Practical Methods of Optimization

Volume 1
Unconstrained Optimization

R. Fletcher

Department of Mathematics
University of Dundee, Scotland, U.K.

A Wiley–Interscience Publication

JOHN WILEY & SONS

Chichester · New York · [illegible] · Toronto

British Library Cataloguing in Publication Data:

Fletcher, Roger
Practical methods of optimization.
Vol. 1: Unconstrained optimization
1. Mathematical optimization
I. Title
515 QA402.5

ISBN 0 471 27711 8

Typeset by Preface Ltd, Salisbury, Wilts. Printed in the United States of America.

Contents

Preface vii

Chapter 1 Introduction 1
1.1 History and Applications 1
1.2 Mathematical Background 4
Questions for Chapter 1 9

Chapter 2 Structure of Methods 10
2.1 Conditions for Local Minima 10
2.2 Ad hoc Methods 13
2.3 Useful Algorithmic Properties 16
2.4 Descent Methods and Stability 20
2.5 Quadratic Models 23
2.6 Algorithms for the Line Search 25
Questions for Chapter 2 29

Chapter 3 Newton-like Methods 33
3.1 Newton's Method 33
3.2 Quasi-Newton Methods 38
3.3 Invariance and Metrics 45
3.4 The Broyden Family 48
3.5 Numerical Experiments 54
3.6 Other Formulae 58
Questions for Chapter 3 60

Chapter 4 Conjugate Direction Methods 63
4.1 Conjugate Gradient Methods 63
4.2 Direction Set Methods 70
Questions for Chapter 4 75

Chapter 5 Restricted Step Methods 77
5.1 A Model Algorithm 77
5.2 Levenberg–Marquardt Methods 82
Questions for Chapter 5 88

Chapter 6 Sums of Squares and Non-linear Equations 91
6.1 Over-determined Systems 91
6.2 Well-determined Systems of Equations 98
6.3 No-derivative Methods 106
Questions for Chapter 6 110

References 113

Subject Index 118

Chapter 1

Introduction

1.1 Histcry and Applications

Optimization might be defined as the science of determining the 'best' solutions to certain mathematically defined problems, which are often models of physical reality. It involves the study of optimality criteria for problems, the determination of algorithmic methods of solution, the study of the structure of such methods, and computer experimentation with methods both under trial conditions and on real life problems. There is an extremely diverse range of practical applications. Yet the subject can be studied (not here) as a branch of pure mathematics.

Before 1940 relatively little was known about methods for numerical optimization of functions of many variables. There had been some least squares calculations carried out, and steepest descent type methods had been applied in some physics problems. The Newton method in many variables was known, and more sophisticated methods were being attempted such as the self-consistent field for variational problems in theoretical chemistry. Nonetheless anything of any complexity demanded armies of assistants operating desk calculating machines. There is no doubt therefore that the advent of the computer was paramount in the development of optimization methods and indeed in the whole of numerical analysis. The 1940s and 1950s saw the introduction and development of the very important branch of the subject known as linear programming. (The term 'programming' by the way is synonymous with 'optimization' and was originally used to mean optimization in the sense of optimal planning.) All these methods however had a fairly restricted range of application, and again in the postwar period the development of 'hill-climbing' methods took place – methods of wide applicability which did not rely on any special structure in the problem. The latter methods were at first very crude and inefficient, but the subject was again revolutionized in 1959 with the publication of a report by W. C. Davidon which lead to the introduction of variable metric methods. My friend and colleague M. J. D. Powell describes a meeting he attended in 1961 in which the speakers were telling of the difficulty of minimizing functions of ten variables, whereas he had just programmed a method based on Davidon's ideas which had solved problems of 100 variables in a short time. Since that time the development of the subject has proceeded apace and has included methods for a wide variety of

problems. This book describes these developments in what is hoped will be a systematic and comprehensive way.

The applicability of optimization methods is widespread, reaching into almost every activity in which numerical information is processed (Science, Engineering, Mathematics, Economics, Commerce, etc.). To provide a comprehensive account of all these applications would therefore be unrealistic, but a selection might include:

(a) chemical reactor design;
(b) aero-engine or aero-frame design;
(c) structural design – buildings, bridges, etc.;
(d) commerce – resource allocation, scheduling, blending;

and applications to other branches of numerical analysis:

(e) data fitting;
(f) variational principles in p.d.e.s;
(g) nonlinear equations in o.d.e.s; and
(h) penalty functions.

More such applications can be found in the proceedings of a conference on 'Optimization in Action' (Dixon, 1976), and many more of course in the specialized technical literature. However to give some idea of what is involved consider the optimum design of a distillation column, which can be modelled in an idealized way as in Figure 1.1.1. The aim of such a column is to separate out a more voltatile component from a mixture of components in the input stream. An *objective function* to be optimized might therefore be the quantity of the product or the profit from operating the system. The variables would be the rate of flow in the input, the heat rates applied and on each plate the liquid and vapour compositions of each component, and the temperature and vapour pressure. The variables are

