国外数学名著系列（续一）

（影印版）42

J. E. Dennis Jr.   Robert B. Schnabel

Numerical Methods for Unconstrained Optimization and Nonlinear Equations

无约束最优化与非线性方程的数值方法

科学出版社
www.sciencep.com

国外数学名著系列(影印版) 42

Numerical Methods for Unconstrained Optimization and Nonlinear Equations

无约束最优化与非线性方程的数值方法

J. E. Dennis Jr. Robert B. Schnabel

科 学 出 版 社

北 京

图字：01-2008-5108

Original American edition published by:
SIAM: Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania

图书在版编目 (CIP) 数据

无约束最优化与非线性方程的数值方法 = Numerical Methods for Unconstrained Optimization and Nonlinear Equations / (美)丹尼斯(Dennis, J. E.)等著. —影印版. —北京：科学出版社，2009
(国外数学名著系列；42)
ISBN 978-7-03-023482-7
I. 无… II. 丹… III. ①最优化算法-数值计算-英文 ②非线性方程-数值计算-英文 IV. O242.23 O241.7
中国版本图书馆 CIP 数据核字(2008) 第 186201 号

责任编辑：范庆奎 / 责任印刷：钱玉芬 / 封面设计：黄华斌

科学出版社 出版
北京东黄城根北街 16 号
邮政编码：100717
http://www.sciencep.com
北京佳信达艺术印刷有限公司 印刷
科学出版社发行 各地新华书店经销
*
2009 年 1 月第 一 版 开本：B5(720 × 1000)
2009 年 1 月第一次印刷 印张：25
印数：1—2 500 字数：476 000

定价：86.00 元

如有印装质量问题，我社负责调换

Preface to the Classics Edition

We are delighted that SIAM is republishing our original 1983 book after what many in the optimization field have regarded as "premature termination" by the previous publisher. At 12 years of age, the book may be a little young to be a "classic," but since its publication it has been well received in the numerical computation community. We are very glad that it will continue to be available for use in teaching, research, and applications.

We set out to write this book in the late 1970s because we felt that the basic techniques for solving small to medium-sized nonlinear equations and unconstrained optimization problems had matured and converged to the point where they would remain relatively stable. Fortunately, the intervening years have confirmed this belief. The material that constitutes most of this book—the discussion of Newton-based methods, globally convergent line search and trust region methods, and secant (quasi-Newton) methods for nonlinear equations, unconstrained optimization, and nonlinear least squares—continues to represent the basis for algorithms and analysis in this field. On the teaching side, a course centered around Chapters 4 to 9 forms a basic, in-depth introduction to the solution of nonlinear equations and unconstrained optimization problems. For researchers or users of optimization software, these chapters give the foundations of methods and software for solving small to medium-sized problems of these types.

We have not revised the 1983 book, aside from correcting all the typographical errors that we know of. (In this regard, we especially thank Dr. Oleg Burdakov who, in the process of translating the book for the Russian edition published by Mir in 1988, found numerous typographical errors.) A main reason for not revising the book at this time is that it would have delayed its republication substantially. A second reason is that there appear to be relatively few places where the book needs updating. But inevitably there are some. In our opinion, the main developments in the solution of small to medium-sized unconstrained optimization and nonlinear equations problems since the publication of this book, which a current treatment should include, are

1. improved algorithms and analysis for trust region methods for unconstrained optimization in the case when the Hessian matrix is indefinite [1, 2] and
2. improved global convergence analysis for secant (quasi-Newton) methods [3].

A third, more recent development is the field of automatic (or computational) differentiation [4]. Although it is not yet fully mature, it is clear that this development is increasing the availability of analytic gradients and Jacobians and therefore reducing the cases where finite difference approximations to these derivatives are needed. A fourth, more minor but still significant development is a new, more stable modified Cholesky factorization method [5, 6]. Far more progress has been made in the solution of large nonlinear equations and unconstrained optimization problems. This includes the development or improvement of conjugate gradient, truncated-Newton, Krylov-subspace, and limited-memory methods. Treating these fully would go beyond the scope of this book even if it were revised, and fortunately some excellent new references are emerging, including [7]. Another important topic that is related to but not within the scope of this book is that of new derivative-free methods for unconstrained optimization [8].

