

Numerical Methods for Unconstrained Optimization

an introduction

M.A. Wolfe

*Department of Applied Mathematics
University of St. Andrews*

© M.A. Wolfe, 1978

ISBN 0 442 30214 2 cloth
0 442 30217 7 paper

All rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means – graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage or retrieval systems – without written permission of the publishers

**Published by Van Nostrand Reinhold Company Ltd.,
Molly Millars Lane, Wokingham, Berkshire, England**

*Published in 1978 by Van Nostrand Reinhold Company
A Division of Litton Educational Publishing, Inc.,
450 West 33rd Street, New York, N.Y. 10001, U.S.A.*

*Van Nostrand Reinhold Limited
1410 Birchmount Road, Scarborough, Ontario, M1P 2E7,
Canada*

*Van Nostrand Reinhold Australia Pty. Limited
17 Queen Street, Mitcham, Victoria 3132, Australia*

Library of Congress Cataloging in Publication Data

Wolfe, Michael Anthony.

Numerical methods for unconstrained optimization.

Bibliography: p.

Includes index.

1. Mathematical optimization. 2. Approximation theory. I. Title.

QA402.5.W64 519.4 77-15945

ISBN 0-442-30214-2

ISBN 0-442-30217-7 pbk.

Typeset by Alden Press Oxford, London & Northampton and
printed in Great Britain by Henry Ling Ltd., the Dorset Press,
Dorchester, Dorset.

Contents

Preface	v
Chapter 1 Fundamentals	1
1.1 Introduction	1
1.2 Fundamental Linear Algebra	6
1.3 Fundamental Analysis	12
1.4 The Classification of Minimizers	20
1.5 Sufficient Conditions for a Strong Local Minimizer	25
1.6 Convex Functions	29
1.7 Lagrange Multipliers	36
1.8 Convergence	40
Exercises	43
Chapter 2 Line Search Techniques	53
2.1 Introduction	53
2.2 Quadratic Interpolation	62
2.3 Cubic Interpolation	73
Exercises	78
Chapter 3 The Steepest Descent and Newton Methods	82
3.1 The Method of Steepest Descent	82
3.2 The Method of Newton	89
3.3 Safeguarding the Method of Newton	92
3.4 Murray's Modification of Newton's Method	99
3.5 Numerical Approximation of the Hessian	108
Exercises	110
Chapter 4 Conjugate Direction Methods	113
4.1 Introduction	113
4.2 Conjugate Directions	114
4.3 Conjugate Direction Methods	116
4.4 Conjugate Gradient Methods	120
Exercises	126
Chapter 5 Direct Search Methods	130
5.1 Introduction	130
5.2 The Method of Nelder and Mead	131
5.3 The Method of Powell	141
Exercises	157

Chapter 6 Quasi-Newton Methods	158
6.1 Introduction	158
6.2 The Method of Fletcher and Powell	161
6.3 The Method of Stewart	167
6.4 Quadratically Convergent Methods	176
6.5 A Rank One Updating Formula	193
6.6 Some Rank Two Updating Formulae	196
6.7 Gill and Murrys' Implementation of Quasi-Newton Methods	209
Exercises	215
Chapter 7 The Least Squares Problem	218
7.1 Introduction	218
7.2 The Linear Least Squares Problem	220
7.3 The Gauss-Newton Method	225
7.4 The Method of Hartley	230
7.5 The Methods of Levenberg and Marquardt	233
7.6 Fletcher's Method	239
7.7 The Method of Meyer and Roth	247
7.8 Approximation of Derivatives	261
Exercises	263
Appendices	267
A. Derivation of the Cubic Interpolation Formula	267
B. Objective Functions	270
C. The Cholesky Decomposition of a Symmetric Positive Definite Matrix	275
D. The Connection Between Orthogonal and Oblique Co-ordinates	278
E. A Proof of Hadamard's Inequality	280
F. The Sherman-Morrison-Woodbury Formula	283
G. Updating the Cholesky Decomposition of a Positive Definite Matrix	284
H. Householder Transformations	293
I. The Eigenvalues of the Sum of Two Real Symmetric Matrices	299
J. A Criterion for the Convergence of a Sequence in a Compact Set of $\{R^n, \ \cdot\ \}$	306
References	307
Index	311

Fundamentals

1.1 Introduction

This book is an introduction to the study of numerical methods for unconstrained optimization. In order to understand the nature of unconstrained optimization and to gain some idea of its importance in industrial and scientific contexts, we consider four problems. The first problem serves to distinguish constrained optimization from unconstrained optimization and the other three problems are intended to give some idea, however slight, of the many uses to which unconstrained optimization can be put.

Problem 1: A Manufacturing Process

Suppose that in a factory a chemical is manufactured in bulk by using a given process. The cost of manufacturing unit weight of the chemical depends upon the amount of each constituent raw material required, and upon various parameters associated with the manufacturing process such as temperatures and pressures at various points in the plant. By empirical or theoretical means the cost of manufacturing unit weight of the chemical is expressible as a function f of a number of parameters x_1, \dots, x_n which are measures of weights of raw materials, temperatures, pressures, and the like. In general, the ranges of values of the parameters x_1, \dots, x_n are dictated by engineering and economic considerations and cannot be arbitrary; for example, the weights of raw materials required cannot be negative. We say that *constraints* are imposed upon the values of x_1, \dots, x_n . It is clearly of interest to the manufacturer of the chemical to find values of x_1, \dots, x_n , within the limitations imposed by the constraints, for which the cost $f(x_1, \dots, x_n)$ of manufacturing unit weight of the chemical has the smallest possible value. He has therefore to solve the following optimization problem.