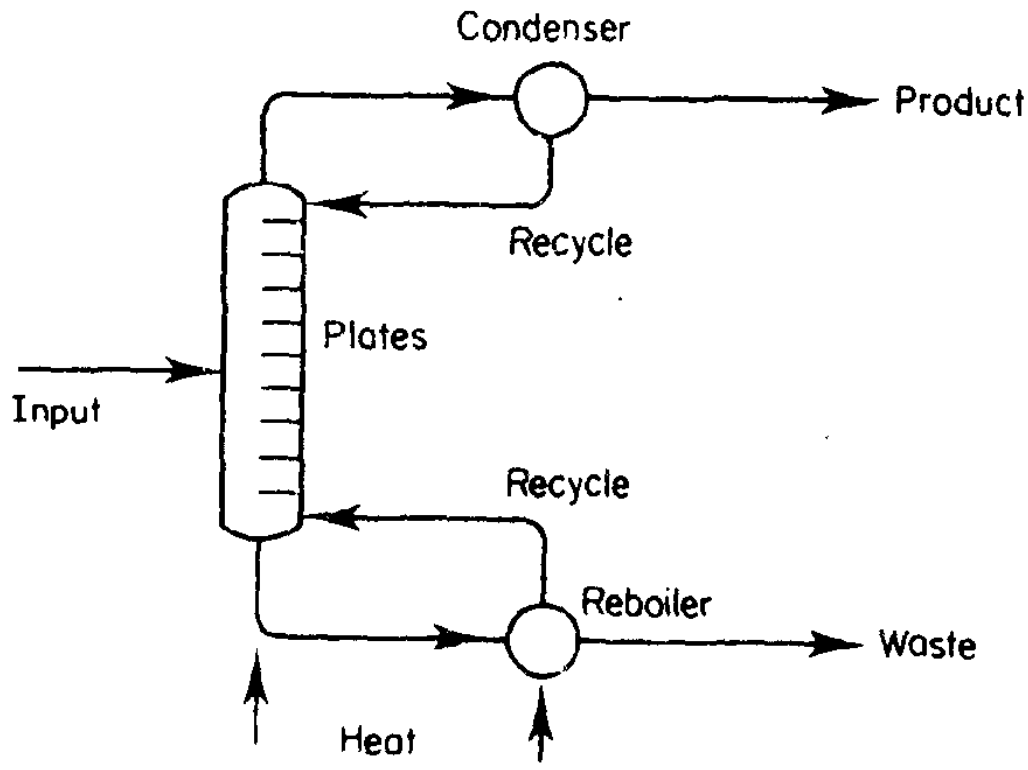

Figure 1.1.1 A model distillation column

subject to restrictions or inter-relations of many kinds, which are referred to as *constraints*. For instance compositions and flows must be non-negative ($x_i \geqslant 0$) and temperatures must not exceed certain upper bounds ($T_i \leqslant T_{max}$). Relationships such as the unit sum of percentage compositions must be included explicitly ($\Sigma_i x_i = 1$). More complicated constraints state how components interact physically, for instance vapour and liquid compositions are related by $v_i = l_i \phi(T_i)$, where $\phi(T_i)$ is a given but highly nonlinear function of temperature. A more difficult situation arises if the number of plates in the column is allowed to vary, and this is an example of an *integer variable* which can take on only integer values.

This book however is not concerned with applications, except insofar as they indicate the different types of optimization problem which arise. It is possible to categorize these into a relatively small number of standard problems and to state algorithms for each one. The user's task is to discover into what category his problem fits, and then to call up the appropriate optimization subroutine on a computer. This subroutine will specify to the user how the problem data is to be presented, for example non-linear functions usually have to be programmed in a user-written subroutine in a certain standard format. It is also as well to remember that in practice the solution of an optimization problem is not the only information that the user might need. He will often be interested in the *sensitivity* of the solution to changes in the parameters, especially so if the mathematical model is not a close approximation to reality, or if he cannot build his design to the same accuracy as the solution. He may indeed be interested in the variation of the solution obtained by varying some parameters over wide ranges, and it is often possible to provide this information without re-solving the problem numerous times.

This book therefore is concerned with some of the various standard optimization problems which can arise. In fact the material is divided into Volume I and Volume II. This volume, Volume I, is devoted to the subject of *unconstrained optimization*, in which the optimum value is sought of an objective function of many variables, without any constraints. This problem is important in its own right and also as a major tool in solving some constrained problems. Also many of the ideas carry over into constrained optimization. The special case of sums of squares functions, which arise in data fitting problems, is also considered. This also includes the solution of sets of simultaneous nonlinear equations, which is an important problem in its own right, but which is often solved by optimization methods. Volume II is devoted to *constrained optimization* in which the additional complication arises of the various types of constraint referred to above. It is planned to include material on

linear programming,
integer programming,
first and second order conditions,
convexity and duality,
quadratic programming,
linearly constrained programming,
nonlinear programming and penalty functions,
geometric programming, and
nondifferentiable optimization.

In these volumes a selection has had to be made amongst the extensive literature about optimization methods. I have been concerned to present *practical* methods (and associated theory) which have been implemented and for which a body of satisfactory numerical experience exists. I am equally concerned about reliability of algorithms and whether there is proof or good reason to think that convergence to a solution will occur at a reasonably rapid rate. However I shall also be trying to point out which new ideas in the subject I feel are significant and which might lead to future developments. Many people may read this book seeking a particular algorithm which best solves their specific problem. Such advice is not easy to give, especially in that the decision is not as clear-cut as it may seem. There are many special cases which should be taken into account, for instance the relative ease of computing the function and its derivatives. Similarly, considerations of how best to pose the problem in the first instance are relevant to the choice of method. Finally, and of most importance, the decision is subject to the availability of computer subroutines or packages which implement the methods. However some program libraries now give a decision tree in the documentation to help the user choose his method. Whilst these are valuable, they should only be used as a rough guide, and never as a substitute for common sense or the advice of a specialist in optimization techniques.

1.2 Mathematical Background

The book relies heavily on the concepts and techniques of matrix algebra and numerical linear algebra, which are not set out here (see Broyden, 1975, for example), although brief explanations are given in passing in certain cases. A *vector* is represented by a lower case bold letter (e.g. **a**) and refers to a column vector. A *matrix* is referred to by a bold upper case letter (e.g. **B**). That is

$$\mathbf{a} = \begin{pmatrix} a_1 \\ a_2 \\ . \\ . \\ . \\ a_n \end{pmatrix}; \qquad \mathbf{B} = \begin{bmatrix} B_{11} & B_{12} & \dots & B_{1s} \\ B_{21} & B_{22} & \dots & B_{2s} \\ . & . & & . \\ . & . & & . \\ . & . & & . \\ B_{r1} & B_{r2} & \dots & B_{rs} \end{bmatrix}.$$

Transposition is referred to by superscript T so that $\mathbf{a}^T$ is a row vector and $\mathbf{a}^T\mathbf{z}$ for instance is the *scalar product* $\mathbf{a}^T\mathbf{z} = \mathbf{z}^T\mathbf{a} = \sum_i a_i z_i$.

The ideas of vector spaces are also used, although often only in a simple minded way. A *point* x in n-dimensional space ($\mathbb{R}^n$) is the vector $(x_1, x_2, \dots, x_n)^T$, where x_1 is the component in the first coordinate direction, and so on. Most of the methods to be described are *iterative methods* which generate a *sequence* of points, $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}, \dots$ say, or $\{\mathbf{x}^{(k)}\}$ (the superscripts denoting iteration number), hopefully converging to a fixed point $\mathbf{x}^*$ which is the solution of the problem (see Figure 1.2.2). The idea of a *line* is important, and is the set of points

$$\mathbf{x}(=\mathbf{x}(\alpha)) = \mathbf{x}' + \alpha\mathbf{s} \tag{1.2.1}$$

for all α (sometimes for all $\alpha \geq 0$; this is strictly a half-line), in which $\mathbf{x}'$ is a fixed

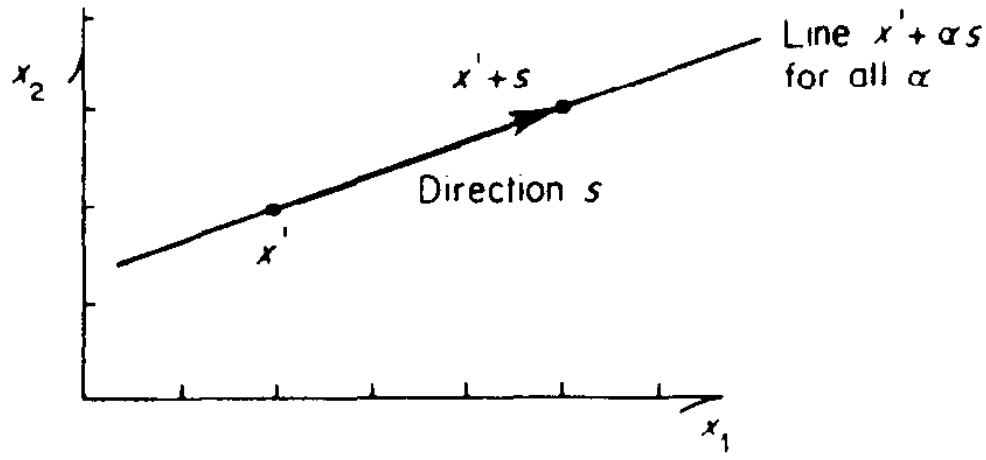

Figure 1.2.1 A line in two dimensions

point along the line (corresponding to $\alpha = 0$), and $\mathbf{s}$ is the *direction* of the line. For instance in Figure 1.2.1 $\mathbf{x}'$ is the point $\begin{pmatrix} 2 \\ 2 \end{pmatrix}$ and $\mathbf{s}$ the direction $\begin{pmatrix} 3 \\ 1 \end{pmatrix}$. The vector $\mathbf{s}$ is indicated by the arrow. Sometimes it is convenient to *normalize* $\mathbf{s}$ so that for instance $\mathbf{s}^T\mathbf{s} = \Sigma_i s_i^2 = 1$; this does not change the line, but only the value of α associated with any point.

