Chapter 11
Constrained Minimization Conditions

We turn now, in this final part of the book, to the study of minimization problems having constraints. We begin by studying in this chapter the necessary and sufficient conditions satisfied at solution points. These conditions, aside from their intrinsic value in characterizing solutions, define Lagrange multipliers and a certain Hessian matrix which, taken together, form the foundation for both the development and analysis of algorithms presented in subsequent chapters.

The general method used in this chapter to derive necessary and sufficient conditions is a straightforward extension of that used in Chap. 7 for unconstrained problems. In the case of equality constraints, the feasible region is a curved surface embedded in E^n. Differential conditions satisfied at an optimal point are derived by considering the value of the objective function along curves on this surface passing through the optimal point. Thus the arguments run almost identically to those for the unconstrained case; families of curves on the constraint surface replacing the earlier artifice of considering feasible directions. There is also a theory of zero-order conditions that is presented in the final section of the chapter.

11.1 Constraints

We deal with general nonlinear programming problems of the form

$$\begin{aligned} &\text{minimize } f(\mathbf{x}) \\ &\text{subject to } h_1(\mathbf{x}) = 0, \quad g_1(\mathbf{x}) \leqslant 0 \\ &\qquad\qquad h_2(\mathbf{x}) = 0, \quad g_2(\mathbf{x}) \leqslant 0 \\ &\qquad\qquad \vdots \qquad\qquad\quad \vdots \\ &\qquad\qquad h_m(\mathbf{x}) = 0, \quad g_p(\mathbf{x}) \leqslant 0 \\ &\qquad\qquad \mathbf{x} \in \Omega \subset \mathrm{E}^n, \end{aligned} \tag{11.1}$$

D.G. Luenberger, Y. Ye, *Linear and Nonlinear Programming*, International Series in Operations Research & Management Science 228,
DOI 10.1007/978-3-319-18842-3_11

where $m \leqslant n$ and the functions f, $h_i, i = 1, 2, \ldots, m$ and g_j, $j = 1, 2, \ldots, p$ are continuous, and usually assumed to possess continuous second partial derivatives. For notational simplicity, we introduce the vector-valued functions $\mathbf{h} = (h_1, h_2, \ldots, h_m)$ and $\mathbf{g} = (g_1, g_2, \ldots, g_P)$ and rewrite (11.1) as

$$\begin{aligned} &\text{minimize } f(\mathbf{x}) \\ &\text{subject to } \mathbf{h}(\mathbf{x}) = \mathbf{0},\ \mathbf{g}(\mathbf{x}) \leqslant \mathbf{0} \\ &\qquad\qquad \mathbf{x} \in \Omega. \end{aligned} \tag{11.2}$$

The constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, $\mathbf{g}(\mathbf{x}) \leqslant \mathbf{0}$ are referred to as *functional constraints*, while the constraint $\mathbf{x} \in \Omega$ is a *set constraint*. As before we continue to de-emphasize the set constraint, assuming in most cases that either Ω is the whole space E^n or that the solution to (11.2) is in the interior of Ω. A point $\mathbf{x} \in \Omega$ that satisfies all the functional constraints is said to be *feasible*.

A fundamental concept that provides a great deal of insight as well as simplifying the required theoretical development is that of an *active constraint*. An inequality constraint $g_i(\mathbf{x}) \leqslant 0$ is said to be *active* at a feasible point $\mathbf{x}$ if $g_i(\mathbf{x}) = 0$ and *inactive* at $\mathbf{x}$ if $g_i(\mathbf{x}) < 0$. By convention we refer to any equality constraint $h_i(\mathbf{x}) = 0$ as *active* at any feasible point. The constraints active at a feasible point $\mathbf{x}$ restrict the domain of feasibility in neighborhoods of $\mathbf{x}$, while the other, inactive constraints, have no influence in neighborhoods of $\mathbf{x}$. Therefore, in studying the properties of a local minimum point, it is clear that attention can be restricted to the active constraints. This is illustrated in Fig. 11.1 where local properties satisfied by the solution $\mathbf{x}^*$ obviously do not depend on the inactive constraints g_2 and g_3.

It is clear that, if it were known a priori which constraints were active at the solution to (11.1), the solution would be a local minimum point of the problem defined by ignoring the inactive constraints and treating all active constraints as equality constraints. Hence, with respect to local (or relative) solutions, the problem could be regarded as having equality constraints only. This observation suggests that the

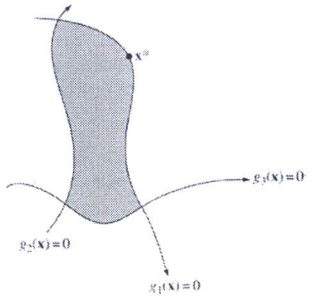

Fig. 11.1 Example of inactive constraints

majority of insight and theory applicable to (11.1) can be derived by consideration of equality constraints alone, later making additions to account for the selection of the active constraints. This is indeed so. Therefore, in the early portion of this chapter we consider problems having only equality constraints, thereby both economizing on notation and isolating the primary ideas associated with constrained problems. We then extend these results to the more general situation.

11.2 Tangent Plane

A set of equality constraints on E^n

$$\begin{aligned} h_1(\mathbf{x}) &= 0 \\ h_2(\mathbf{x}) &= 0 \\ &\vdots \\ h_m(\mathbf{x}) &= 0 \end{aligned} \qquad (11.3)$$

defines a subset of E^n which is best viewed as a hypersurface. If the constraints are everywhere regular, in a sense to be described below, this hypersurface is of dimension $n - m$. If, as we assume in this section, the functions h_i, $i = 1, 2, \ldots, m$ belong to C^1, the surface defined by them is said to be *smooth*.