Find values x_1^*, \dots, x_n^* of x_1, \dots, x_n , which satisfy a given set of constraints, such that $f(x_1^*, \dots, x_n^*) \leq f(x_1, \dots, x_n)$ for all values of x_1, \dots, x_n which satisfy the constraints; that is, *minimize* f subject to the given set of constraints.

The function f to be minimized is called the *objective* function. Typical constraints are of the form

$$a_i \leq x_i \leq b_i \quad (i = 1, \dots, n) \quad (1.1.1)$$

where the a_i and the b_i are given real numbers.

More generally, a problem in constrained optimization is of the following form.

Problem 1

Minimize

$$f(x_1, \dots, x_n) \quad (1.1.2a)$$

Subject to

$$\phi_i(x_1, \dots, x_n) = 0 \quad (i = 1, \dots, l) \quad (1.1.2b)$$

$$\psi_j(x_1, \dots, x_n) \leq 0 \quad (j = 1, \dots, m) \quad (1.1.2c)$$

$$(x_1, \dots, x_n) \in D \subset R^n \quad (1.1.2d)$$

in which ϕ_i and ψ_j are given real functions of n real variables, where $l \leq n$, R^n is the

set of all n -tuples of real numbers, and D is a subset of R^n , such as the subset defined by (1.1.1) \square

It is convenient to introduce vector notation in which \mathbf{x} , $\boldsymbol{\phi}(\mathbf{x})$, and $\boldsymbol{\psi}(\mathbf{x})$ are column vectors defined by

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \quad \boldsymbol{\phi}(\mathbf{x}) = \begin{bmatrix} \phi_1(\mathbf{x}) \\ \vdots \\ \phi_l(\mathbf{x}) \end{bmatrix} \quad \boldsymbol{\psi}(\mathbf{x}) = \begin{bmatrix} \psi_1(\mathbf{x}) \\ \vdots \\ \psi_m(\mathbf{x}) \end{bmatrix} \quad (1.1.3a)$$

or

$$\begin{aligned} \mathbf{x} &= [x_1, \dots, x_n]^T \quad \boldsymbol{\phi}(\mathbf{x}) = [\phi_1(\mathbf{x}), \dots, \phi_l(\mathbf{x})]^T \\ \boldsymbol{\psi}(\mathbf{x}) &= [\psi_1(\mathbf{x}), \dots, \psi_m(\mathbf{x})]^T \end{aligned} \quad (1.1.3b)$$

It is also convenient to introduce vector inequalities. Let

$$\mathbf{a} = [a_1, \dots, a_n]^T \quad \mathbf{b} = [b_1, \dots, b_n]^T$$

Then we write $\mathbf{a} \geq \mathbf{b}$ if and only if $a_i \geq b_i$ ($i = 1, \dots, n$). By using the vector notation we may express the general constrained optimization problem as follows.

Problem 1

Minimize

$$f(\mathbf{x}) \quad (1.1.4a)$$

Subject to

$$\boldsymbol{\phi}(\mathbf{x}) = \mathbf{0} \quad (1.1.4b)$$

$$\boldsymbol{\psi}(\mathbf{x}) \leq \mathbf{0} \quad (1.1.4c)$$

$$\mathbf{x} \in D \subset R^n \quad (1.1.4d)$$

\square

The constraints (1.1.2b) or (1.1.4b) are *equality* constraints, while the constraints (1.1.2c) or (1.1.4c) are *inequality* constraints, and (1.1.2d) or (1.1.4d) are *set* constraints. Any point \mathbf{x} which satisfies (1.1.2b), (1.1.2c), and (1.1.2d) or (1.1.4b), (1.1.4c), and (1.1.4d) is called a *feasible* point.

The problem (1.1.2) or (1.1.4) is usually referred to as the *general nonlinear programming problem*. An example is as follows.

Minimize

$$- [25 - (x_1 - 5)^2 - (x_2 - 5)^2]^{1/2}$$

Subject to

$$5x_2 - 4x_1 \leq 0$$

$$5x_2 + 4x_1 - 40 \leq 0$$

$$x_1 \geq 0, \quad x_2 \geq 0$$

This problem is trivial in that it can be solved by noticing that geometrically the surface

$$(x_1 - 5)^2 + (x_2 - 5)^2 + y^2 = 25$$

is a sphere with centre (5, 5, 0) and radius 5 and that the surfaces

$$5x_2 - 4x_1 = 0$$

$$5x_2 + 4x_1 - 40 = 0$$

are planes parallel to the y -axis which intersect the surface of the sphere and each other at the points $(5, 4, \sqrt{24})$ and $(5, 4, -\sqrt{24})$. Hence the minimum value of the objective function which satisfies the given constraints is $-\sqrt{24}$ and this occurs at $x_1^* = 5, x_2^* = 4$.

The nonlinear programming problems which are encountered in industrial, economic, administrative and scientific contexts are seldom as trivial as the preceding problem, and cannot usually be solved analytically or geometrically. A number of numerical methods exist, however, for solving nonlinear programming problems but a satisfactory account of them is beyond the scope of this book.