The calculus of any *function* of many variables, $f(\mathbf{x})$ say, is clearly important. Some pictorial intuition for two variable problems is often gained by drawing *contours* (surfaces along which $f(\mathbf{x})$ is constant). A well-known test function for optimization methods is Rosenbrock's function

$$f(\mathbf{x}) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2 \tag{1.2.2}$$

the contours for which are shown in Figure 1.2.2. Some other contours are

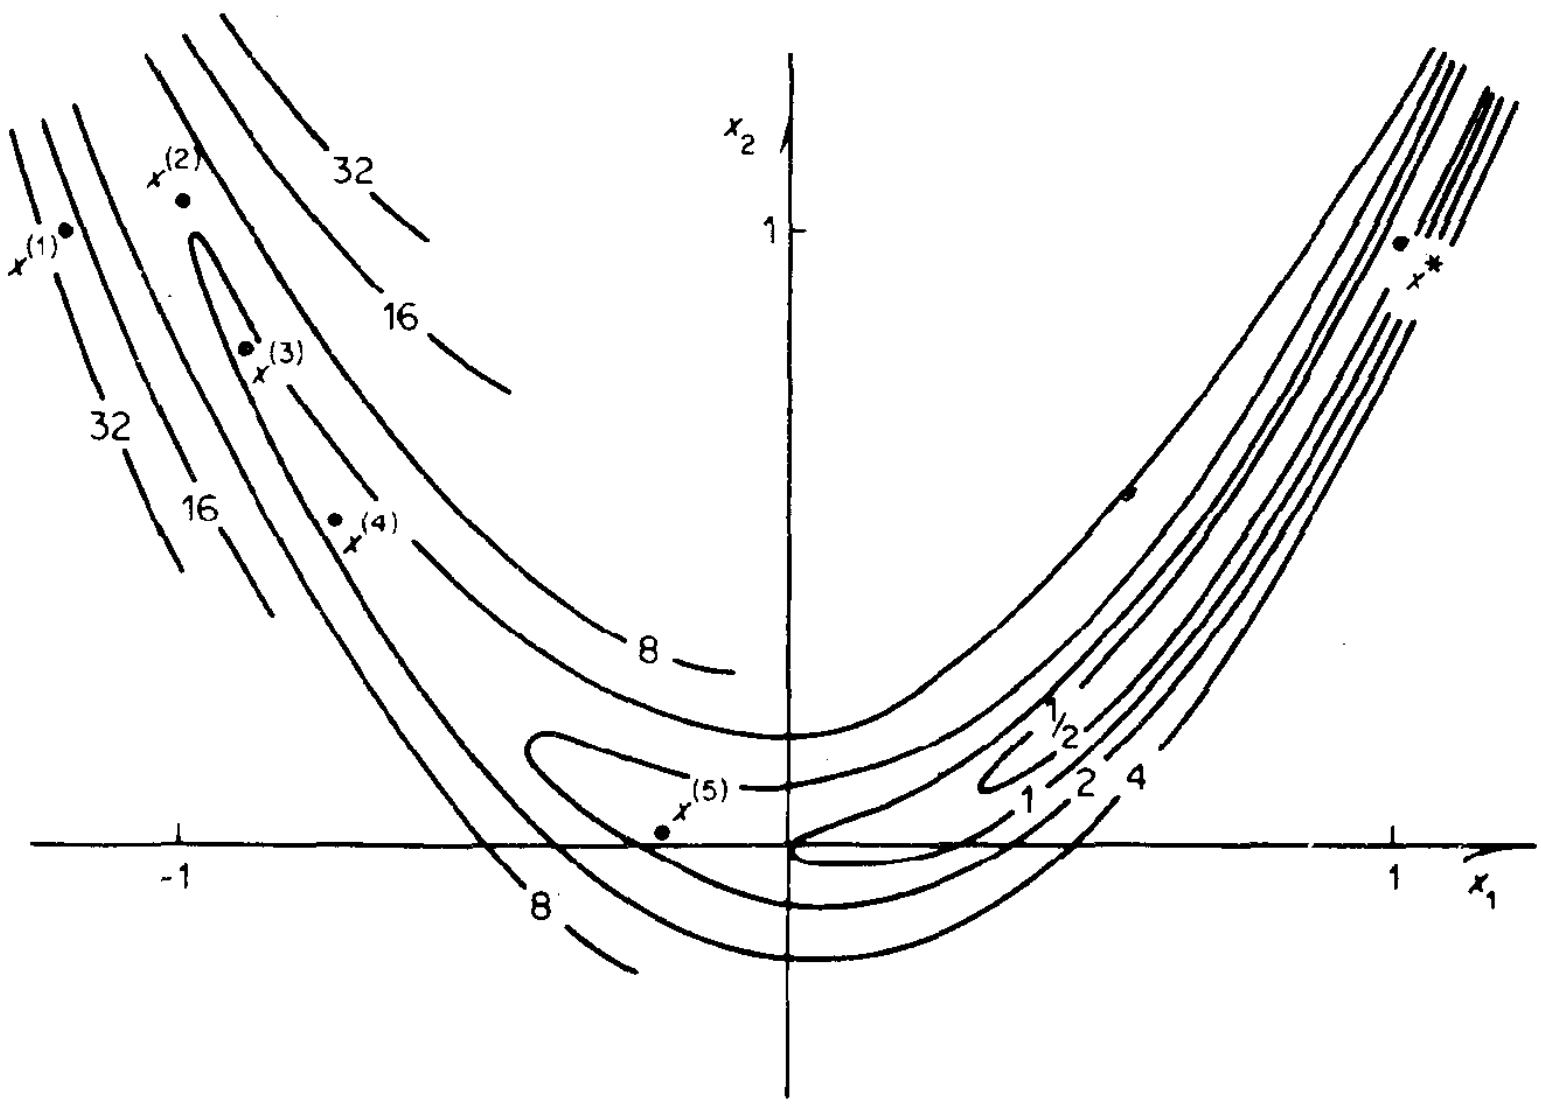

Figure 1.2.2 Contours for Rosenbrock's function, equation (1.2.2)

illustrated in Figure 6.2.2 in Chapter 6. In general it will be assumed that the problem functions which arise are *smooth*, that is continuous and continuously (Fréchet) differentiable ($\mathbb{C}^1$). Therefore for a function $f(\mathbf{x})$ at any point $\mathbf{x}$ there is a *vector of first partial derivatives*, or *gradient vector*