Associated with a point on a smooth surface is the *tangent plane* at that point, a term which in two or three dimensions has an obvious meaning. To formalize the general notion, we begin by defining curves on a surface. A *curve* on a surface S is a family of points $\mathbf{x}(t) \in S$ continuously parameterized by t for $a \leqslant t \leqslant b$. The curve is *differentiable* if $\dot{\mathbf{x}} \equiv (\mathrm{d}/\mathrm{d}t)\mathbf{x}(t)$ exists, and is *twice differentiable* if $\ddot{\mathbf{x}}(t)$ exists. A curve $\mathbf{x}(t)$ is said to pass through the point $\mathbf{x}^*$ if $\mathbf{x}^* = \mathbf{x}(t^*)$ for some t^*, $a \leqslant t^* \leqslant b$. The derivative of the curve at $\mathbf{x}^*$ is, of course, defined as $\dot{\mathbf{x}}(t^*)$. It is itself a vector in E^n.

Now consider all differentiable curves on S passing through a point $\mathbf{x}^*$. The *tangent plane* at $\mathbf{x}^*$ is defined as the collection of the derivatives at $\mathbf{x}^*$ of all these differentiable curves. The tangent plane is a subspace of E^n.

For surfaces defined through a set of constraint relations such as (11.3), the problem of obtaining an explicit representation for the tangent plane is a fundamental problem that we now address. Ideally, we would like to express this tangent plane in terms of derivatives of functions h_i that define the surface. We introduce the subspace

$$M = \{\mathbf{y} : \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}\}$$

and investigate under what conditions M is equal to the tangent plane at $\mathbf{x}^*$. The key concept for this purpose is that of a *regular point*. Figure 11.2 shows some examples where for visual clarity the tangent planes (which are sub-spaces) are translated to the point $\mathbf{x}^*$.

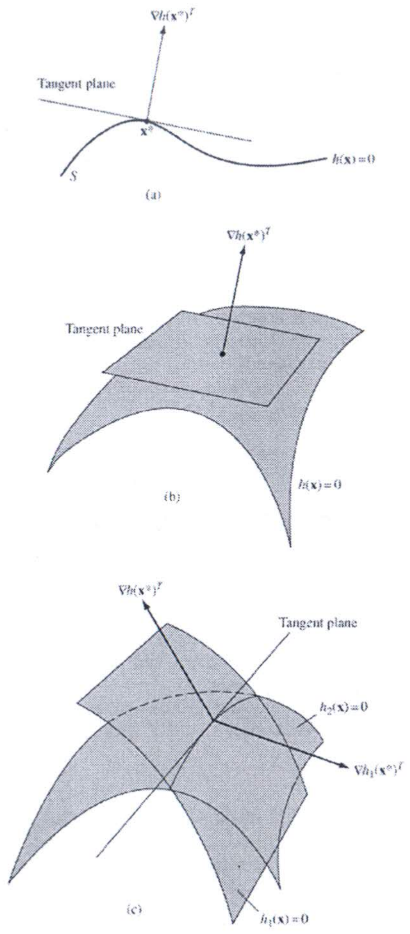

Fig. 11.2 Three examples of tangent planes (translated to $\mathbf{x}^*$)

Definition. A point $\mathbf{x}^*$ satisfying the constraint $\mathbf{h}(\mathbf{x}^*) = \mathbf{0}$ is said to be a *regular point* of the constraint if the gradient vectors $\nabla h_1(\mathbf{x}^*), \nabla h_2(\mathbf{x}^*), \ldots, \nabla h_m(\mathbf{x}^*)$ are linearly independent.

Note that if $\mathbf{h}$ is affine, $\mathbf{h}(\mathbf{x}) = \mathbf{A}\mathbf{x}+\mathbf{b}$, regularity is equivalent to A having rank equal to m, and this condition is independent of $\mathbf{x}$.

In general, at regular points it is possible to characterize the tangent plane in terms of the gradients of the constraint functions.

Theorem. *At a regular point $\mathbf{x}^*$ of the surface S defined by $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ the tangent plane is equal to*

$$M = \{\mathbf{y} : \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}\}.$$

Proof. Let T be the tangent plane at $\mathbf{x}^*$. It is clear that $T \subset M$ whether $\mathbf{x}^*$ is regular or not, for any curve $\mathbf{x}(t)$ passing through $\mathbf{x}^*$ at $t = t^*$ having derivative $\dot{\mathbf{x}}(t^*)$ such that $\nabla \mathbf{h}(\mathbf{x}^*)\dot{\mathbf{x}}(t^*) \neq \mathbf{0}$ would not lie on S.

To prove that $M \subset T$ we must show that if $\mathbf{y} \in M$ then there is a curve on S passing through $\mathbf{x}^*$ with derivative $\mathbf{y}$. To construct such a curve we consider the equations

$$\mathbf{h}(\mathbf{x}^* + t\mathbf{y} + \nabla \mathbf{h}(\mathbf{x}^*)^T \mathbf{u}(t)) = \mathbf{0}, \tag{11.4}$$

where for fixed t we consider $\mathbf{u}(t) \in E^m$ to be the unknown. This is a nonlinear system of m equations and m unknowns, parameterized continuously, by t. At $t = 0$ there is a solution $\mathbf{u}(0) = \mathbf{0}$. The Jacobian matrix of the system with respect to $\mathbf{u}$ at $t = 0$ is the $m \times m$ matrix

$$\nabla \mathbf{h}(\mathbf{x}^*)\nabla \mathbf{h}(\mathbf{x}^*)^T,$$

which is nonsingular, since $\nabla \mathbf{h}(\mathbf{x}^*)$ is of full rank if $\mathbf{x}^*$ is a regular point. Thus, by the Implicit Function Theorem (see Appendix A) there is a continuously differentiable solution $\mathbf{u}(t)$ in some region $-a \leqslant t \leqslant a$.