Problem 2: Fitting Data

Suppose that some variable y is thought to depend upon a variable t through a formula of the form

$$y = Y(t; \mathbf{x}) \quad (1.1.5)$$

where $Y: R^1 \times R^n \rightarrow R^1$ is a known function of t and of a vector \mathbf{x} of n parameters. (To denote that f is a real-valued function of n real variables x_1, \dots, x_n we write $f: R^n \rightarrow R^1$, and to denote that ϕ is a vector-valued function with l components, of the n real variables x_1, \dots, x_n we write $\phi: R^n \rightarrow R^l$. If F is a vector-valued function with n real components, of both a vector with l real components and a vector with m real components, we write $F: R^l \times R^m \rightarrow R^n$.) This situation arises frequently when a physically observable quantity y is predicted on theoretical grounds to depend upon a variable t in a certain way, but the theory, being incomplete, gives only the functional form of the dependence, which contains a number of parameters, the values of which are unknown. For example, if the length y of a rod of copper is regarded as a function of its temperature t , then a satisfactory formula for y is

$$y = Y(t; \mathbf{x}) = y_0(1 + x_1 t + x_2 t^2) \quad (1.1.6)$$

where y_0 , which is supposed to be known, is the length of the rod when $t = 0$ and x_1, x_2 , are parameters which are supposed to be unknown. If a number of measurements of y were made corresponding to an equal number of known values of t then the formula (1.1.6) could be 'fitted' to the experimental data consisting of the values of y and the corresponding values of t , and the values of x_1 and x_2 which ensure the 'best fit' to the data could be obtained.

In the general case corresponding to (1.1.5) suppose that for m values t_1, \dots, t_m of t , measurements y_1, \dots, y_m of y are made. In practice, even if the formula (1.1.5) exactly represents the physical situation, the measured values of y are usually subject to experimental errors, so that a value \mathbf{x}^* of \mathbf{x} such that $Y(t_i; \mathbf{x}^*)$ is exactly equal to the measured value y_i of y for all values of i ($i = 1, \dots, m$) is unlikely to exist. If, however, we take $m \geq n$ and let $f: R^n \rightarrow R^1$ be defined by

$$f(\mathbf{x}) = \sum_{i=1}^m [Y(t_i; \mathbf{x}) - y_i]^2 \quad (\mathbf{x} \in R^n) \quad (1.1.7)$$

then we should expect to be able to 'fit' (1.1.5) to the experimental data y_i ($i = 1, \dots, m$) by minimizing the objective function f . For obvious reasons, a value \mathbf{x}^* of \mathbf{x} obtained by minimizing f defined by (1.1.7) is said to give a *least squares* fit to the experimental data. The general nonlinear least squares problem is as follows.

Problem 2.

Minimize

$$f(\mathbf{x})$$

Subject to

$$\mathbf{x} \in R^n$$

where $f: R^n \rightarrow R^1$ is defined by (1.1.7) \square

Problem 2 is a special case of the general unconstrained optimization problem, which has the form

Minimize

$$f(\mathbf{x}) \tag{1.1.8a}$$

Subject to

$$\mathbf{x} \in R^n \tag{1.1.8b}$$

where $f: R^n \rightarrow R^1$ is a given objective function. \square

The least squares problem corresponding to (1.1.6) is linear because Y is a linear function of \mathbf{x}_1 and \mathbf{x}_2 . In many least squares problems of practical interest Y is not linear in \mathbf{x} and the corresponding least squares problem cannot usually be solved analytically. A large number of numerical methods for solving such problems exist, however, and some of these are described in Chapter 7.

Problem 3: Solving Differential Equations

Many applications of mathematics to physical problems result in the need to solve differential or integral equations. The following mathematical problem is typical.

Problem 3

Solve the boundary value problem

$$F(t, y, y^{(1)}, \dots, y^{(r)}) = 0 \tag{1.1.9a}$$

where $F: [a, b] \times R^1 \times \dots \times R^1 \rightarrow R^1$ is a given function and

$$y^{(k)}(t) = \frac{d^k}{dt^k} y(t) \quad (k = 1, \dots, r)$$

subject to a given set of r boundary conditions chosen from

$$y(a) = A \quad y(b) = B \tag{1.1.9b}$$

$$y^{(k)}(a) = A_k \quad y^{(k)}(b) = B_k \quad (k = 1, \dots, r) \quad \square \tag{1.1.9c}$$

Suppose that $\phi_k: [a, b] \rightarrow R^1$ ($k = 1, 2, \dots$) are given functions which are r times continuously differentiable on $[a, b]$. If the ϕ_k are suitable we could attempt to find a solution $Y(t; \mathbf{x}^*)$ of the boundary value problem (1.1.9) by defining $Y(t; \mathbf{x})$ by

$$Y(t; \mathbf{x}) = \sum_{k=1}^n x_k \phi_k(t) \quad (a \leq t \leq b) \tag{1.1.10}$$

for some integer $n \geq 1$ in such a way that $Y(t; \mathbf{x})$ satisfies the given set of r boundary conditions selected from (1.1.9b) and (1.1.9c) for all values of \mathbf{x} , and then determining \mathbf{x}^* by requiring that

$$F(t, Y(t; \mathbf{x}^*), \dots, \frac{d^r}{dt^r} Y(t; \mathbf{x}^*))$$

be in some sense as close as possible to zero throughout $[a, b]$. For example, we could choose $a = t_1 \leq t_2 \leq \dots \leq t_s = b$ ($s \geq n$) and define $f: R^n \rightarrow R^1$ by

$$J(x) = \sum_{i=1}^s \left[F \left(t_i, Y(t_i; x), \frac{d}{dt} Y(t_i; x), \dots, \frac{d^r}{dt^r} Y(t_i; x) \right) \right]^2 \quad (1.1.11)$$

Then $x^* \in R^n$ could be determined by minimizing f without constraints since it is supposed that we have chosen the ϕ_k so that the r boundary conditions are satisfied for all values of x .