$$\begin{pmatrix} \partial f/\partial x_1 \\ \partial f/\partial x_2 \\ \vdots \\ \partial f/\partial x_n \end{pmatrix}_{\mathbf{x}} = \nabla f(\mathbf{x}) \tag{1.2.3}$$

where ∇ denotes the gradient operator $(\partial/\partial x_1, \ldots, \partial/\partial x_n)^{\mathrm{T}}$. If $f(\mathbf{x})$ is twice continuously differentiable ($\mathbb{C}^2$) then there exists a *matrix of second partial derivatives* or *Hessian matrix*, written $\nabla^2 f(\mathbf{x})$, for which the i,jth element is $\partial^2 f/(\partial x_i \partial x_j)$. This matrix is square and symmetric. Since any column (the jth, say) is $\nabla(\partial f/\partial x_j)$, the matrix can strictly be written as $\nabla(\nabla f^{\mathrm{T}})$. For example, in (1.2.2)

$$\nabla f(\mathbf{x}) = \begin{pmatrix} -400x_1(x_2 - x_1^2) - 2(1 - x_1) \\ 200(x_2 - x_1^2) \end{pmatrix}$$

$$\nabla^2 f(\mathbf{x}) = \begin{bmatrix} 1200x_1^2 - 400x_2 + 2 & -400x_1 \\ -400x_1 & 200 \end{bmatrix} \tag{1.2.4}$$

and this illustrates that ∇f and $\nabla^2 f$ will in general depend upon $\mathbf{x}$, and vary from point to point. Thus at $\mathbf{x}' = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, $\nabla f(\mathbf{x}') = \begin{pmatrix} -2 \\ 0 \end{pmatrix}$ and $\nabla^2 f(\mathbf{x}') = \begin{bmatrix} 2 & 0 \\ 0 & 200 \end{bmatrix}$ by substitution into (1.2.4).

These expressions can be used to determine the derivatives of f along any line $\mathbf{x}(\alpha)$ in (1.2.1). By the chain rule

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} = \sum_i \frac{\mathrm{d}}{\mathrm{d}\alpha} x_i(\alpha) \frac{\partial}{\partial x_i} = \sum_i s_i \frac{\partial}{\partial x_i} = \mathbf{s}^{\mathrm{T}} \nabla \tag{1.2.5}$$

so the *slope* of $f(=f(\mathbf{x}(\alpha))$ along the line at any point $\mathbf{x}(\alpha)$ is

$$\frac{\mathrm{d}f}{\mathrm{d}\alpha} = \mathbf{s}^{\mathrm{T}} \nabla f = \nabla f^{\mathrm{T}} \mathbf{s}. \tag{1.2.6}$$

Likewise the *curvature* along the line is

$$\frac{\mathrm{d}^2 f}{\mathrm{d}\alpha^2} = \frac{\mathrm{d}}{\mathrm{d}\alpha} \frac{\mathrm{d}f}{\mathrm{d}\alpha} = \mathbf{s}^{\mathrm{T}} \nabla (\nabla f^{\mathrm{T}} \mathbf{s}) = \mathbf{s}^{\mathrm{T}} \nabla^2 f \, \mathbf{s} \tag{1.2.7}$$

where ∇f and $\nabla^2 f$ are evaluated at $\mathbf{x}(\alpha)$. Note that, writing $\mathbf{G} = \nabla^2 f$, then $\mathbf{Gs}$ is the vector for which $(\mathbf{Gs})_i = \sum_j G_{ij} s_j$, and $\mathbf{s}^{\mathrm{T}}\mathbf{Gs}$ is the scalar product of $\mathbf{s}$ and $\mathbf{Gs}$. For example, for (1.2.2) at $\mathbf{x}' = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, the slope along the line generated by $\mathbf{s} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ (the x_1-axis in Figure 1.2.2) is $\mathbf{s}^{\mathrm{T}} \nabla f = -2$ and the curvature is $\mathbf{s}^{\mathrm{T}}\mathbf{Gs} = 2$ $\left(\text{since } \mathbf{Gs} = \begin{pmatrix} 2 \\ 0 \end{pmatrix} \right)$.

These definitions of slope and curvature depend on the size of **s**, and this ambiguity can be resolved by requiring that $\| \mathbf{s} \| = 1$. (Note: the *norm* $\| \mathbf{s} \|$ is just a measure of the size of **s**; one common norm is the L_2 norm $\| \mathbf{s} \|_2 = \sqrt{(\mathbf{s}^T \mathbf{s})}$.) Denoting $\nabla f(\mathbf{x}')$ by $\mathbf{g}'$, then $\pm \mathbf{g}'/\| \mathbf{g}' \|_2$ are the directions of greatest and least slope, over all directions for which $\| \mathbf{s} \|_2 = 1$, and are orthogonal to the contour and tangent plane of $f(\mathbf{x})$ at $\mathbf{x}'$ (see Figure 1.2.3 and Question 1.4).

Special cases of many variable functions include the general *linear function* which can be written

$$f(\mathbf{x}) = \sum_{i=1}^{n} a_i x_i + b = \mathbf{a}^T \mathbf{x} + b \tag{1.2.8}$$

where **a** and b are constant. If the coordinate vector

$$\mathbf{e}_i = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix} \leftarrow i\text{th position} \tag{1.2.9}$$

is defined, then the identity $\nabla x_i = \mathbf{e}_i$ gives

$$\nabla \mathbf{x}^T = \nabla (x_1, x_2, \ldots, x_n) = [\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n] = \mathbf{I} \tag{1.2.10}$$

since the vectors $\mathbf{e}_i$ are the columns of the *unit matrix* **I**. Thus for (1.2.8), $\nabla f = \mathbf{a}$ is a constant vector, and $\nabla^2 f = \mathbf{0}$ is the *zero matrix*. A general *quadratic function* can be written

$$q(\mathbf{x}) = \tfrac{1}{2} \mathbf{x}^T \mathbf{G} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c \tag{1.2.11}$$

where **G**, **b**, and c are constant and **G** is symmetric, or as

$$q(\mathbf{x}) = \tfrac{1}{2} (\mathbf{x} - \mathbf{x}')^T \mathbf{G} (\mathbf{x} - \mathbf{x}') + c' \tag{1.2.12}$$

where $\mathbf{G}\mathbf{x}' = -\mathbf{b}$ and $c' = c - \tfrac{1}{2} \mathbf{x}'^T \mathbf{G} \mathbf{x}'$. From the rule for differentiating a product,

