual can be used to obtain easily a lower bound on the optimal value of the objective for the primal problem. As remarked above, the dual has also been used to design algorithms for solving the primal problem.

Our results in this section are mostly restricted to the special case of (12.1) in which there are no equality constraints and the objective f and the negatives of the inequality constraints $-c_i$ are all convex functions. For simplicity we assume that there are m inequality constraints labelled $1, 2, \ldots, m$ and rewrite (12.1) as follows:

$$\min_{x \in \mathbb{R}^n} f(x) \qquad \text{subject to } \; c_i(x) \ge 0, \quad i = 1, 2, \ldots, m.$$

If we assemble the constraints into a vector function

$$c(x) \stackrel{\text{def}}{=} (c_1(x), c_2(x), \ldots, c_m(x))^T,$$

we can write the problem as

$$\min_{x \in \mathbb{R}^n} f(x) \qquad \text{subject to } \; c(x) \ge 0, \tag{12.81}$$

for which the Lagrangian function (12.16) with Lagrange multiplier vector $\lambda \in \mathbb{R}^m$ is

$$\mathcal{L}(x, \lambda) = f(x) - \lambda^T c(x).$$

We define the dual objective function $q : \mathbb{R}^n \to \mathbb{R}$ as follows:

$$q(\lambda) \stackrel{\text{def}}{=} \inf_x \mathcal{L}(x, \lambda). \tag{12.82}$$

In many problems, this infimum is $-\infty$ for some values of λ. We define the domain of q as the set of λ values for which q is finite, that is,

$$\mathcal{D} \stackrel{\text{def}}{=} \{\lambda \mid q(\lambda) > -\infty\}. \tag{12.83}$$

Note that calculation of the infimum in (12.82) requires finding the *global* minimizer of the function $\mathcal{L}(\cdot, \lambda)$ for the given λ which, as we have noted in Chapter 2, may be extremely difficult in practice. However, when f and $-c_i$ are convex functions and $\lambda \geq 0$ (the case in which we are most interested), the function $\mathcal{L}(\cdot, \lambda)$ is also convex. In this situation, all local minimizers are global minimizers (as we verify in Exercise 12.4), so computation of $q(\lambda)$ becomes a more practical proposition.

The dual problem to (12.81) is defined as follows:

$$\max_{\lambda \in \mathbb{R}^n} q(\lambda) \qquad \text{subject to} \;\; \lambda \geq 0. \tag{12.84}$$

❑ EXAMPLE 12.10

Consider the problem

$$\min_{(x_1, x_2)} \; 0.5(x_1^2 + x_2^2) \quad \text{subject to } x_1 - 1 \geq 0. \tag{12.85}$$

The Lagrangian is

$$\mathcal{L}(x_1, x_2, \lambda_1) = 0.5(x_1^2 + x_2^2) - \lambda_1(x_1 - 1).$$

If we hold λ_1 fixed, this is a convex function of $(x_1, x_2)^T$. Therefore, the infimum with respect to $(x_1, x_2)^T$ is achieved when the partial derivatives with respect to x_1 and x_2 are zero, that is,

$$x_1 - \lambda_1 = 0, \quad x_2 = 0.$$

By substituting these infimal values into $\mathcal{L}(x_1, x_2, \lambda_1)$ we obtain the dual objective (12.82):

$$q(\lambda_1) = 0.5(\lambda_1^2 + 0) - \lambda_1(\lambda_1 - 1) = -0.5\lambda_1^2 + \lambda_1.$$

Hence, the dual problem (12.84) is

$$\max_{\lambda_1 \geq 0} -0.5\lambda_1^2 + \lambda_1, \tag{12.86}$$

which clearly has the solution $\lambda_1 = 1$. ❐

In the remainder of this section, we show how the dual problem is related to (12.81). Our first result concerns concavity of q.

Theorem 12.10.

The function q defined by (12.82) is concave and its domain $\mathcal{D}$ is convex.