If the boundary value problem (1.1.9) has a solution and if the ϕ_k form a complete set it is reasonable to suppose that as n is increased, $Y(t; x)$ becomes in some sense closer to the solution. Kowalik and Osborne (1968) mention a specific problem of this sort. By using this approach, integral equations and partial differential equations may be solved in principle.

Problem 4: The Calculus of Variations

Let $u: [a, b] \subset R^1 \rightarrow R^1$ be continuously differentiable on $[a, b]$ and let $F: [a, b] \times R^1 \times R^1 \rightarrow R^1$ be such that $F(t, u(t), u^{(1)}(t))$ is a continuous function of t on $[a, b]$. Let

$$I(u) = \int_a^b F(t, u(t), u^{(1)}(t)) dt \quad (1.1.12)$$

and let S be the set of functions u which are continuously differentiable on $[a, b]$ and for which $u(a) = A, u(b) = B$, where A and B are given. Then the solution of the following problem is of great importance in Physics and in Engineering.

Problem 4

Find $u^* \in S$ such that $I(u^*) \leq I(u)$ ($u \in S$), where I is defined by (1.1.12) \square

Problem 4 is the *basic problem* of variational calculus. It can be shown that the function u satisfies a certain differential equation subject to the boundary conditions $u(a) = A, u(b) = B$, but in general the differential equation is nonlinear. It is therefore of interest to consider a more direct approach to the solution of Problem 4.

Let u_1, \dots, u_n be n continuously differentiable functions on $[a, b]$ and suppose that $u_i(a) = u_i(b) = 0$ ($i = 1, \dots, n$). Suppose also that u_1, \dots, u_n are linearly independent so that

$$\sum_{k=1}^n \alpha_k u_k(t) = 0 \quad (a \leq t \leq b)$$

implies that $\alpha_k = 0$ ($k = 1, \dots, n$). Let

$$\phi(t; x) = \sum_{k=1}^n x_k u_k(t) + \frac{(A - B)}{(a - b)} t + \frac{(aB - Ab)}{(a - b)} \quad (a \leq t \leq b) \quad (1.1.13)$$

Then

$$\phi(a; x) = A \quad \phi(b; x) = B$$

Let

$$f(x) = I(\phi(t; x))$$

If we can obtain x^* such that $f(x^*) \leq f(x)$ ($x \in R^n$) then we have an approximation to u^* given by

$$u_n^*(t) = \phi(t; x^*)$$

As n increases it is reasonable to suppose that u_n^* is in some sense an increasingly good approximation to u^* .

If $I(u)$ cannot be evaluated analytically then it may be estimated by using a numerical quadrature formula so that Problem 4 may still be solved in principle by minimizing f . Burley (1974) has given several elementary examples of problems similar to Problems 3 and 4.

The three applications of unconstrained optimization considered in this section illustrate only a small part of the whole field of applications, but it is hoped that the reader will gain from studying them a little insight into the nature of unconstrained optimization and its great importance in industrial and scientific contexts. A description of some applications of nonlinear programming to a wide diversity of problems has been given by Bracken and McCormick (1968). See also Dixon (1976).

In all of the problems considered in this section, no efficient general analytical methods for minimizing the objective functions exist. A great deal of effort has, however, been expended since 1960 in constructing efficient numerical methods for solving unconstrained and constrained minimization problems. Many of these methods are easy to implement on a computer using a high level language such as FORTRAN or ALGOL and much can be said about their behaviour without resorting to advanced mathematical analysis. Certain numerical methods for optimization have proved repeatedly to be not only robust but also efficient; that is to say, they seldom fail to obtain a sufficiently accurate estimate of the solution of a given optimization problem and do so with tolerably few evaluations of the objective function and of its derivatives if the latter are required. This is not to say, however, that these methods are universally applicable. There is no panacea for all optimization problems.

In the subsequent chapters of this book some methods for unconstrained optimization are described in sufficient detail to permit them to be implemented on a computer. Computer programs are not given, because many have been published, and some of these are readily available from sources which are given where appropriate in the text. The reader is, however, urged to write programs and to run them in order to gain that insight into the methods which are described in this book which can be acquired only from numerical experience.

1.2 Fundamental Linear Algebra

In this section some fundamental ideas from linear algebra which are used subsequently are collected, and some notation is described. For greater detail than it is possible to give here the reader should consult, for example, Hadley (1969), Noble (1969), or Franklin (1968). It is not necessary for the reader to understand the content of this section in detail if it is required only to use the methods which are described in subsequent chapters, but an understanding of this section is, however, necessary to appreciate the theory which underlies these methods.