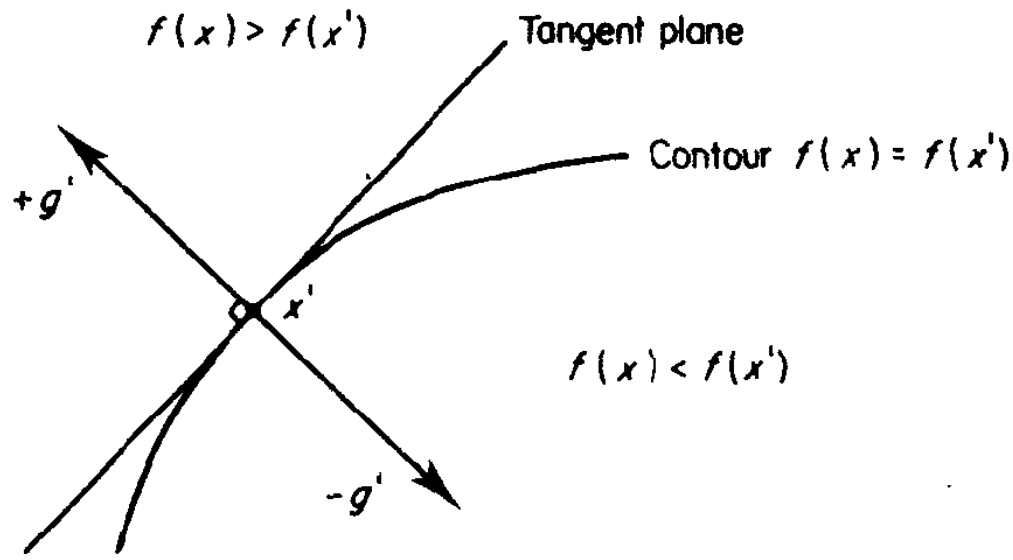

Figure 1.2.3 Properties of the gradient vector

it can be verified that

$$\nabla(\mathbf{u}^T\mathbf{v}) = (\nabla\mathbf{u}^T)\mathbf{v} + (\nabla\mathbf{v}^T)\mathbf{u} \tag{1.2.13}$$

if $\mathbf{u}$ and $\mathbf{v}$ depend upon $\mathbf{x}$. It therefore follows from (1.2.11) (using $\mathbf{u} = \mathbf{x}$, $\mathbf{v} = \mathbf{G}\mathbf{x}$) that

$$\nabla q(\mathbf{x}) = \tfrac{1}{2}(\mathbf{G} + \mathbf{G}^T)\mathbf{x} + \mathbf{b} = \mathbf{G}\mathbf{x} + \mathbf{b} \tag{1.2.14}$$

using the symmetry of $\mathbf{G}$. Likewise $\nabla^2 q = \mathbf{G}$ can be established. Thus $q(\mathbf{x})$ has a constant Hessian matrix $\mathbf{G}$ and its gradient is a linear function of $\mathbf{x}$. A consequence of (1.2.14) is that if $\mathbf{x}'$ and $\mathbf{x}''$ are two given points and if $\mathbf{g}' = \nabla q(\mathbf{x}')$ and $\mathbf{g}'' = \nabla q(\mathbf{x}'')$ then

$$\mathbf{g}'' - \mathbf{g}' = \mathbf{G}(\mathbf{x}'' - \mathbf{x}') \tag{1.2.15}$$

that is the Hessian matrix maps differences in position into differences in gradient. This result is used widely.

An indispensable technique for handling more general smooth functions of many variables is the *Taylor series*. For functions of one variable the infinite series is

$$f(\alpha) = f(0) + \alpha f'(0) + \tfrac{1}{2}\alpha^2 f''(0) + \ldots \tag{1.2.16}$$

although the series may be truncated after the term in α^p, replacing $f^{(p)}(0)$ by $f^{(p)}(\xi)$ where $\xi \in [0, \alpha]$ (the closed interval $0 \leqslant x \leqslant \alpha$). An integral form of the remainder can also be used. Now let $f(\alpha) = f(\mathbf{x}(\alpha))$ be the value of a function of many variables along the line $\mathbf{x}(\alpha)$ (see (1.2.1)). Then using (1.2.6) and (1.2.7) in (1.2.16)

$$f(\mathbf{x}' + \alpha\mathbf{s}) = f(\mathbf{x}') + \alpha\mathbf{s}^T\nabla f(\mathbf{x}') + \tfrac{1}{2}\alpha^2\mathbf{s}^T[\nabla^2 f(\mathbf{x}')]\mathbf{s} + \ldots \tag{1.2.17}$$

or by writing $\mathbf{h} = \alpha\mathbf{s}$

$$f(\mathbf{x}' + \mathbf{h}) = f(\mathbf{x}') + \mathbf{h}^T\nabla f(\mathbf{x}') + \tfrac{1}{2}\mathbf{h}^T[\nabla^2 f(\mathbf{x}')]\mathbf{h} + \ldots. \tag{1.2.18}$$

These are two forms of the many variable Taylor series. Furthermore, consider applying (1.2.18) to the function $\partial f(\mathbf{x})/\partial x_i$. Since $\nabla(\partial f(\mathbf{x})/\partial x_i)$ is the ith column of the Hessian matrix $\nabla^2 f$, it follows that

$$\nabla f(\mathbf{x}' + \mathbf{h}) = \nabla f(\mathbf{x}') + [\nabla^2 f(\mathbf{x}')]\mathbf{h} + \ldots \tag{1.2.19}$$

which is a Taylor series expansion for the gradient of f. Neglecting the higher terms in the limit $\mathbf{h} \to \mathbf{0}$, then this reduces to (1.2.15) showing that a general function behaves like a quadratic function in a sufficiently small neighbourhood of $\mathbf{x}'$.

It is hoped that a grasp of simple mathematical concepts such as these will enable the reader to follow most of the developments in the book. In certain places more complicated mathematics is used without detailed explanation. This is usually in an attempt to establish important results rigorously; however they often can be skipped over without losing the thread of the explanation.

Questions for Chapter 1

1. Obtain expressions for the gradient vector and Hessian matrix for the functions of n variables:
   (i) $\mathbf{a}^T\mathbf{x}$: $\mathbf{a}$ constant;
   (ii) $\mathbf{x}^T\mathbf{A}\mathbf{x}$: $\mathbf{A}$ unsymmetric and constant;
   (iii) $\frac{1}{2}\mathbf{x}^T\mathbf{A}\mathbf{x} + \mathbf{b}^T\mathbf{x}$: $\mathbf{A}$ symmetric, $\mathbf{A}$, $\mathbf{b}$ constant;
   (iv) $\mathbf{f}^T\mathbf{f}$: $\mathbf{f}$ is an m-vector depending on $\mathbf{x}$ and $\nabla \mathbf{f}^T$ is denoted by $\mathbf{A}$ which is not constant.
2. Write down the Taylor expansion for the gradient $\mathbf{g}(\mathbf{x}' + \boldsymbol{\delta})$ about $\mathbf{x}'$, neglecting terms of order $\| \boldsymbol{\delta} \|^2$. Hence show that if $f(\mathbf{x})$ is a quadratic function with Hessian $\mathbf{G}$, then $\boldsymbol{\gamma} = \mathbf{G}\boldsymbol{\delta}$, where $\boldsymbol{\delta}$ is the difference between any two points and $\boldsymbol{\gamma}$ is the corresponding difference in gradients.
3. Write down the Taylor expansion for the m-vector $\mathbf{f}(\mathbf{x})$ about $\mathbf{x}'$, where $\nabla \mathbf{f}^T$ is denoted by $\mathbf{A}$.
4. At a point $\mathbf{x}'$ for which $\mathbf{g}' \neq \mathbf{0}$, show that the direction vector $\mathbf{s} = \mathbf{g}'/\|\mathbf{g}'\|_2$ has the greatest slope, over all vectors for which $\mathbf{s}^T\mathbf{s} = 1$. (The *steepest ascent* vector.)
5. At a point $\mathbf{x}'$ for which $\mathbf{g}' \neq \mathbf{0}$, show that the direction vectors $\pm\mathbf{g}'$ are orthogonal to the contour and the tangent plane surface at $\mathbf{x}'$.
6. If $\mathbf{x}(\alpha)$ is any twice differentiable arc, if $f(\mathbf{x}(\alpha))$ is regarded as $f(\alpha)$, and if $d\mathbf{x}(\alpha_0)/d\alpha = \mathbf{s}$ and $d^2\mathbf{x}(\alpha_0)/d\alpha^2 = \mathbf{t}$, use the chain rule to obtain expressions for $df(\alpha_0)/d\alpha$ and $d^2f(\alpha_0)/d\alpha^2$ in terms of $\mathbf{s}$, $\mathbf{t}$ and the derivatives of $f(\mathbf{x})$ evaluated at $\mathbf{x}(\alpha_0)$.