PROOF. For any λ^0 and λ^1 in $\mathbb{R}^m$, any $x \in \mathbb{R}^n$, and any $\alpha \in [0, 1]$, we have

$$\mathcal{L}(x, (1-\alpha)\lambda^0 + \alpha\lambda^1) = (1-\alpha)\mathcal{L}(x, \lambda^0) + \alpha\mathcal{L}(x, \lambda^1).$$

By taking the infimum of both sides in this expression, using the definition (12.82), and using the results that the infimum of a sum is greater than or equal to the sum of infimums, we obtain

$$q((1-\alpha)\lambda^0 + \alpha\lambda^1) \geq (1-\alpha)q(\lambda^0) + \alpha q(\lambda^1),$$

confirming concavity of q. If both λ^0 and λ^1 belong to $\mathcal{D}$, this inequality implies that $q((1-\alpha)\lambda^0 + \alpha\lambda^1) \geq -\infty$ also, and therefore $(1-\alpha)\lambda^0 + \alpha\lambda^1 \in \mathcal{D}$, verifying convexity of $\mathcal{D}$. □

The optimal value of the dual problem (12.84) gives a lower bound on the optimal objective value for the primal problem (12.81). This observation is a consequence of the following *weak duality* result.

Theorem 12.11 (Weak Duality).

For any $\bar{x}$ feasible for (12.81) and any $\bar{\lambda} \geq 0$, we have $q(\bar{\lambda}) \leq f(\bar{x})$.

PROOF.

$$q(\bar{\lambda}) = \inf_x f(x) - \bar{\lambda}^T c(x) \leq f(\bar{x}) - \bar{\lambda}^T c(\bar{x}) \leq f(\bar{x}),$$

where the final inequality follows from $\bar{\lambda} \geq 0$ and $c(\bar{x}) \geq 0$. □

For the remaining results, we note that the KKT conditions (12.34) specialized to (12.81) are as follows:

$$\nabla f(\bar{x}) - \nabla c(\bar{x})\bar{\lambda} = 0, \tag{12.87a}$$

$$c(\bar{x}) \geq 0, \tag{12.87b}$$

$$\bar{\lambda} \geq 0, \tag{12.87c}$$

$$\bar{\lambda}_i c_i(\bar{x}) = 0, \quad i = 1, 2, \ldots, m, \tag{12.87d}$$

where $\nabla c(x)$ is the $n \times m$ matrix defined by $\nabla c(x) = [\nabla c_1(x), \nabla c_2(x), \ldots, \nabla c_m(x)]$.

The next result shows that optimal Lagrange multipliers for (12.81) are solutions of the dual problem (12.84) under certain conditions. It is essentially due to Wolfe [309].

Theorem 12.12.

Suppose that $\bar{x}$ is a solution of (12.81) and that f and $-c_i$, $i = 1, 2, \ldots, m$ are convex functions on $\mathbb{R}^n$ that are differentiable at $\bar{x}$. Then any $\bar{\lambda}$ for which $(\bar{x}, \bar{\lambda})$ satisfies the KKT conditions (12.87) is a solution of (12.84).

PROOF. Suppose that $(\bar{x}, \bar{\lambda})$ satisfies (12.87). We have from $\bar{\lambda} \geq 0$ that $\mathcal{L}(\cdot, \bar{\lambda})$ is a convex and differentiable function. Hence, for any x, we have

$$\mathcal{L}(x, \bar{\lambda}) \geq \mathcal{L}(\bar{x}, \bar{\lambda}) + \nabla_x \mathcal{L}(\bar{x}, \bar{\lambda})^T (x - \bar{x}) = \mathcal{L}(\bar{x}, \bar{\lambda}),$$

where the last equality follows from (12.87a). Therefore, we have

$$q(\bar{\lambda}) = \inf_x \mathcal{L}(x, \bar{\lambda}) = \mathcal{L}(\bar{x}, \bar{\lambda}) = f(\bar{x}) - \bar{\lambda}^T c(\bar{x}) = f(\bar{x}),$$

where the last equality follows from (12.87d). Since from Theorem 12.11, we have $q(\lambda) \leq f(\bar{x})$ for all $\lambda \geq 0$ it follows immediately from $q(\bar{\lambda}) = f(\bar{x})$ that $\bar{\lambda}$ is a solution of (12.84). $\square$

Note that if the functions are continuously differentiable and a constraint qualification such as LICQ holds at $\bar{x}$, then an optimal Lagrange multiplier is guaranteed to exist, by Theorem 12.1.