The curve $\mathbf{x}(t) = \mathbf{x}^* + t\mathbf{y} + \nabla \mathbf{h}(\mathbf{x}^*)^T \mathbf{u}(t)$ is thus, by construction, a curve on S. By differentiating the system (11.4) with respect to t at $t = 0$ we obtain

$$\mathbf{0} = \frac{d}{dt}\mathbf{h}(\mathbf{x}(t))\bigg|_{t=0} = \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} + \nabla \mathbf{h}(\mathbf{x}^*)\nabla \mathbf{h}(\mathbf{x}^*)^T \dot{\mathbf{u}}(0).$$

By definition of $\mathbf{y}$ we have $\nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}$ and thus, again since $\nabla \mathbf{h}(\mathbf{x}^*)\nabla \mathbf{h}(\mathbf{x}^*)^T$ is nonsingular, we conclude that $\dot{\mathbf{x}}(0) = \mathbf{0}$. Therefore

$$\dot{\mathbf{x}}(0) = \mathbf{y} + \nabla \mathbf{h}(\mathbf{x}^*)^T \dot{\mathbf{x}}(0) = \mathbf{y},$$

and the constructed curve has derivative $\mathbf{y}$ at $\mathbf{x}^*$. ∎

It is important to recognize that the condition of being a regular point is not a condition on the constraint surface itself but on its representation in terms of an $\mathbf{h}$. The tangent plane is defined independently of the representation, while M is not.

Example. In E^2 let $h(x_1, x_2) = x_1$. Then $h(\mathbf{x}) = 0$ yields the x_2 axis, and every point on that axis is regular. If instead we put $h(x_1, x_2) = x_1^2$, again S is the x_2 axis but now no point on the axis is regular. Indeed in this case $M = E^2$, while the tangent plane is the x_2 axis.

11.3 First-Order Necessary Conditions (Equality Constraints)

The derivation of necessary and sufficient conditions for a point to be a local minimum point subject to equality constraints is fairly simple now that the representation of the tangent plane is known. We begin by deriving the first-order necessary conditions.

Lemma. *Let $\mathbf{x}^*$ be a regular point of the constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ and a local extremum point (a minimum or maximum) of f subject to these constraints.*

Then all $\mathbf{y} \in E^n$ satisfying

$$\nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0} \tag{11.5}$$

must also satisfy

$$\nabla f(\mathbf{x}^*)\mathbf{y} = 0. \tag{11.6}$$

Proof. Let $\mathbf{y}$ be any vector in the tangent plane at $\mathbf{x}^*$ and let $\mathbf{x}(t)$ be any smooth curve on the constraint surface passing through $\mathbf{x}^*$ with derivative $\mathbf{y}$ at $\mathbf{x}^*$: that is, $\mathbf{x}(0) = \mathbf{x}^*$, $\dot{\mathbf{x}}(0) = \mathbf{y}$, and $\mathbf{h}(\mathbf{x}(\mathbf{t})) = \mathbf{0}$ for $-a \leqslant t \leqslant a$ for some $a > 0$.

Since $\mathbf{x}^*$ is a regular point, the tangent plane is identical with the set of $\mathbf{y}$'s satisfying $\nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}$. Then, since $\mathbf{x}^*$ is a constrained local extremum point of f, we have

$$\frac{d}{dt}f(\mathbf{x}(t))\bigg|_{t=0} = 0,$$

or equivalently,

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

The above Lemma says that $\nabla f(\mathbf{x}^*)$ is orthogonal to the tangent plane. Next we conclude that this implies that $\nabla f(\mathbf{x}^*)$ is a linear combination of the gradients of $\mathbf{h}$ at $\mathbf{x}^*$, a relation that leads to the introduction of Lagrange multipliers. As in much of nonlinear programming, the Lagrange multiplier vector is often labeled λ rather than $\mathbf{y}$ in linear programming, and this convention is followed here.

Theorem. *Let $\mathbf{x}^*$ be a local extremum point of f subject to the constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$. Assume further that $\mathbf{x}^*$ is a regular point of these constraints. Then there is a $\boldsymbol{\lambda} \in E^m$ such that*

$$\nabla f(\mathbf{x}^*) + \lambda^T \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}. \tag{11.7}$$

Proof. From the Lemma we may conclude that the value of the linear program

$$\begin{aligned} &\text{maximize} && \nabla f(\mathbf{x}^*)\mathbf{y} \\ &\text{subject to} && \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0} \end{aligned}$$

is zero. Thus, by the Duality Theorem of linear programming (Sect. 4.2) the dual problem is feasible. Specifically, there is $\boldsymbol{\lambda} \in E^m$ such that $\nabla f(\mathbf{x}^*) + \lambda^T \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$. ∎

It should be noted that the first-order necessary conditions

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

together with the constraints

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

give a total of $n+m$ (generally nonlinear) equations in the $n+m$ variables comprising $\mathbf{x}^*$, $\boldsymbol{\lambda}$. Thus the necessary conditions are a complete set since, at least locally, they determine a unique solution.

It is convenient to introduce the *Lagrangian* associated with the constrained problem, defined as

$$l(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{h}(\mathbf{x}). \tag{11.8}$$

The necessary conditions can then be expressed in the form

$$\nabla_{\mathbf{x}} l(\mathbf{x}, \boldsymbol{\lambda}) = \mathbf{0} \tag{11.9}$$

$$\nabla_{\boldsymbol{\lambda}} l(\mathbf{x}, \boldsymbol{\lambda}) = \mathbf{0}, \tag{11.10}$$

the second of these being simply a restatement of the constraints.

11.4 Examples

We digress briefly from our mathematical development to consider some examples of constrained optimization problems. We present five simple examples that can be treated explicitly in a short space and then briefly discuss a broader range of applications.