(Some other questions which partly refer to the material of Section 1.2 are given at the end of Chapter 2.)

Chapter 2

Structure of Methods

2.1 Conditions for Local Minima

In the following chapters the problem of finding a local solution to the problem

$$\text{minimize } f(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^n \tag{2.1.1}$$

is considered. The minimizing point or *minimizer* is referred to as $\mathbf{x}^*$. Note first of all that it is only generally practicable to search for *local* solutions rather than *global* solutions (see Figure 2.1.1 for example and also Figure 6.2.2 in Chapter 6). The possibility of finding global solutions has been considered; see Dixon and Szego (1975), for example, and also Section 6.2. However there are considerable difficulties.

Other difficulties are caused when *non-smooth* minima exist (see Figure 2.1.1) and it is convenient that this case is excluded by the differentiability assumptions of Section 1.2. However some problems of this type do arise in practice and it is hoped to mention them in Volume II. It is not difficult however to extend (2.1.1) to solve *maximization problems* through the simple transformation

$$\max_{\mathbf{x}} f(\mathbf{x}) = -\min_{\mathbf{x}} -f(\mathbf{x}). \tag{2.1.2}$$

The existence of first and usually second derivatives of f will be assumed and to simplify the notation

$$\mathbf{g}(\mathbf{x}) = \nabla f(\mathbf{x}); \qquad \mathbf{G}(\mathbf{x}) = \nabla^2 f(\mathbf{x})$$

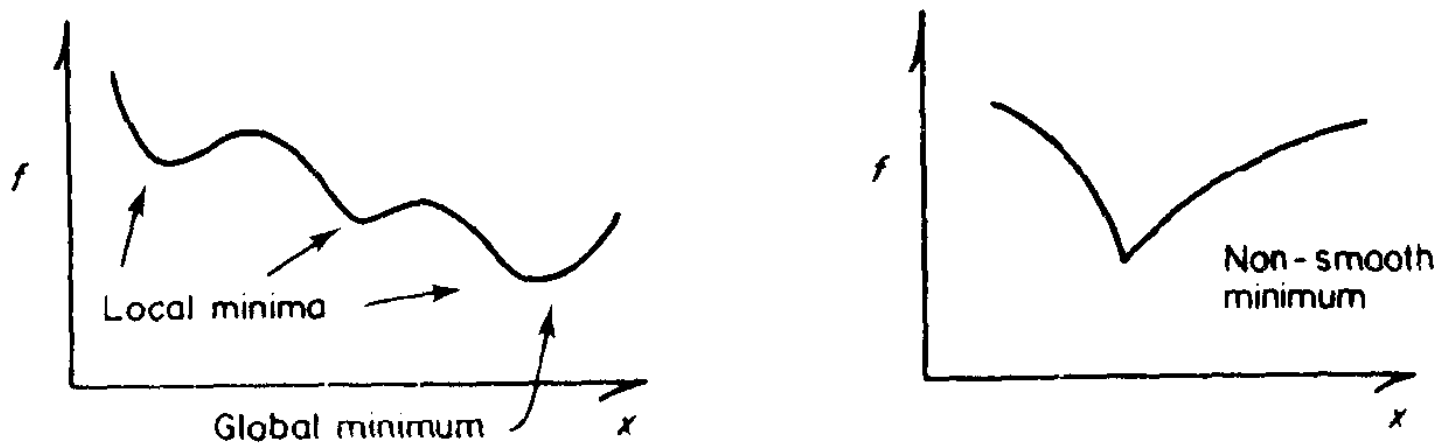

Figure 2.1.1 Types of minima

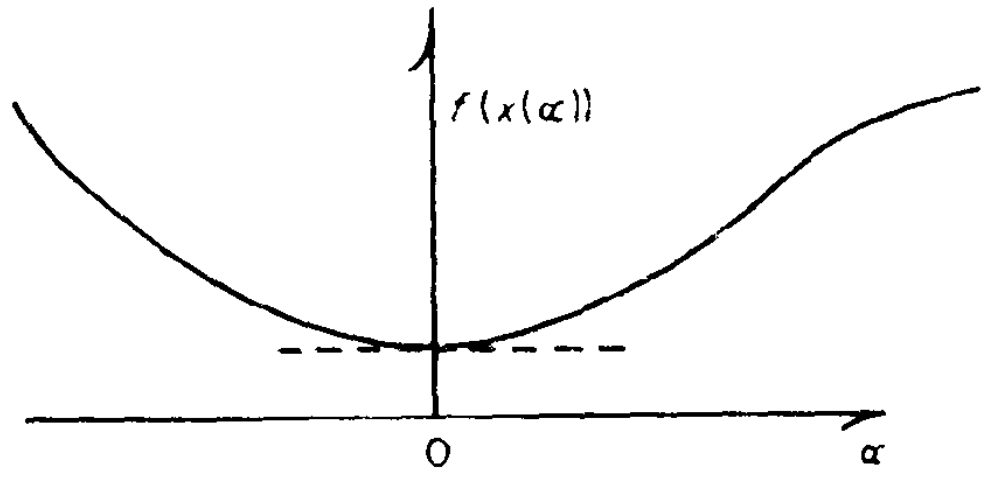

Figure 2.1.2 Zero slope and positive curvature at $\alpha = 0$

is written. Also $f^* = f(\mathbf{x}^*)$, $\mathbf{g}^* = \mathbf{g}(\mathbf{x}^*)$, etc, is used for quantities derived from $\mathbf{x}^*$, likewise $f^{(k)} = f(\mathbf{x}^{(k)})$, $\mathbf{g}^{(k)} = \mathbf{g}(\mathbf{x}^{(k)})$ and so on.