In Example 12.10, we see that $\lambda_1 = 1$ is both an optimal Lagrange multiplier for the problem (12.85) and a solution of (12.86). Note too that the optimal objective for both problems is 0.5.

We prove a partial converse of Theorem 12.12, which shows that solutions to the dual problem (12.84) can sometimes be used to derive solutions to the original problem (12.81). The essential condition is strict convexity of the function $\mathcal{L}(\cdot, \hat{\lambda})$ for a certain value $\hat{\lambda}$. We note that this condition holds if either f is strictly convex (as is the case in Example 12.10) or if c_i is strictly convex for some $i = 1, 2, \ldots, m$ with $\hat{\lambda}_i > 0$.

Theorem 12.13.

Suppose that f and $-c_i$, $i = 1, 2, \dots, m$ are convex and continuously differentiable on $\mathbb{R}^n$. Suppose that $\bar{x}$ is a solution of (12.81) at which LICQ holds. Suppose that $\hat{\lambda}$ solves (12.84) and that the infimum in $\inf_x \mathcal{L}(x, \hat{\lambda})$ is attained at $\hat{x}$. Assume further than $\mathcal{L}(\cdot, \hat{\lambda})$ is a strictly *convex function. Then $\bar{x} = \hat{x}$ (that is, $\hat{x}$ is the unique solution of (12.81)), and $f(\bar{x}) = \mathcal{L}(\hat{x}, \hat{\lambda})$.*

PROOF. Assume for contradiction that $\bar{x} \neq \hat{x}$. From Theorem 12.1, because of the LICQ assumption, there exists $\bar{\lambda}$ satisfying (12.87). Hence, from Theorem 12.12, we have that $\bar{\lambda}$ also solves (12.84), so that

$$\mathcal{L}(\bar{x}, \bar{\lambda}) = q(\bar{\lambda}) = q(\hat{\lambda}) = \mathcal{L}(\hat{x}, \hat{\lambda}).$$

Because $\hat{x} = \arg\min_x \mathcal{L}(x, \hat{\lambda})$, we have from Theorem 2.2 that $\nabla_x \mathcal{L}(\hat{x}, \hat{\lambda}) = 0$. Moreover, by strict convexity of $\mathcal{L}(\cdot, \hat{\lambda})$, it follows that

$$\mathcal{L}(\bar{x}, \hat{\lambda}) - \mathcal{L}(\hat{x}, \hat{\lambda}) > \nabla_x \mathcal{L}(\hat{x}, \hat{\lambda})^T (\bar{x} - \hat{x}) = 0.$$

Hence, we have

$$\mathcal{L}(\bar{x}, \hat{\lambda}) > \mathcal{L}(\hat{x}, \hat{\lambda}) = \mathcal{L}(\bar{x}, \bar{\lambda}),$$

so in particular we have

$$-\hat{\lambda}^T c(\bar{x}) > -\bar{\lambda}^T c(\bar{x}) = 0,$$

where the final equality follows from (12.87d). Since $\hat{\lambda} \geq 0$ and $c(\bar{x}) \geq 0$, this yields the contradiction, and we conclude that $\hat{x} = \bar{x}$, as claimed. □

In Example 12.10, at the dual solution $\lambda_1 = 1$, the infimum of $\mathcal{L}(x_1, x_2, \lambda_1)$ is achieved at $(x_1, x_2) = (1, 0)^T$, which is the solution of the original problem (12.85).

An slightly different form of duality that is convenient for computations, known as the *Wolfe dual* [309], can be stated as follows:

$$\max_{x,\lambda} \mathcal{L}(x, \lambda) \tag{12.88a}$$

$$\text{subject to } \nabla_x \mathcal{L}(x, \lambda) = 0, \quad \lambda \geq 0. \tag{12.88b}$$

The following results explains the relationship of the Wolfe dual to (12.81).

Theorem 12.14.

Suppose that f and $-c_i$, $i = 1, 2, \dots, m$ are convex and continuously differentiable on $\mathbb{R}^n$. Suppose that $(\bar{x}, \bar{\lambda})$ is a solution pair of (12.81) at which LICQ holds. Then $(\bar{x}, \bar{\lambda})$ solves the problem (12.88).