By R^n is meant the set of $n \times 1$ vectors $x = [x_1, \dots, x_n]^T$ in which $x_i \in R^1$ ($i = 1, \dots, n$) where R^1 is the set of real numbers. If we admit the usual rules for adding $n \times 1$ vectors $x, y \in R^n$ and multiplying them by real numbers α , namely

$$x + y = [x_1 + y_1, \dots, x_n + y_n]^T \quad (1.2.1)$$

and

$$\alpha x = [\alpha x_1, \dots, \alpha x_n]^T \quad (1.2.2)$$

then R^n together with the operations defined by (1.2.1) and (1.2.2) constitute a linear space. We shall subsequently refer to the linear space as R^n . A function $\|\cdot\|: R^n \rightarrow R^1$ such that

$$\|x\| \geq 0, \quad (\forall x \in R^n)$$

$$(\|x\| = 0) \Leftrightarrow (x = 0)$$

$$\|\alpha x\| = |\alpha| \|x\| \quad (\forall x \in R^n; \forall \alpha \in R^1)$$

and

$$\|x + y\| \leq \|x\| + \|y\| \quad (\forall x, y \in R^n)$$

is called a *vector norm* on R^n . Commonly used vector norms on R^n are the *maximum* norm $\|\cdot\|_M$ defined by

$$\|x\|_M = \max_{1 \leq i \leq n} \{|x_i|\} \quad (1.2.3)$$

the *taxicab* norm $\|\cdot\|_T$ defined by

$$\|x\|_T = \sum_{i=1}^n |x_i| \quad (1.2.4)$$

and the *Euclidean* norm $\|\cdot\|_E$ defined by

$$\|x\|_E = \left(\sum_{i=1}^n x_i^2 \right)^{1/2} \quad (1.2.5)$$

where $x = [x_1, \dots, x_n]^T$

It can be shown that all vector norms on R^n are equivalent in the sense that if $\|\cdot\|_A$ and $\|\cdot\|_B$ are any two such norms then there exist $\alpha, \beta \in (0, \infty)$ such that

$$\alpha \|x\|_A \leq \|x\|_B \leq \beta \|x\|_A \quad (\forall x \in R^n)$$

The linear space R^n together with a vector norm $\|\cdot\|$ on R^n is called a *normed linear space*; we denote it by $\{R^n, \|\cdot\|\}$.

The *scalar product* $x^T y$ of the vectors $x, y \in R^n$ is defined by

$$x^T y = \sum_{i=1}^n x_i y_i \quad (1.2.6)$$

In the normed linear space $\{R^n, \|\cdot\|_E\}$ we have Schwarz' inequality, namely

$$|x^T y| \leq \|x\|_E \|y\|_E \quad (\forall x, y \in R^n) \quad (1.2.7)$$

with equality if and only if there exists $\lambda \in R^1$ such that $y = \lambda x$.

The vectors $x, y \in R^n$ are *orthogonal* if and only if

$$x^T y = 0 \quad (1.2.8)$$

and the vector $x \in R^n$ is *normalized* (to unity) if and only if

$$x^T x = 1 \quad (1.2.9)$$

The set of vectors $\{x_i \in R^n | i = 1, \dots, m; m \leq n\}$ is an *orthonormal* set if and only if

$$x_i^T x_j = \delta_{ij} \quad (i, j = 1, \dots, m) \quad (1.2.10)$$

where the Kronecker delta δ_{ij} is defined by

$$\delta_{ij} = \begin{cases} 1 & (i = j) \\ 0 & (i \neq j) \end{cases} \quad (1.2.11)$$

The set of vectors $\{\mathbf{x}_i \in R^n \mid i = 1, \dots, m; m \leq n\}$ is a *linearly independent* set if and only if

$$\sum_{i=1}^m c_i \mathbf{x}_i = \mathbf{0} \quad (1.2.12)$$

implies that $c_i = 0$ ($i = 1, \dots, m$), and is *linearly dependent* otherwise. A set of vectors $\{\mathbf{u}_i \in R^n \mid i = 1, \dots, m; m \geq n\}$ *spans* or *generates* the linear space R^n if and only if every vector $\mathbf{x} \in R^n$ is expressible in the form

$$\mathbf{x} = \sum_{i=1}^m c_i \mathbf{u}_i \quad (1.2.13)$$

A *basis* for the linear space R^n is a linearly independent set of vectors in R^n which spans R^n . The set of vectors $\{\mathbf{e}_i \in R^n \mid i = 1, \dots, n\}$ where

$$\mathbf{e}_i = [0, \dots, 0, \overset{(i)}{1}, 0, \dots, 0]^T \quad (i = 1, \dots, n) \quad (1.2.14)$$

in which \mathbf{e}_i has all components equal to zero save the i th which is equal to unity, clearly forms a basis for R^n since if $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T$ then

$$\mathbf{x} = \sum_{i=1}^n x_i \mathbf{e}_i \quad (1.2.15)$$

It can be shown that every basis for R^n contains exactly n vectors, and that every set of n orthogonal vectors in R^n excluding $\mathbf{0}$ forms a basis for R^n . An *orthonormal* basis for R^n is a set of n orthonormal vectors. For example, $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ defined by (1.2.14) is an orthonormal basis for R^n .