The main aim of this section is to state and discuss some simple conditions which hold at a local minimizer $\mathbf{x}^*$. These arise from the observation that along any line $\mathbf{x}(\alpha) = \mathbf{x}^* + \alpha\mathbf{s}$ through $\mathbf{x}^*$, then $f(= f(\mathbf{x}(\alpha)))$ has both *zero slope* and *positive curvature* at $\mathbf{x}^*$. This is illustrated in Figure 2.1.2 and is the usual condition for a local minimum of a function of one variable. From (1.2.6) and (1.2.7) it follows for all $\mathbf{s}$ that both $\mathbf{s}^T\mathbf{g}^* = 0$ and $\mathbf{s}^T\mathbf{G}^*\mathbf{s} \geqslant 0$. Since these conditions are implied by $\mathbf{x}^*$ being a local minimizer, they are *necessary conditions* for a local solution and can be stated equivalently as

$$\mathbf{g}^* = \mathbf{0} \tag{2.1.3}$$

and

$$\mathbf{s}^T\mathbf{G}^*\mathbf{s} \geqslant 0, \qquad \forall \mathbf{s} \tag{2.1.4}$$

($\forall$ means 'for all'). Condition (2.1.3) follows because $\mathbf{s}^T\mathbf{g}^* = 0$ for all $\mathbf{s}$ can be true if and only if $\mathbf{g}^* = \mathbf{0}$, and is referred to as a *first order* necessary condition, since it only involves first derivatives. Condition (2.1.4) is a *second order* necessary condition, and is the condition that $\mathbf{G}^*$ is a *positive semi-definite matrix*, by definition of this property.

It is also possible to derive *sufficient conditions* (those which imply that $\mathbf{x}^*$ is a local minimizer). Here the result is stated in the form of a theorem.

Theorem 2.1.1

Sufficient conditions for an isolated local minimizer $\mathbf{x}^*$ *are that (2.1.3) holds and that* $\mathbf{G}^*$ *is positive definite, that is*

$$\mathbf{s}^T\mathbf{G}^*\mathbf{s} > 0 \qquad \forall \mathbf{s} \neq \mathbf{0}. \tag{2.1.5}$$

Proof

Consider any point $\mathbf{x}^* + \boldsymbol{\delta}$, $\boldsymbol{\delta} \neq \mathbf{0}$. A Taylor series about $\mathbf{x}^*$ and (2.1.3) imply that

$$f(\mathbf{x}^* + \boldsymbol{\delta}) = f^* + \tfrac{1}{2}\boldsymbol{\delta}^T\mathbf{G}^*\boldsymbol{\delta} + o(\boldsymbol{\delta}^T\boldsymbol{\delta})$$

(using $o(\cdot)$ and $O(\cdot)$ strictly in the sense of Hardy, 1960, although a rough interpretation is that $o(\cdot)$ means 'negligible relative to $(\cdot)$' and $O(\cdot)$ means 'of order $(\cdot)$'). Now (2.1.5) implies that there exists an $a > 0$ such that $\boldsymbol{\delta}^T\mathbf{G}^*\boldsymbol{\delta} \geqslant a\boldsymbol{\delta}^T\boldsymbol{\delta}$ (a is the smallest eigenvalue of $\mathbf{G}^*$) and hence that

$$f(\mathbf{x}^* + \boldsymbol{\delta}) \geqslant f^* + (\tfrac{1}{2}a + o(1))\boldsymbol{\delta}^T\boldsymbol{\delta}.$$

As $\boldsymbol{\delta} \to \mathbf{0}$, $o(1) \to 0$ and $a > 0$ is fixed, so it follows that $f(\mathbf{x}^* + \boldsymbol{\delta}) > f^*$, and hence $\mathbf{x}^*$ is an isolated local minimizer. □

These sufficient conditions are convenient in that they are readily checked numerically. For instance, if $f(\mathbf{x})$ is given by (1.2.2), then at $\mathbf{x}^* = (1, 1)^T$, $\mathbf{g}^* = (0, 0)^T$ and $\mathbf{G}^* = \begin{bmatrix} 802 & 400 \\ 400 & 200 \end{bmatrix}$ which is positive definite (see below), so that it follows from theorem 2.1.1 that $\mathbf{x}^*$ is an isolated local minimizer. (In fact since $f^* = 0$ and $f(\mathbf{x}) \geqslant 0$ it is clear that $\mathbf{x}^*$ is also a global minimizer.) The necessary conditions (2.1.3) and (2.1.4), and the sufficient conditions of theorem 2.1.1 are almost necessary and sufficient, and there is only a 'gap' in the case of zero curvature. Examples which satisfy the necessary but not the sufficient conditions are $f(x) = x^3$ and $f(x) = x^4$. $x^* = 0$ is a local minimizer of the second function but not the first.

The notion of a positive definite matrix $\mathbf{G}^*$ may be unfamiliar to some readers, and the definition (2.1.5) does not help in that it cannot be checked numerically. However there are several different equivalent definitions which can be checked, namely