PROOF. From the KKT conditions (12.87) we have that $(\bar{x}, \bar{\lambda})$ satisfies (12.88b), and that $\mathcal{L}(\bar{x}, \bar{\lambda}) = f(\bar{x})$. Therefore for any pair (x, λ) that satisfies (12.88b) we have that

$$\begin{aligned} \mathcal{L}(\bar{x}, \bar{\lambda}) &= f(\bar{x}) \\ &\geq f(\bar{x}) - \lambda^T c(\bar{x}) \\ &= \mathcal{L}(\bar{x}, \lambda) \\ &\geq \mathcal{L}(x, \lambda) + \nabla_x \mathcal{L}(x, \lambda)^T (\bar{x} - x) \\ &= \mathcal{L}(x, \lambda), \end{aligned}$$

where the second inequality follows from the convexity of $\mathcal{L}(\cdot, \lambda)$. We have therefore shown that $(\bar{x}, \bar{\lambda})$ maximizes $\mathcal{L}$ over the constraints (12.88b), and hence solves (12.88). □

❑ **EXAMPLE 12.11** (LINEAR PROGRAMMING)

An important special case of (12.81) is the linear programming problem

$$\min c^T x \text{ subject to } Ax - b \geq 0, \tag{12.89}$$

for which the dual objective is

$$q(\lambda) = \inf_x \left[c^T x - \lambda^T (Ax - b) \right] = \inf_x \left[(c - A^T \lambda)^T x + b^T \lambda \right].$$

If $c - A^T \lambda \neq 0$, the infimum is clearly $-\infty$ (we can set x to be a large negative multiple of $-(c - A^T \lambda)$ to make q arbitrarily large and negative). When $c - A^T \lambda = 0$, on the other hand, the dual objective is simply $b^T \lambda$. In maximizing q, we can exclude λ for which $c - A^T \lambda \neq 0$ from consideration (the maximum obviously cannot be attained at a point λ for which $q(\lambda) = -\infty$). Hence, we can write the dual problem (12.84) as follows:

$$\max_\lambda b^T \lambda \quad \text{subject to } A^T \lambda = c, \; \lambda \geq 0. \tag{12.90}$$

The Wolfe dual of (12.89) can be written as

$$\max_\lambda c^T x - \lambda^T (Ax - b) \quad \text{subject to } A^T \lambda = c, \; \lambda \geq 0,$$

and by substituting the constraint $A^T \lambda - c = 0$ into the objective we obtain (12.90) again.

For some matrices A, the dual problem (12.90) may be computationally easier to solve than the original problem (12.89). We discuss the possibilities further in Chapter 13. ❐

❑ **EXAMPLE 12.12** (CONVEX QUADRATIC PROGRAMMING)

Consider

$$\min \frac{1}{2} x^T G x + c^T x \quad \text{subject to} \quad Ax - b \geq 0, \tag{12.91}$$

where G is a symmetric positive definite matrix. The dual objective for this problem is

$$q(\lambda) = \inf_x \mathcal{L}(x, \lambda) = \inf_x \frac{1}{2} x^T G x + c^T x - \lambda^T (Ax - b). \tag{12.92}$$

Since G is positive definite, since $\mathcal{L}(\cdot, \lambda)$ is a strictly convex quadratic function, the infimum is achieved when $\nabla_x \mathcal{L}(x, \lambda) = 0$, that is,

$$Gx + c - A^T \lambda = 0. \tag{12.93}$$

Hence, we can substitute for x in the infimum expression and write the dual objective explicitly as follows:

$$q(\lambda) = -\frac{1}{2} (A^T \lambda - c)^T G^{-1} (A^T \lambda - c)^T + b^T \lambda.$$

Alternatively, we can write the Wolfe dual form (12.88) by retaining x as a variable and including the constraint (12.93) explicitly in the dual problem, to obtain

$$\begin{aligned} \max_{(\lambda, x)} \quad & \frac{1}{2} x^T G x + c^T x - \lambda^T (Ax - b) \qquad (12.94) \\ \text{subject to} \quad & Gx + c - A^T \lambda = 0, \quad \lambda \geq 0. \end{aligned}$$

To make it clearer that the objective is concave, we can use the constraint to substitute $(c - A^T \lambda)^T x = -x^T G x$ in the objective, and rewrite the dual formulation as follows:

$$\max_{(\lambda, x)} -\frac{1}{2} x^T G x + \lambda^T b, \quad \text{subject to} \quad Gx + c - A^T \lambda = 0, \ \lambda \geq 0. \tag{12.95}$$

❐

Note that the Wolfe dual form requires only positive *semidefiniteness* of G.