Example 1. Consider the problem

$$\begin{aligned} &\text{minimize} \quad x_1x_2 + x_2x_3 + x_1x_3 \\ &\text{subject to} \quad x_1 + x_2 + x_3 = 3. \end{aligned}$$

The necessary conditions become

$$\begin{aligned} x_2 + x_3 + \lambda &= 0 \\ x_1 \quad + x_3 + \lambda &= 0 \\ x_1 + x_2 \quad + \lambda &= 0. \end{aligned}$$

These three equations together with the one constraint equation give four equations that can be solved for the four unknowns x_1, x_2, x_3, λ. Solution yields $x_1 = x_2 = x_3 = 1$, $\lambda = -2$.

Example 2 (Maximum Volume). Let us consider an example of the type that is now standard in textbooks and which has a structure similar to that of the example above.

We seek to construct a cardboard box of maximum volume, given a fixed area of cardboard.

Denoting the dimensions of the box by x, y, z, the problem can be expressed as

$$\begin{aligned} &\text{maximize} \quad xyz \\ &\text{subject to} \quad (xy + yz + xz) = \frac{c}{2}, \end{aligned} \tag{11.11}$$

where $c > 0$ is the given area of cardboard. Introducing a Lagrange multiplier, the first-order necessary conditions are easily found to be

$$\begin{aligned} yz + \lambda(y + z) &= 0 \\ xz + \lambda(x + z) &= 0 \\ xy + \lambda(x + y) &= 0 \end{aligned} \tag{11.12}$$

together with the constraint. Before solving these, let us note that the sum of these equations is $(xy + yz + xz) + 2\lambda(x + y + z) = 0$. Using the constraint this becomes $c/2 + 2\lambda(x + y + z) = 0$. From this it is clear that $\lambda \neq 0$. Now we can show that x, y, and z are nonzero. This follows because $x = 0$ implies $z = 0$ from the second equation and $y = 0$ from the third equation. In a similar way, it is seen that if either x, y, or z are zero, all must be zero, which is impossible.

To solve the equations, multiply the first by x and the second by y, and then subtract the two to obtain

$$\lambda(x - y)z = 0.$$

Operate similarly on the second and third to obtain

$$\lambda(y - z)x = 0.$$

Since no variables can be zero, it follows that $x = y = z = \sqrt{c/6}$ is the unique solution to the necessary conditions. The box must be a cube.

Example 3 (Entropy). optimization problems often describe natural phenomena. An example is the characterization of naturally occurring probability distributions as maximum entropy distributions.

As a specific example consider a discrete probability density corresponding to a measured value taking one of n values x_1, x_2, ..., x_n. The probability associated with x_i is p_i. The p_i's satisfy $p_i \geqslant 0$ and $\sum_{i=1}^{n} p_i = 1$.

The *entropy* of such a density is

$$\varepsilon = -\sum_{i=1}^{n} p_i \log(p_i).$$

The *mean value* of the density is $\sum_{i=1}^{n} x_i p_i$.

If the value of mean is known to be m (by the physical situation), the maximum entropy argument suggests that the density should be taken as that which solves the following problem:

$$\begin{aligned} &\text{maximize} \quad -\sum_{i=1}^{n} p_i \log(p_i) \\ &\text{subject to} \quad \sum_{i=1}^{n} p_i = 1 \\ &\qquad \sum_{i=1}^{n} x_i p_i = m \\ &\qquad p_i \geq 0, \quad i = 1, 2, \ldots, n. \end{aligned} \tag{11.13}$$

We begin by ignoring the nonnegativity constraints, believing that they may be inactive. Introducing two Lagrange multipliers, λ and μ, the Lagrangian is

$$l = \sum_{i=1}^{n} \{-p_i \log p_i + \lambda p_i + \mu x_i p_i\} - \lambda - \mu m.$$

The necessary conditions are immediately found to be

$$-\log p_i - 1 + \lambda + \mu x_i = 0, \quad i = 1, 2, \ldots, n.$$

This leads to

$$p_i = \exp\{(\lambda - 1) + \mu x_i\}, \quad i = 1, 2, \ldots, n. \tag{11.14}$$

We note that $p_i > 0$, so the nonnegativity constraints are indeed inactive. The result (11.14) is known as an exponential density. The Lagrange multipliers λ and μ are parameters that must be selected so that the two equality constraints are satisfied.

Example 4 (Hanging Chain). A chain is suspended from two thin hooks that are 16 ft apart on a horizontal line as shown in Fig. 11.3. The chain itself consists of 20 links of stiff steel. Each link is one foot in length (measured inside). We wish to formulate the problem to determine the equilibrium shape of the chain.

The solution can be found by minimizing the potential energy of the chain. Let us number the links consecutively from 1 to 20 starting with the left end. We let link i span an x distance of x_i and a y distance of y_i. Then $x_i^2 + y_i^2 = 1$. The potential energy of a link is its weight times its vertical height (from some reference). The potential energy of the chain is the sum of the potential energies of each link. We may take the top of the chain as reference and assume that the mass of each link is concentrated at its center. Assuming unit weight, the potential energy is then

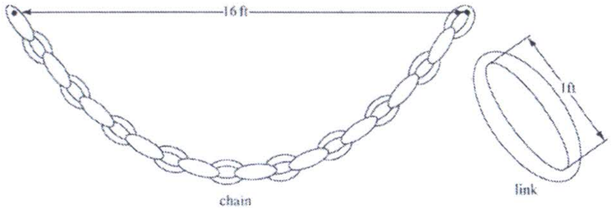

Fig. 11.3 A hanging chain

$$\frac{1}{2}y_1 + \left(y_1 + \frac{1}{2}y_2\right) + \left(y_1 + y_2 + \frac{1}{2}y_3\right) + \cdots + \left(y_1 + y_2 + \cdots + y_{n-1} + \frac{1}{2}y_n\right) = \sum_{i=1}^{n}\left(n - i + \frac{1}{2}\right)y_i,$$

where $n = 20$ in our example.