(i) all eigenvalues of $\mathbf{G}^* > 0$,

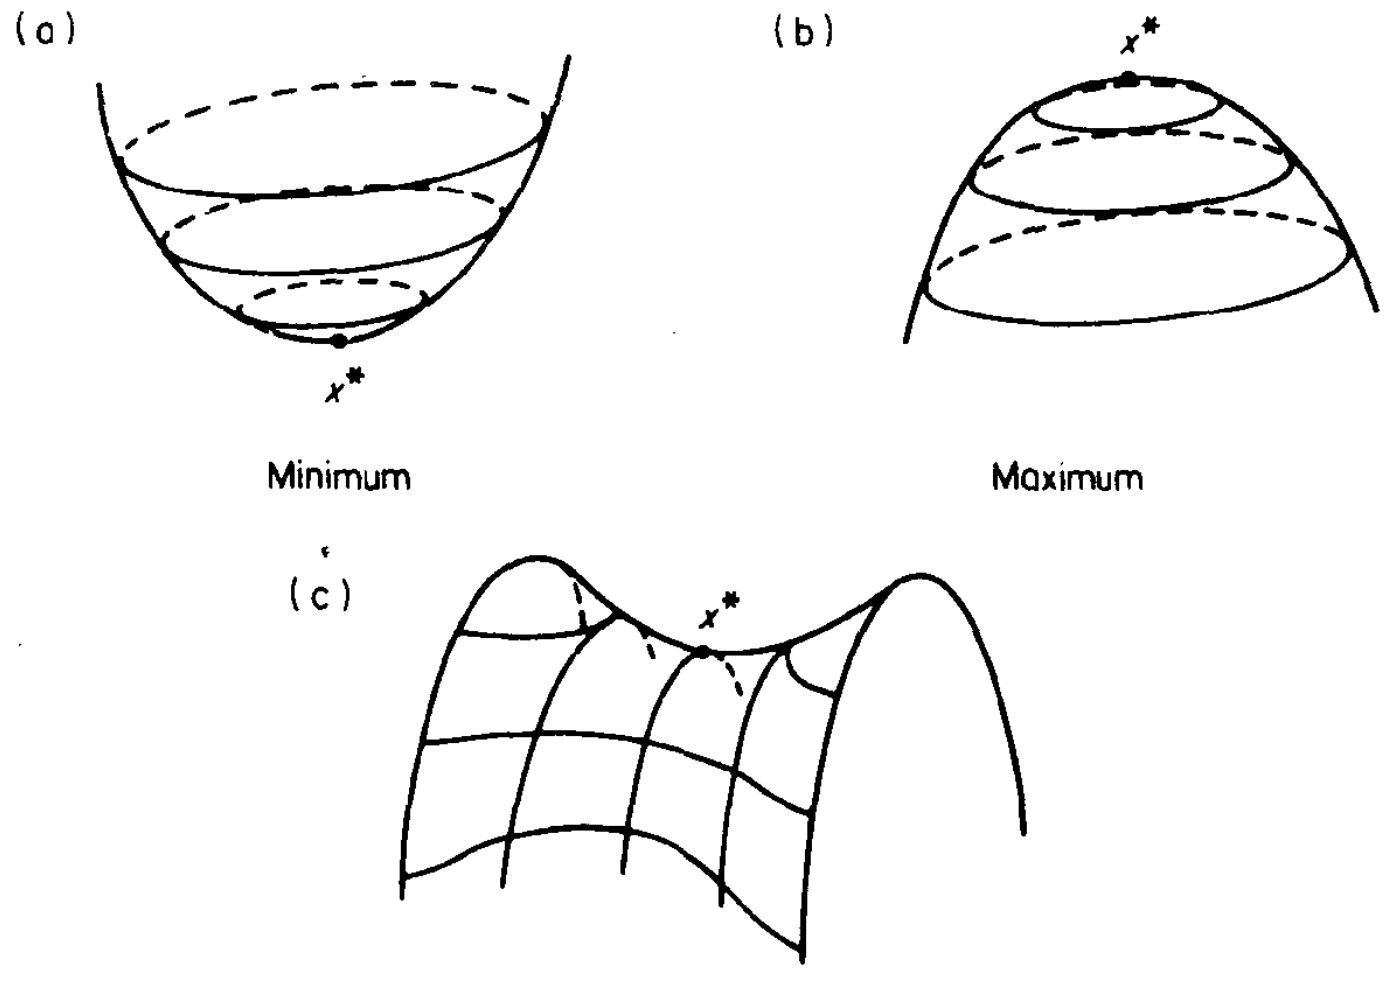

Figure 2.1.3 Types of stationary point

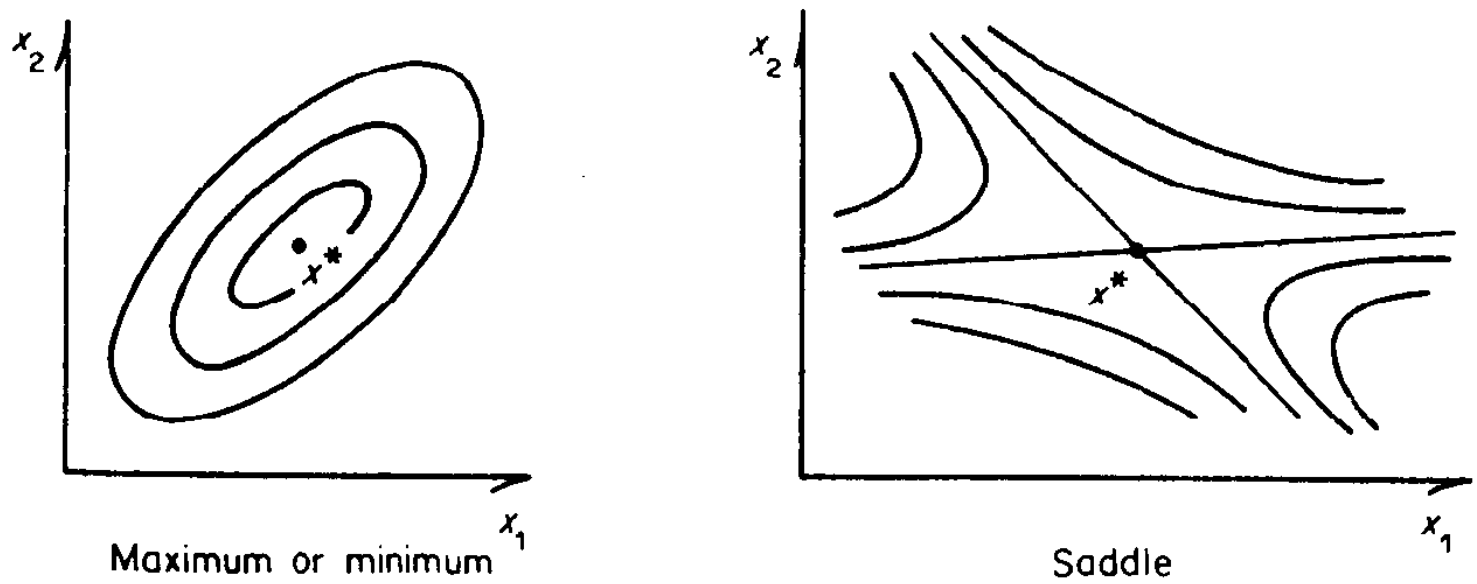

Figure 2.1.4 Contours for stationary points

(ii) $\mathbf{LL}^T$ (Choleski) factors of $\mathbf{G}^*$ exist with $l_{ii} > 0$,
(iii) $\mathbf{LDL}^T$ factors exist with $l_{ii} = 1$ and $d_{ii} > 0$,
(iv) all pivots > 0 in Gaussian elimination without pivoting, and
(v) all principal minors of $\mathbf{G}^* > 0$.

The matrices **L** and **D** are lower triangular and diagonal, respectively. For small n ($\leqslant 3$), condition (v) is most readily checked (for $\mathbf{G}^*$ above we have $\det(802) > 0$ and $\det\begin{bmatrix} 802 & -400 \\ -400 & 200 \end{bmatrix} > 0$), but in general conditions (ii) or (iii) are the most efficient and they also enable linear equations with coefficient matrix **G** to be solved subsequently (see Section 3.1).

In fact many minimization methods are based only upon trying to locate a point $\mathbf{x}^*$ such that $\mathbf{g}(\mathbf{x}^*) = \mathbf{0}$. This may not be a local minimizer and in general is referred to as a *stationary point*. Different types of stationary point are illustrated in Figure 2.1.3 and their contours in Figure 2.1.4. Note that in Figure 2.1.3, whilst all the graphs have zero slope at $\mathbf{x}^*$, for (a) there is positive curvature in every direction, for (b) negative curvature in every direction, whereas for (c) there is negative curvature across the saddle and positive curvature along the saddle. Thus usually a minimizer corresponds to a positive definite Hessian matrix, a maximizer to a negative definite matrix, and a saddle point to an indefinite matrix (that is one in which the eigenvalues have both positive and negative sign). Numerical methods for finding stationary points which are not minimizers are occasionally of interest (see Sinclair and Fletcher, 1974) and some possibilities are described in Question 4.5.

2.2 Ad hoc Methods

Many early methods which were suggested for minimization were based on rough and ready ideas without very much theoretical background. It is instructive for the reader to think about how he or she would go about the problem, given that values of $f(\mathbf{x})$ only can be evaluated for any $\mathbf{x}$. If the problem is in only two or three variables then it is likely that some sort of repeated bisection in each one of the variables might be tried so as to establish a region in which the minimum