NOTES AND REFERENCES

The theory of constrained optimization is discussed in many books on numerical optimization. The discussion in Fletcher [101, Chapter 9] is similar to ours, though a little terser, and includes additional material on duality. Bertsekas [19, Chapter 3] emphasizes the role of duality and discusses sensitivity of the solution with respect to the active constraints in some detail. The classic treatment of Mangasarian [198] is particularly notable for its thorough description of constraint qualifications. It also has an extensive discussion of theorems of the alternative [198, Chapter 2], placing Farkas' Lemma firmly in the context of other related results.

The KKT conditions were described in a 1951 paper of Kuhn and Tucker [185], though they were derived earlier (and independently) in an unpublished 1939 master's thesis of W. Karush. Lagrange multipliers and optimality conditions for general problems (including nonsmooth problems) are described in the deep and wide-ranging article of Rockafellar [270].

Duality theory for nonlinear programming is described in the books of Rockafellar [198] and Bertsekas [19]; the latter treatment is particularly extensive and general. The material in Section 12.9 is adapted from these sources.

We return to our claim that the set N defined by

$$N = \{By + Ct \mid y \geq 0\},$$

(where B and C are matrices of dimension $n \times m$ and $n \times p$, respectively, and y and t are vectors of appropriate dimensions; see (12.45)) is a closed set. This fact is needed in the proof of Lemma 12.4 to ensure that the solution of the projection subproblem (12.47) is well-defined. The following technical result is well known; the proof given below is due to R. Byrd.

Lemma 12.15.

The set N is closed.

PROOF. By splitting t into positive and negative parts, it is easy to see that

$$N = \left\{ \begin{bmatrix} B & C & -C \end{bmatrix} \begin{bmatrix} y \\ t^+ \\ t^- \end{bmatrix} \middle| \begin{bmatrix} y \\ t^+ \\ t^- \end{bmatrix} \geq 0 \right\}.$$

Hence, we can assume without loss of generality that N has the form

$$N = \{By \mid y \geq 0\}.$$

Suppose that B has dimensions $n \times m$.

First, we show that for any $s \in N$, we can write $s = B_I y_I$ with $y_I \geq 0$, where $I \subset \{1, 2, \ldots, m\}$, B_I is the column submatrix of B indexed by I with full column rank, and I has minimum cardinality. To prove this claim, we assume for contradiction that $K \subset \{1, 2, \ldots, m\}$ is an index set with minimal cardinality such that $s = B_K y_K$, $y_K \geq 0$, yet the columns of B_K are linearly dependent. Since K is minimal, y_K has no zero components. We then have a nonzero vector w such that $B_K w = 0$. Since $s = B_K(y_K + \tau w)$ for any τ, we can increase or decrease τ from 0 until one or more components of $y_K + \tau w$ become zero, while the other components remain positive. We define $\bar{K}$ by removing the indices from K that correspond to zero components of $y_K + \tau w$, and define $\bar{y}_{\bar{K}}$ to be the vector of strictly positive components of $y_K + \tau w$. We then have that $s = B_{\bar{K}} \bar{y}_{\bar{K}}$ and $\bar{y}_{\bar{K}} \geq 0$, contradicting our assumption that K was the set of minimal cardinality with this property.