The chain is subject to two constraints: The total y displacement is zero, and the total x displacement is 16. Thus the equilibrium shape is the solution of

$$\begin{aligned} &\text{minimize} \quad \sum_{i=1}^{n}\left(n - i + \frac{1}{2}\right)y_i \\ &\text{subject to} \quad \sum_{i=1}^{n} y_i = 0 \\ &\qquad \sum_{i=1}^{n} \sqrt{1 - y_i^2} = 16. \end{aligned} \tag{11.15}$$

The first-order necessary conditions are

$$\left(n - i + \frac{1}{2}\right) + \lambda - \frac{\mu y_i}{\sqrt{1 - y_i^2}} = 0 \tag{11.16}$$

for $i = 1, 2, \ldots, n$. This leads directly to

$$y_i = -\frac{n - i + \frac{1}{2} + \lambda}{\sqrt{\mu^2 + \left(n - i + \frac{1}{2} + \lambda\right)^2}}. \tag{11.17}$$

As in Example 2 the solution is determined once the Lagrange multipliers are known. They must be selected so that the solution satisfies the two constraints.

It is useful to point out that problems of this type may have local minimum points. The reader can examine this by considering a short chain of, say, four links and v and w configurations.

Example 5 (Portfolio Design). Suppose there are n securities indexed by $i = 1, 2, \ldots, n$. Each security i is characterized by its random rate of return r_i which has mean value $\bar{r}_i$. Its covariances with the rates of return of other securities are σ_{ij}, for $j = 1, 2, \ldots, n$. The portfolio problem is to allocate total available wealth among these n securities, allocating a fraction w_i of wealth to the security i.

The overall rate of return of a portfolio is $r = \sum_{i=1}^n w_i\bar{r}_i$ and variance $\sigma^2 = \sum_{i,j=1}^n w_i\sigma_{ij}w_j$.

Markowitz introduced the concept of devising *efficient* portfolios which for a given expected rate of return $\bar{r}$ have minimum possible variance. Such a portfolio is the solution to the problem

$$\begin{aligned} \min_{w_1, w_2, \ldots, w_n} & \sum_{i,j=1}^n w_i\sigma_{ij}w_j \\ \text{subject to} & \sum_{i=1}^n w_i\bar{r}_i = \bar{r} \\ & \sum_{i=1}^n w_i = 1. \end{aligned}$$

The second constraint forces the sum of the weights to equal one. There may be the further restriction that each $w_i \geq 0$ which would imply that the securities must not be shorted (that is, sold short).

Introducing Lagrange multipliers λ and μ for the two constraints leads easily to the $n + 2$ linear equations

$$\sum_{j=1}^n \sigma_{ij}w_j + \lambda\bar{r}_i + \mu = 0 \quad \text{for } i = 1, 2, \ldots, n$$

$$\sum_{i=1}^n w_i\bar{r}_i = \bar{r}$$

$$\sum_{i=1}^n w_i = 1$$

in the $n + 2$ unknowns (the w_i's, λ and μ).

Large-Scale Applications

The problems that serve as the primary motivation for the methods described in this part of the book are actually somewhat different in character than the problems represented by the above examples, which by necessity are quite simple. Larger, more complex, nonlinear programming problems arise frequently in modern applied analysis in a wide variety of disciplines. Indeed, within the past few decades nonlinear programming has advanced from a relatively young and primarily analytic subject to a substantial general tool for problem solving.

Large nonlinear programming problems arise in problems of mechanical structures, such as determining optimal configurations for bridges, trusses, and so forth. Some mechanical designs and configurations that in the past were found by solving differential equations are now often found by solving suitable optimization problems. An example that is somewhat similar to the hanging chain problem is the determination of the shape of a stiff cable suspended between two points and supporting a load.

A wide assortment of large-scale optimization problems arise in a similar way as methods for solving partial differential equations. In situations where the underlying continuous variables are defined over a two- or three-dimensional region, the continuous region is replaced by a grid consisting of perhaps several thousand discrete points. The corresponding discrete approximation to the partial differential equation is then solved indirectly by formulating an equivalent optimization problem. This approach is used in studies of plasticity, in heat equations, in the flow of fluids, in atomic physics, and indeed in almost all branches of physical science.

Problems of optimal control lead to large-scale nonlinear programming problems. In these problems a dynamic system, often described by an ordinary differential equation, relates control variables to a trajectory of the system state. This differential equation, or a discretized version of it, defines one set of constraints. The problem is to select the control variables so that the resulting trajectory satisfies various additional constraints and minimizes some criterion. An early example of such a problem that was solved numerically was the determination of the trajectory of a rocket to the moon that required the minimum fuel consumption.

There are many examples of nonlinear programming in industrial operations and business decision making. Many of these are nonlinear versions of the kinds of examples that were discussed in the linear programming part of the book. Nonlinearities can arise in production functions, cost curves, and, in fact, in almost all facets of problem formulation.

Portfolio analysis, in the context of both stock market investment and evaluation of a complex project within a firm, is an area where nonlinear programming is becoming increasingly useful. These problems can easily have thousands of variables.

In many areas of model building and analysis, optimization formulations are increasingly replacing the direct formulation of systems of equations. Thus large economic forecasting models often determine equilibrium prices by minimizing an objective termed *consumer surplus*. Physical models are often formulated as minimization of energy. Decision problems are formulated as maximizing expected utility. Data analysis procedures are based on minimizing an average error or maximizing a probability. As the methodology for solution of nonlinear programming improves, one can expect that this trend will continue.