Given any basis for R^n , an orthonormal basis for R^n may be constructed by using the Gram-Schmidt procedure. Let $\{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be a basis for R^n and let $\mathbf{v}_i, \mathbf{w}_i$ ($i = 1, \dots, n$) be generated recursively from

$$\mathbf{v}_1 = \mathbf{u}_1 \quad \mathbf{w}_1 = \frac{\mathbf{v}_1}{(\mathbf{v}_1^T \mathbf{v}_1)^{1/2}} \quad (1.2.16a)$$

$$\mathbf{v}_i = \mathbf{u}_i - \sum_{k=1}^{i-1} (\mathbf{w}_k^T \mathbf{u}_i) \mathbf{w}_k \quad \mathbf{w}_i = \frac{\mathbf{v}_i}{(\mathbf{v}_i^T \mathbf{v}_i)^{1/2}} \quad (i = 2, \dots, n) \quad (1.2.16b)$$

Then it is easy to show (Exercise 1.2.7) that

$$\mathbf{w}_i^T \mathbf{w}_j = \delta_{ij} \quad (i, j = 1, \dots, n) \quad (1.2.17)$$

so that $\{\mathbf{w}_1, \dots, \mathbf{w}_n\}$ is an orthonormal basis for R^n .

A set $S \subset R^n$ is a *subspace* of the linear space R^n if and only if for each pair $\mathbf{x}', \mathbf{x}''$ of vectors in S , the vector $\alpha \mathbf{x}' + \beta \mathbf{x}''$ is also in S where α, β are any real numbers. As with R^n , the dimensionality m of a subspace S of R^n is the number of vectors in any basis for S . Clearly $m \leq n$. Let $\{\mathbf{u}_i \in R^n \mid i = 1, \dots, m\}$ be a basis for the m -dimensional subspace S of R^n . Then clearly an orthonormal basis for S can be constructed by using the Gram-Schmidt procedure (1.2.16), where i runs from 2 to m .

Let S be a subspace of R^n and let $\hat{x} \in R^n$ be given. Then the set V defined by

$$V = \{v \in R^n | v = \hat{x} + y, y \in S\}$$

is called a *linear variety*. It is convenient to use the notation

$$V = \hat{x} + S$$

In conformity with the notation used to represent column vectors, the *transpose* A^T of the $m \times n$ matrix A with elements a_{ij} ($i = 1, \dots, m; j = 1, \dots, n$) is the $n \times m$ matrix B with elements b_{ij} ($i = 1, \dots, n; j = 1, \dots, m$) such that

$$b_{ij} = a_{ji} \quad (i = 1, \dots, n; j = 1, \dots, m)$$

An $n \times n$ matrix A is *symmetric* if and only if

$$A^T = A \quad (1.2.18)$$

and is *skew-symmetric* or *anti-symmetric* if and only if

$$A^T = -A \quad (1.2.19)$$

Any $n \times n$ matrix A may be expressed as the sum of a symmetric matrix B and a skew-symmetric matrix C where

$$B = \frac{1}{2}(A + A^T) \quad C = \frac{1}{2}(A - A^T) \quad (1.2.20)$$

If A and B are any $n \times n$ matrices, then it follows from the definitions of matrix multiplication and transpose that

$$(AB)^T = B^T A^T \quad (1.2.21)$$

We write

$$A = [a_{ij}] \quad (i = 1, \dots, m; j = 1, \dots, n)$$

to denote that A is an $m \times n$ matrix with elements a_{ij} , and we write

$$A = \text{Diag}(a_{11}, \dots, a_{nn})$$

to denote that A is the diagonal $n \times n$ matrix with diagonal elements a_{ii} ($i = 1, \dots, n$).

It is assumed that the reader is familiar with the more elementary properties of the determinant, $\text{Det}(A)$, of an $n \times n$ matrix A , and with the elementary properties of the inverse A^{-1} of A . In particular, if A and B are $n \times n$ matrices then

$$\text{Det}(A^T) = \text{Det}(A) \quad (1.2.22)$$

and

$$\text{Det}(AB) = \text{Det}(A) \text{Det}(B) \quad (1.2.23)$$

Also, A is singular, so that A^{-1} does not exist, if and only if $\text{Det}(A) = 0$, and if A and B are nonsingular then

$$(AB)^{-1} = B^{-1} A^{-1} \quad (1.2.24)$$

It is also assumed that the reader is familiar with the idea of *partitioning* of matrices. For example, if A is an $m \times n$ matrix with elements a_{ij} ($i = 1, \dots, m; j = 1, \dots, n$) and columns a_j ($j = 1, \dots, n$) then we write

$$A = \begin{bmatrix} a_1 & \dots & a_n \end{bmatrix}$$

to denote the partitioning of A into columns. More generally, to denote partitioning of

A into submatrices A_{ij} ($i = 1, \dots, p \leq m; j = 1, \dots, q \leq n$) we use a similar notation. For example, if $p = 2$ and $q = 3$ we write

$$A = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \end{bmatrix}$$

A real symmetric $n \times n$ matrix A has n orthonormal eigenvectors u_i ($i = 1, \dots, n$) corresponding to n real eigenvalues λ_i ($i = 1, \dots, n$) (which need not be distinct) so that

$$Au_i = \lambda_i u_i \quad (i = 1, \dots, n) \quad (1.2.25)$$

and

$$u_i^T u_j = \delta_{ij} \quad (i, j = 1, \dots, n) \quad (1.2.26)$$

An $n \times n$ matrix B is *orthogonal* if and only if

$$BB^T = B^T B = I \quad (1.2.27)$$

It is easy to see that if A is an $n \times n$ real symmetric matrix then there exists an orthogonal matrix P such that $P^T A P$ is a diagonal matrix. Indeed, P is the matrix with columns u_1, \dots, u_n where the u_i are the normalized eigenvectors of A which satisfy (1.2.25) and (1.2.26).