Now let $\{s^k\}$ be a sequence with $s^k \in N$ for all k and $s^k \to s$. We prove the lemma by showing that $s \in N$. By the claim of the previous paragraph, for all k we can write $s^k = B_{I_k} y^k_{I_k}$ with $y^k_{I_k} \geq 0$, I_k is minimal, and the columns of B_{I_k} are linearly independent. Since there only finitely many possible choices of index set I_k, at least one index set occurs infinitely often in the sequence. By choosing such an index set I, we can take a subsequence if necessary and assume without loss of generality that $I_k \equiv I$ for all k. We then have that $s^k = A_I y^k_I$ with $y^k_I \geq 0$ and A_I has full column rank. Because of the latter property, we have that $A_I^T A_I$ is invertible, so that y^k_I is defined uniquely as follows:

$$y^k_I = (A_I^T A_I)^{-1} A_I^T s^k, \qquad k = 0, 1, 2, \ldots.$$

By taking limits and using $s^k \to s$, we have that

$$y^k_I \to y_I \stackrel{\text{def}}{=} (A_I^T A_I)^{-1} A_I^T s,$$

and moreover $y_I \geq 0$, since $y^k_I \geq 0$ for all k. Hence we can write $s = B_I y_I$ with $y_I \geq 0$, and therefore $s \in N$. □

EXERCISES

12.1 The following example from [268] with a single variable $x \in \mathbb{R}$ and a single equality constraint shows that strict local solutions are not necessarily isolated. Consider

$$\min_x \; x^2 \qquad \text{subject to } c(x) = 0, \text{ where } \quad c(x) = \begin{cases} x^6 \sin(1/x) = 0 & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases} \tag{12.96}$$

(a) Show that the constraint function is twice continuously differentiable at all x (including at $x = 0$) and that the feasible points are $x = 0$ and $x = 1/(k\pi)$ for all nonzero integers k.

(b) Verify that each feasible point except $x = 0$ is an isolated local solution by showing that there is a neighborhood $\mathcal{N}$ around each such point within which it is the only feasible point.

(c) Verify that $x = 0$ is a global solution and a strict local solution, but not an isolated local solution

✐ **12.2** Is an isolated local solution necessarily a strict local solution? Explain.

✐ **12.3** Does problem (12.4) have a finite or infinite number of local solutions? Use the first-order optimality conditions (12.34) to justify your answer.

✐ **12.4** If f is convex and the feasible region Ω is convex, show that local solutions of the problem (12.3) are also global solutions. Show that the set of global solutions is convex. (Hint: See Theorem 2.5.)

✐ **12.5** Let $v : \mathbb{R}^n \to \mathbb{R}^m$ be a smooth vector function and consider the unconstrained optimization problems of minimizing $f(x)$ where

$$f(x) = \|v(x)\|_\infty, \qquad f(x) = \max_{i=1,2,\dots,m} v_i(x).$$

Reformulate these (generally nonsmooth) problems as smooth constrained optimization problems.

✐ **12.6** Can you perform a smooth reformulation as in the previous question when f is defined by

$$f(x) = \min_{i=1,2,\dots,m} f_i(x)?$$

(N.B. "min" not "max.") Why or why not?

✐ **12.7** Show that the vector defined by (12.15) satisfies (12.14) when the first-order optimality condition (12.10) is not satisfied.

✐ **12.8** Verify that for the sequence $\{z_k\}$ defined by (12.30), the function $f(x) = x_1 + x_2$ satisfies $f(z_{k+1}) > f(z_k)$ for $k = 2, 3, \dots$. (Hint: Consider the trajectory $z(s) = (-\sqrt{2 - 1/s^2}, -1/s)^T$ and show that the function $h(s) \stackrel{\text{def}}{=} f(z(s))$ has $h'(s) > 0$ for all $s \geq 2$.)

✐ **12.9** Consider the problem (12.9). Specify *two* feasible sequences that approach the maximizing point $(1, 1)^T$, and show that neither sequence is a decreasing sequence for f.

✐ **12.10** Verify that neither the LICQ nor the MFCQ holds for the constraint set defined by (12.32) at $x^* = (0, 0)^T$.

✐ **12.11** Consider the feasible set Ω in $\mathbb{R}^2$ defined by $x_2 \geq 0$, $x_2 \leq x_1^2$.

(a) For $x^* = (0, 0)^T$, write down $T_\Omega(x^*)$ and $\mathcal{F}(x^*)$.

(b) Is LICQ satisfied at x^*? Is MFCQ satisfied?