11.5 Second-Order Conditions

By an argument analogous to that used for the unconstrained case, we can also derive the corresponding second-order conditions for constrained problems. Throughout this section it is assumed that f, $\mathbf{h} \in C^2$.

Second-Order Necessary Conditions. *Suppose that $\mathbf{x}^*$ is a local minimum of f subject to $\mathbf{h}(\mathbf{x}) = 0$ and that $\mathbf{x}^*$ is a regular point of these constraints. Then there is a $\lambda \in E^m$ such that*

$$\nabla f(\mathbf{x}^*) + \lambda^T \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}. \tag{11.18}$$

If we denote by M the tangent plane $M = \{\mathbf{y} : \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}\}$, then the matrix

$$\mathbf{L}(\mathbf{x}^*) = \mathbf{F}(\mathbf{x}^*) + \lambda^T \mathbf{H}(\mathbf{x}^*) \tag{11.19}$$

is positive semidefinite on M, that is, $\mathbf{y}^T \mathbf{L}(\mathbf{x}^*)\mathbf{y} \geqslant 0$ for all $\mathbf{y} \in M$.

Proof. From elementary calculus it is clear that for every twice differentiable curve on the constraint surface S through $\mathbf{x}^*$ (with $\mathbf{x}(0) = \mathbf{x}^*$) we have

$$\left.\frac{d^2}{dt^2} f(\mathbf{x}(t))\right|_{t=0} \geqslant 0. \tag{11.20}$$

By definition

$$\left.\frac{d^2}{dt^2} f(\mathbf{x}(t))\right|_{t=0} = \dot{\mathbf{x}}(0)^T \mathbf{F}(\mathbf{x}^*)\dot{\mathbf{x}}(0) + \nabla f(\mathbf{x}^*)\ddot{\mathbf{x}}(0). \tag{11.21}$$

Furthermore, differentiating the relation $\lambda^T \mathbf{h}(\mathbf{x}(t)) = 0$ twice, we obtain

$$\dot{\mathbf{x}}(0)^T \lambda^T \mathbf{H}(\mathbf{x}^*)\dot{\mathbf{x}}(0) + \lambda^T \nabla \mathbf{h}(\mathbf{x}^*)\ddot{\mathbf{x}}(\mathbf{0}) = 0. \tag{11.22}$$

Adding (11.22) to (11.21), while taking account of (11.20), yields the result

$$\left.\frac{d^2}{dt^2} f(\mathbf{x}(t))\right|_{t=0} = \dot{\mathbf{x}}(0)^T \mathbf{L}(\mathbf{x}^*)\dot{\mathbf{x}}(0) \geqslant 0.$$

Since $\dot{\mathbf{x}}(0)$ is arbitrary in M, we immediately have the stated conclusion. ∎

The above theorem is our first encounter with the matrix $\mathbf{L} = \mathbf{F} + \lambda^T \mathbf{H}$ which is the matrix of second partial derivatives, with respect to $\mathbf{x}$, of the Lagrangian l. (See Appendix A, Sect. A.6, for a discussion of the notation $\lambda^T \mathbf{H}$ used here.) This matrix is the backbone of the theory of algorithms for constrained problems, and it is encountered often in subsequent chapters.

We next state the corresponding set of sufficient conditions.

Second-Order Sufficiency Conditions. *Suppose there is a point $\mathbf{x}^*$ satisfying $\mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, and a $\lambda \in E^m$ such that*

$$\nabla f(\mathbf{x}^*) + \lambda^T \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}. \tag{11.23}$$

Suppose also that the matrix $\mathbf{L}(\mathbf{x}^*) = \mathbf{F}(\mathbf{x}^*) + \lambda^T \mathbf{H}(\mathbf{x}^*)$ is positive definite on $M = \{\mathbf{y} : \nabla \mathbf{h}(\mathbf{x}^*)\mathbf{y} = \mathbf{0}\}$, that is, for $\mathbf{y} \in M$, $\mathbf{y} \neq \mathbf{0}$ there holds $\mathbf{y}^T \mathbf{L}(\mathbf{x}^*)\mathbf{y} > 0$. Then $\mathbf{x}^*$ is a strict local minimum of f subject to $\mathbf{h}(\mathbf{x}) = \mathbf{0}$.

Proof. If $\mathbf{x}^*$ is not a strict relative minimum point, there exists a sequence of feasible points $\{\mathbf{y}_k\}$ converging to $\mathbf{x}^*$ such that for each k, $f(\mathbf{y}_k) \leqslant f(\mathbf{x}^*)$. Write each $\mathbf{y}_k$ in the form $\mathbf{y}_k = \mathbf{x}^* + \delta_k \mathbf{s}_k$ where $\mathbf{s}_k \in \mathbf{E}^n$, $|\mathbf{s}_k| = 1$, and $\delta_k > 0$ for each k. Clearly, $\delta_k \to 0$ and the sequence $\{\mathbf{s}_k\}$, being bounded, must have a convergent subsequence converging to some $\mathbf{s}^*$. For convenience of notation, we assume that the sequence $\{\mathbf{s}_k\}$ is itself convergent to $\mathbf{s}^*$. We also have $\mathbf{h}(\mathbf{y}_k) - \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, and dividing by δ_k and letting $k \to \infty$ we see that $\nabla \mathbf{h}(\mathbf{x}^*)\mathbf{s}^* = \mathbf{0}$.