Let A be an $n \times n$ matrix and let x be an $n \times 1$ vector. Then, irrespective of whether A is symmetric,

$$\begin{aligned} x^T A x &= \sum_{i=1}^n \sum_{j=1}^n x_i a_{ij} x_j \\ &= \sum_{i=1}^n \sum_{j=1}^n x_i a_{ji} x_j \end{aligned}$$

so that

$$x^T A x = x^T B x$$

where the symmetric matrix B is defined in (1.2.20). Therefore when considering quadratic forms $x^T A x$ we may always suppose that the corresponding matrix A is symmetric.

An $n \times n$ matrix A is *positive definite* if and only if

$$x^T A x > 0 \quad (\forall x \in R^n, x \neq 0) \quad (1.2.28a)$$

and is *positive semi-definite* if and only if

$$x^T A x \geq 0 \quad (\forall x \in R^n) \quad (1.2.29a)$$

Similarly A is *negative definite* if and only if

$$x^T A x < 0 \quad (\forall x \in R^n, x \neq 0) \quad (1.2.28b)$$

and A is *negative semi-definite* if and only if

$$x^T A x \leq 0 \quad (\forall x \in R^n) \quad (1.2.29b)$$

Setting $x = e_i$ where e_i is defined by (1.2.14), it follows immediately from (1.2.28a) that if A is positive definite then $a_{ii} > 0$ ($i = 1, \dots, n$). Similar results hold when A is positive semi-definite, negative semi-definite, and negative definite, respectively.

It can be shown that if A is symmetric positive definite, then there exists a symmetric positive definite matrix M such that

$$M^2 = A \quad (1.2.30)$$

Furthermore, a matrix A is positive definite if and only if all of its eigenvalues are positive.

Let B be an $m \times n$ matrix. Then the *rank* of B , which we write $r(B)$, is the maximum number of linearly independent columns of B . The $n \times n$ matrix A is nonsingular if and only if $r(A) = n$. Hence if $r(A) = n$, then $Au = 0$ if and only if $u = 0$. Furthermore if C is an $m \times n$ real matrix with $r(C) = n$ then $Cu = 0$ if and only if $u = 0$, while if $r(C) < n$, then $Cu = 0$ has a nonzero solution. If $m < n$ then necessarily $r(C) < n$ and so $Cu = 0$ always has a nonzero solution if $m < n$. Finally in connection with the concept of rank, we note that if A and B are any two matrices such that the product AB exists, then

$$r(AB) \leq \min \{r(A), r(B)\} \quad (1.2.31)$$

Let $\{R^n, \|\cdot\|\}$ be a normed linear space, and let A be any $n \times n$ matrix. Then the *matrix norm* $\|\cdot\|$ corresponding to the vector norm $\|\cdot\|$ is defined by

$$\|A\| = \max_{x \neq 0} \frac{\|Ax\|}{\|x\|} \quad (1.2.32)$$

which implies that

$$\|Ax\| \leq \|A\| \|x\| \quad (\forall x \in R^n) \quad (1.2.33)$$

and that if A and B are any $n \times n$ matrices then

$$\|AB\| \leq \|A\| \|B\| \quad (1.2.34)$$

Note that the same symbol $\|\cdot\|$ is used for a vector norm and its corresponding matrix norm, but in practice no confusion results because which norm is intended is always clear from the context in which it appears.

The matrix norms corresponding to the vector norms $\|\cdot\|_M$, $\|\cdot\|_T$, and $\|\cdot\|_E$ are given by

$$\|A\|_M = \max_{1 \leq i \leq n} \left\{ \sum_{j=1}^n |a_{ij}| \right\} \quad (1.2.35)$$

$$\|A\|_T = \max_{1 \leq j \leq n} \left\{ \sum_{i=1}^n |a_{ij}| \right\} \quad (1.2.36)$$

$$\|A\|_E = \left(\max_{1 \leq i \leq n} \{\lambda_i\} \right)^{1/2} \quad (1.2.37)$$

in which A is an $n \times n$ matrix with elements a_{ij} ($i, j = 1, \dots, n$) and $\lambda_1, \dots, \lambda_n$ are the n real eigenvalues of the matrix $A^H A$, where A^H is the Hermitian conjugate matrix of A defined by

$$A^H = (\bar{A})^T$$

\bar{A} denoting the complex conjugate of A , so that \bar{A} has elements \bar{a}_{ij} ($i, j = 1, \dots, n$). If A is real symmetric then $A^H = A$ and so the eigenvalues of $A^H A$ are μ_1^2, \dots, μ_n^2 where

μ_1, \dots, μ_n are the eigenvalues of A , which are all real. Hence when A is an $n \times n$ real symmetric matrix then

$$\|A\|_E = \max_{1 \leq i \leq n} \{|\mu_i|\} \quad (1.2.38)$$

Finally, let $\|\cdot\|$ be any matrix norm. Then we have the following lemma, which is used in Chapter 7.

Lemma 1.2.1

- If 1. A and B are $n \times n$ matrices;
 2. A is nonsingular and $\|A^{-1}\| \leq \alpha$;
 3. $\|B - A\| \leq \beta$;
 4. $\alpha\beta < 1$,

then B is nonsingular and

$$\|B^{-1}\| \leq \frac{\alpha}{(1 - \alpha\beta)}$$

□

For a proof of Lemma 1.2.1 see Exercise 1.2.14.