(c) If the objective function is $f(x) = -x_2$, verify that that KKT conditions (12.34) are satisfied at x^*.

(d) Find a feasible sequence $\{z_k\}$ approaching x^* with $f(z_k) < f(x^*)$ for all k.

12.12 It is trivial to construct an example of a feasible set and a feasible point x^* at which the LICQ is satisfied but the constraints are nonlinear. Give an example of the reverse situation, that is, where the active constraints are linear but the LICQ is not satisfied.

12.13 Show that for the feasible region defined by

$$\begin{aligned}(x_1 - 1)^2 + (x_2 - 1)^2 &\le 2,\\ (x_1 - 1)^2 + (x_2 + 1)^2 &\le 2,\\ x_1 &\ge 0,\end{aligned}$$

the MFCQ is satisfied at $x^* = (0, 0)^T$ but the LICQ is not satisfied.

12.14 Consider the half space defined by $H = \{x \in \mathbb{R}^n \mid a^T x + \alpha \ge 0\}$ where $a \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$ are given. Formulate and solve the optimization problem for finding the point x in H that has the smallest Euclidean norm.

12.15 Consider the following modification of (12.36), where t is a parameter to be fixed prior to solving the problem:

$$\min_x \left(x_1 - \frac{3}{2}\right)^2 + (x_2 - t)^4 \qquad \text{s.t.} \qquad \begin{bmatrix} 1 - x_1 - x_2 \\ 1 - x_1 + x_2 \\ 1 + x_1 - x_2 \\ 1 + x_1 + x_2 \end{bmatrix} \ge 0. \qquad (12.97)$$

(a) For what values of t does the point $x^* = (1, 0)^T$ satisfy the KKT conditions?

(b) Show that when $t = 1$, only the first constraint is active at the solution, and find the solution.

12.16 (Fletcher [101]) Solve the problem

$$\min_x \; x_1 + x_2 \;\; \text{subject to } x_1^2 + x_2^2 = 1$$

by eliminating the variable x_2. Show that the choice of sign for a square root operation during the elimination process is critical; the "wrong" choice leads to an incorrect answer.

12.17 Prove that when the KKT conditions (12.34) and the LICQ are satisfied at a point x^*, the Lagrange multiplier λ^* in (12.34) is unique.

✐ **12.18** Consider the problem of finding the point on the parabola $y = \frac{1}{5}(x-1)^2$ that is closest to $(x, y) = (1, 2)$, in the Euclidean norm sense. We can formulate this problem as

$$\min\ f(x, y) = (x-1)^2 + (y-2)^2 \quad \text{subject to } (x-1)^2 = 5y.$$

(a) Find all the KKT points for this problem. Is the LICQ satisfied?

(b) Which of these points are solutions?

(c) By directly substituting the constraint into the objective function and eliminating the variable x, we obtain an unconstrained optimization problem. Show that the solutions of this problem cannot be solutions of the original problem.

✐ **12.19** Consider the problem

$$\min_{x \in \mathbb{R}^2} f(x) = -2x_1 + x_2 \quad \text{subject to} \quad \begin{cases} (1-x_1)^3 - x_2 & \geq & 0 \\ x_2 + 0.25x_1^2 - 1 & \geq & 0. \end{cases}$$

The optimal solution is $x^* = (0, 1)^T$, where both constraints are active.

(a) Do the LICQ hold at this point?

(b) Are the KKT conditions satisfied?

(c) Write down the sets $\mathcal{F}(x^*)$ and $\mathcal{C}(x^*, \lambda^*)$.

(d) Are the second-order necessary conditions satisfied? Are the second-order sufficient conditions satisfied?

✐ **12.20** Find the minima of the function $f(x) = x_1 x_2$ on the unit circle $x_1^2 + x_2^2 = 1$. Illustrate this problem geometrically.

✐ **12.21** Find the *maxima* of $f(x) = x_1 x_2$ over the unit disk defined by the inequality constraint $1 - x_1^2 - x_2^2 \geq 0$.

✐ **12.22** Show that for (12.1), the feasible set Ω is convex if c_i, $i \in \mathcal{E}$ are linear functions and $-c_i$, $i \in \mathcal{I}$ are convex functions.