Now by Taylor's theorem, we have for each j

$$0 = h_j(\mathbf{y}_k) = h_j(\mathbf{x}^*) + \delta_k \nabla h_j(\mathbf{x}^*)\mathbf{s}_k + \frac{\delta_k^2}{2}\mathbf{s}_k^T \nabla^2 h_j(\eta_j)\mathbf{s}_k \tag{11.24}$$

and

$$0 \geqslant f(\mathbf{y}_k) - f(\mathbf{x}^*) = \delta_k \nabla f(\mathbf{x}^*)\mathbf{s}_k + \frac{\delta_k^2}{2}\mathbf{s}_k^T \nabla^2 f(\eta_0)\mathbf{s}_k, \tag{11.25}$$

where each η_j is a point on the line segment joining $\mathbf{x}^*$ and $\mathbf{y}_k$. Multiplying (11.24) by λ_j and adding these to (11.25) we obtain, on accounting for (11.23),

$$0 \geqslant \frac{\delta_k^2}{2}\mathbf{s}_k^T \left\{ \nabla^2 f(\eta_0) + \sum_{i=1}^{m} \lambda_i \nabla^2 h_i(\eta_i) \right\} \mathbf{s}_k,$$

which yields a contradiction as $k \to \infty$. ∎

Example 1. Consider the problem

$$\begin{array}{ll} \text{maximize} & x_1x_2 + x_2x_3 + x_1x_3 \\ \text{subject to} & x_1 + x_2 + x_3 = 3. \end{array}$$

In Example 1 of Sect. 11.4 it was found that $x_1 = x_2 = x_3 = 1$, $\lambda = -2$ satisfy the first-order conditions. The matrix $\mathbf{F} + \lambda^T \mathbf{H}$ becomes in this case

$$L = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix},$$

which itself is neither positive nor negative definite. On the subspace $M = \{\mathbf{y}: y_1 + y_2 + y_3 = 0\}$, however, we note that

$$\begin{aligned} \mathbf{y}^T \mathbf{L}\mathbf{y} &= y_1(y_2 + y_3) + y_2(y_1 + y_3) + y_3(y_1 + y_2) \\ &= -(y_1^2 + y_2^2 + y_3^2), \end{aligned}$$

and thus $\mathbf{L}$ is negative definite on M. Therefore, the solution we found is at least a local maximum.

11.6 Eigenvalues in Tangent Subspace

In the last section it was shown that the matrix **L** restricted to the subspace M that is tangent to the constraint surface plays a role in second-order conditions entirely analogous to that of the Hessian of the objective function in the unconstrained case. It is perhaps not surprising, in view of this, that the structure of **L** restricted to M also determines rates of convergence of algorithms designed for constrained problems in the same way that the structure of the Hessian of the objective function does for unconstrained algorithms. Indeed, we shall see that the eigenvalues of **L** restricted to M determine the natural rates of convergence for algorithms designed for constrained problems. It is important, therefore, to understand what these restricted eigenvalues represent. We first determine geometrically what we mean by the restriction of **L** to M which we denote by $\mathbf{L}_M$. Next we define the eigenvalues of the operator $\mathbf{L}_M$. Finally we indicate how these various quantities can be computed.

Given any vector $\mathbf{y} \in M$, the vector $\mathbf{Ly}$ is in E^n but not necessarily in M. We project $\mathbf{Ly}$ orthogonally back onto M, as shown in Fig. 11.4, and the result is said to be the restriction of **L** to M operating on **y**. In this way we obtain a linear transformation from M to M. The transformation is determined somewhat implicitly, however, since we do not have an explicit matrix representation.

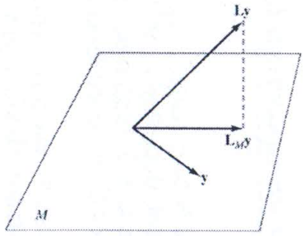

Fig. 11.4 Definition of $\mathbf{L}_M$

A vector $\mathbf{y} \in M$ is an *eigenvector* of $\mathbf{L}_M$ if there is a real number λ such that $\mathbf{L}_M\mathbf{y} = \lambda\mathbf{y}$; the corresponding λ is an *eigenvalue* of $\mathbf{L}_M$. This coincides with the standard definition. In terms of **L** we see that **y** is an eigenvector of $\mathbf{L}_M$ if $\mathbf{Ly}$ can be written as the sum of $\lambda\mathbf{y}$ and a vector orthogonal to M. See Fig. 11.5.

To obtain a matrix representation for $\mathbf{L}_M$ it is necessary to introduce a basis in the subspace M. For simplicity it is best to introduce an orthonormal basis, say $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{n-m}$. Define the matrix **E** to be the $n \times (n-m)$ matrix whose columns consist of the vectors $\mathbf{e}_i$. Then any vector **y** in M can be written as $\mathbf{y} = \mathbf{Ez}$ for some $\mathbf{z} \in E^{n-m}$ and, of course, $\mathbf{LEz}$ represents the action of **L** on such a vector. To project this result back into M and express the result in terms of the basis $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{n-m}$,

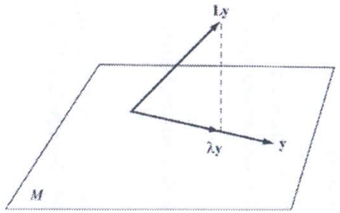

Fig. 11.5 Eigenvector of $\mathbf{L}_M$

we merely multiply by $\mathbf{E}^T$. Thus $\mathbf{E}^T\mathbf{LEz}$ is the vector whose components give the representation in terms of the basis; and, correspondingly, the $(n-m) \times (n-m)$ matrix $\mathbf{E}^T\mathbf{LE}$ is the matrix representation of **L** restricted to M.

The eigenvalues of **L** restricted to M can be found by determining the eigenvalues of $\mathbf{E}^T\mathbf{LE}$. These eigenvalues are independent of the particular orthonormal basis **E**.

Example 1. In the last section we considered

$$L = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

restricted to $M = \{\mathbf{y} : y_1 + y_2 + y_3 = 0\}$. To obtain an explicit matrix representation on M let us introduce the orthonormal basis:

$$\mathbf{e}_1 = \frac{1}{\sqrt{2}}(1, 0, -1)$$
$$\mathbf{e}_2 = \frac{1}{\sqrt{6}}(1, -2, 1).$$

This gives, upon expansion,

$$\mathbf{E}^T\mathbf{LE} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix},$$

and hence **L** restricted to M acts like the negative of the identity.