1.3 Fundamental Analysis

In this section we collect some results from the calculus of functions of several variables and describe some notation which is used subsequently. For more detail than it is possible to give here the reader should consult, for example, Russell (1970), Burkill (1970), and Mangasarian (1969). As with Section 1.2 it is not necessary to understand in detail the content of this section if the reader is not primarily interested in the theory underlying the numerical methods which are described subsequently in the text.

It is very convenient to express many mathematical ideas connected with the theory of optimization in a geometrical form. We speak of $x, y \in R^n$ as *points* in the linear space R^n , and we define a *distance* $\rho(x, y)$ between the points x, y in the normed linear space $\{R^n, \|\cdot\|\}$ by

$$\rho(x, y) = \|x - y\| \quad (1.3.1)$$

We generalize the idea of open and closed intervals in R^1 by means of open and closed *balls* in $\{R^n, \|\cdot\|\}$.

Definition 1.3.1

The set of points $x \in R^n$ such that $\|x - \hat{x}\| \leq r$, where $r > 0$ is a given real number and \hat{x} is a given point in R^n is called the *closed ball* $B[\hat{x}, r]$ with centre \hat{x} and radius r . Thus

$$B[\hat{x}, r] = \{x \in R^n \mid \|x - \hat{x}\| \leq r\}$$

The *open ball* $B(\hat{x}, r)$ with centre \hat{x} and radius r is defined by

$$B(\hat{x}, r) = \{x \in R^n \mid \|x - \hat{x}\| < r\}$$

□

Using the idea of distance between points in the normed linear space $\{R^n, \|\cdot\|\}$ we can generalize the idea of convergence of a sequence $\{x^{(k)}\}$ in $\{R^1, |\cdot|\}$.

Definition 1.3.2

The sequence $\{x^{(k)}\}$ in the normed linear space $\{R^n, \|\cdot\|\}$ converges to $x^* \in R^n$ if and only if

$$\lim_{k \rightarrow \infty} \|x^{(k)} - x^*\| = 0 \quad \square$$

Because, as we have seen in Section 1.2, all norms on R^n are equivalent, we have

$$\left(\lim_{k \rightarrow \infty} \|x^{(k)} - x^*\| = 0 \right) \Leftrightarrow \left(\lim_{k \rightarrow \infty} \|x^{(k)} - x^*\|' = 0 \right)$$

where $\|\cdot\|$ and $\|\cdot\|'$ are any two vector norms on R^n . Hence convergence of $\{x^{(k)}\}$ to x^* with respect to any norm on R^n implies convergence of $\{x^{(k)}\}$ to x^* with respect to any other norm on R^n .

In conformity with the definition of a Cauchy sequence of real numbers, we say that the sequence $\{x^{(k)}\}$ in the normed linear space $\{R^n, \|\cdot\|\}$ is a *Cauchy sequence* if and only if given $\epsilon > 0$, there exists an integer l such that $\|x^{(j)} - x^{(k)}\| < \epsilon$ whenever $j > l$ and $k > l$ simultaneously. It can be shown that if $\{x^{(k)}\}$ is a Cauchy sequence in $\{R^n, \|\cdot\|\}$ then $\{x^{(k)}\}$ has a limit $x^* \in R^n$.

We can generalize the idea of continuity of $f: R^1 \rightarrow R^1$ as follows.

Definition 1.3.3

A function $f: D \subseteq R^n \rightarrow R^m$ is *continuous* at $\hat{x} \in D$ if and only if

$$\|f(x) - f(\hat{x})\|_m \rightarrow 0$$

whenever

$$\|x - \hat{x}\|_n \rightarrow 0$$

where $\|\cdot\|_m$ is a norm on R^m and $\|\cdot\|_n$ is a norm on R^n . \square

By definition 1.3.3, $f: R^n \rightarrow R^1$ is continuous at $\hat{x} \in R^n$ if $|f(x) - f(\hat{x})| \rightarrow 0$ whenever $\|x - \hat{x}\| \rightarrow 0$, where $\|\cdot\|$ is any norm on R^n .

We shall denote the first partial derivatives of $f: R^n \rightarrow R^1$ by $\partial_i f(x)$ ($i = 1, \dots, n$) so that

$$\partial_i f(x) = \frac{\partial}{\partial x_i} f(x) \quad (i = 1, \dots, n) \quad (1.3.2)$$

We denote the second partial derivatives of f by $\partial_i \partial_j f(x)$ ($i, j = 1, \dots, n$) so that

$$\partial_i \partial_j f(x) = \frac{\partial^2}{\partial x_i \partial x_j} f(x) \quad (i, j = 1, \dots, n) \quad (1.3.3)$$

Let $f: R^n \rightarrow R^1$ have first partial derivatives in R^n . Then the *gradient vector* $g(x)$ of f at $x \in R^n$ is defined by

$$g(x) = [\partial_1 f(x), \dots, \partial_n f(x)]^T. \quad (1.3.4)$$

It is sometimes convenient to use the notation $\nabla f(x)$ for the gradient vector of f at x .

Let $f: R^n \rightarrow R^1$ have second partial derivatives in R^n . Then the *Hessian matrix* $G(x)$ of f at x is defined by

$$G(x) = [\partial_i \partial_j f(x)] \quad (i, j = 1, \dots, n) \quad (1.3.5)$$

If f has continuous second partial derivatives in R^n then

$$\partial_i \partial_j f(x) = \partial_j \partial_i f(x) \quad (i, j = 1, \dots, n)$$