Example 2. Let us consider the problem

$$\begin{aligned} &\text{extremize} && x_1 + x_2^2 + x_2x_3 + 2x_3^2 \\ &\text{subject to} && \frac{1}{2}(x_1^2 + x_2^2 + x_3^2) = 1. \end{aligned}$$

The first-order necessary conditions are

$$\begin{aligned} 1 + \quad \lambda x_1 &= 0 \\ 2x_2 + x_3 + \lambda x_2 &= 0 \\ x_2 + 4x_3 + \lambda x_3 &= 0. \end{aligned}$$

One solution to this set is easily seen to be $x_1 = 1$, $x_2 = 0$, $x_3 = 0$, $\lambda = -1$. Let us examine the second-order conditions at this solution point. The Lagrangian matrix there is

$$\mathbf{L} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 3 \end{bmatrix},$$

and the corresponding subspace M is

$$M = \{\mathbf{y} : y_1 = 0\}.$$

In this case M is the subspace spanned by the second two basis vectors in E^3 and hence the restriction of $\mathbf{L}$ to M can be found by taking the corresponding submatrix of $\mathbf{L}$. Thus, in this case,

$$\mathbf{E}^T\mathbf{L}\mathbf{E} = \begin{bmatrix} 1 & 1 \\ 1 & 3 \end{bmatrix}.$$

The characteristic polynomial of this matrix is

$$\det \begin{bmatrix} 1-\lambda & 1 \\ 1 & 3-\lambda \end{bmatrix} = (1-\lambda)(3-\lambda) - 1 = \lambda^2 - 4\lambda + 2.$$

The eigenvalues of $\mathbf{L}_M$ are thus $\lambda = 2 \pm \sqrt{2}$, and $\mathbf{L}_M$ is positive definite.

Since the $\mathbf{L}_M$ matrix is positive definite, we conclude that the point found is a relative minimum point. This example illustrates that, in general, the restriction of $\mathbf{L}$ to M can be thought of as a submatrix of $\mathbf{L}$, although it can be read directly from the original matrix only if the subspace M is spanned by a subset of the original basis vectors.

Projected Hessians

The above approach for determining the eigenvalues of $\mathbf{L}$ projected onto M is quite direct and relatively simple. There is another approach, however, that is useful in some theoretical arguments and convenient for simple applications. It is based on constructing matrices and determinants of order n rather than $n - m$, but there is no need to find the orthonormal basis $\mathbf{E}$. For simplicity, let $\mathbf{A} = \nabla\mathbf{h}$ which has full row rank.

Any $\mathbf{x}$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{0}$ can be expressed as

$$\mathbf{x} = (\mathbf{I} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A})\mathbf{z}$$

for some $\mathbf{z}$ (and the converse is also true), where $\mathbf{P}_A = (\mathbf{I} - \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)^{-1}\mathbf{A})$ is the so called projection matrix to the null space of $\mathbf{A}$ (that is, M). If $\mathbf{x}^T\mathbf{L}\mathbf{x} \geq 0$ for all $\mathbf{x}$ in this null space, then $\mathbf{z}^T\mathbf{P}_A\mathbf{L}\mathbf{P}_A\mathbf{z} \geq 0$ for all $\mathbf{z} \in E^n$, or the n-dimensional symmetric matrix $\mathbf{P}_A\mathbf{L}\mathbf{P}_A$ is positive semidefinite. Furthermore, if $\mathbf{P}_A\mathbf{L}\mathbf{P}_A$ has rank $n - m$, then $\mathbf{L}_M$ is positive definite, which results the following test.

Projected Hessian Test. *The matrix* $\mathbf{L}$ *is positive definite on the subspace* $M = \{\mathbf{x} : \nabla\mathbf{h}\mathbf{x} = \mathbf{0}\}$ *if and only if the projected Hessian matrix to the null space of* $\nabla\mathbf{h}$ *is positive semidefinite and has rank* $n - m$.

Example 3. Approaching Example 2 in this way and noting $\mathbf{A} = \nabla\mathbf{h} = (1, 0, 0)$ we have

$$\mathbf{P}_A = \mathbf{I} - \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}^T = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Then

$$\mathbf{P}_A\mathbf{L}\mathbf{P}_A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 3 \end{bmatrix}$$

which is clearly positive semidefinite and has rank 2.

11.7 Sensitivity

The Lagrange multipliers associated with a constrained minimization problem have an interpretation as prices, similar to the prices associated with constraints in linear programming. In the nonlinear case the multipliers are associated with the particular solution point and correspond to incremental or marginal prices, *that* is, prices associated with small variations in the constraint requirements.

Suppose the problem

$$\begin{aligned} &\text{minimize} \quad f(\mathbf{x}) \\ &\text{subject to} \quad \mathbf{h}(\mathbf{x}) = \mathbf{0} \end{aligned} \tag{11.26}$$

has a solution at the point $\mathbf{x}^*$ which is a regular point of the constraints. Let $\boldsymbol{\lambda}$ be the corresponding Lagrange multiplier vector. Now consider the family of problems

$$\begin{aligned} &\text{minimize} \quad f(\mathbf{x}) \\ &\text{subject to} \quad \mathbf{h}(\mathbf{x}) = \mathbf{c}, \end{aligned} \tag{11.27}$$

where $\mathbf{c} \in E^m$. For a sufficiently small range of $\mathbf{c}$ near the zero vector, the problem will have a solution point $\mathbf{x}(\mathbf{c})$ near $\mathbf{x}(\mathbf{0}) \equiv \mathbf{x}^*$. For each of these solutions there is a