The appendix of this book has had an impact on software in this field. The IMSL library created their unconstrained optimization code from this appendix, and the UNCMIN software [9] created in conjunction with this appendix has been and continues to be a widely used package for solving unconstrained optimization problems. This software also has been included in a number of software packages and other books. The UNCMIN software continues to be available from the second author (bobby@cs.colorado.edu).

Finally, one of the most important developments in our lives since 1983 has been the emergence of a new generation: a granddaughter for one of us, a daughter and son for the other. This new edition is dedicated to them in recognition of the immense joy they have brought to our lives and with all our hopes and wishes for the lives that lay ahead for them.

[1] J. J. Moré and D. C. Sorensen, *Computing a trust region step,* SIAM J. Sci. Statist. Comput., 4 (1983), pp. 553–572.

[2] G. A. Shultz, R. B. Schnabel, and R. H. Byrd, *A family of trust region based algorithms for unconstrained minimization with strong global convergence properties,* SIAM J. Numer. Anal., 22 (1985), pp. 47–67.

[3] R. H. Byrd, J. Nocedal, and Y. Yuan, *Global convergence of a class of quasi-Newton methods on convex problems,* SIAM J. Numer. Anal., 24 (1987), pp. 1171–1189.

[4] A. Griewank and G. F. Corliss, eds., *Automatic Differentiation of Algorithms: Theory, Implementation, and Application,* Society for Industrial and Applied Mathematics, Philadelphia, PA, 1991.

[5] R. B. Schnabel and E. Eskow, *A new modified Cholesky factorization,* SIAM J. Sci. Statist. Comput., 11 (1990), pp. 1136–1158.

[6] E. Eskow and R. B. Schnabel, *Software for a new modified Cholesky factorization,* ACM Trans. Math. Software, 17 (1991), pp. 306–312.

[7] C. T. Kelley, *Iterative Methods for Linear and Nonlinear Equations,* Society for Industrial and Applied Mathematics, Philadelphia, PA, 1995.

[8] J. E. Dennis, Jr. and V. Torczon, *Direct search methods on parallel computers,* SIAM J. Optim., 1 (1991), pp. 448–474.

[9] R. B. Schnabel, J. E. Koontz, and B. E. Weiss, *A modular system of algorithms for unconstrained minimization,* ACM Trans. Math. Software, 11 (1985), pp. 419–440.

Preface

This book offers a careful introduction, at a low level of mathematical and computational sophistication, to the numerical solution of problems in unconstrained optimization and systems of nonlinear equations. We have written it, beginning in 1977, because we feel that the algorithms and theory for small-to-medium-size problems in this field have reached a mature state, and that a comprehensive reference will be useful. The book is suitable for graduate or upper-level undergraduate courses, but also for self-study by scientists, engineers, and others who have a practical interest in such problems.

The minimal background required for this book would be calculus and linear algebra. The reader should have been at least exposed to multivariable calculus, but the necessary information is surveyed thoroughly in Chapter 4. Numerical linear algebra or an elementary numerical methods course would be helpful; the material we use is covered briefly in Section 1.3 and Chapter 3.

The algorithms covered here are all based on Newton's method. They are often called **Newton-like**, but we prefer the term **quasi-Newton**. Unfortunately, this term is used by specialists for the subclass of these methods covered in our Chapters 8 and 9. Because this subclass consists of sensible multidimensional generalizations of the secant method, we prefer to call them secant methods. Particular secant methods are usually known by the proper names of their discoverers, and we have included these servings of alphabet soup, but we have tried to suggest other descriptive names commensurate with their place in the overall scheme of our presentation.

The heart of the book is the material on computational methods for multidimensional unconstrained optimization and nonlinear equation problems covered in Chapters 5 through 9. Chapter 1 is introductory and will be more useful for students in pure mathematics and computer science than for readers with some experience in scientific applications. Chapter 2, which covers the one-dimensional version of our problems, is an overview of our approach to the subject and is essential motivation. Chapter 3 can be omitted by readers who have studied numerical linear algebra, and Chapter 4 can be omitted by those who have a good background in multivariable calculus. Chapter 10 gives a fairly complete treatment of algorithms for nonlinear least squares, an important type of unconstrained optimization problem that, owing to its special structure, is solved by special methods. It draws heavily on the chapters that precede it. Chapter 11 indicates some research directions in which the field is headed; portions of it are more difficult than the preceding material.

We have used the book for undergraduate and graduate courses. At the lower level, Chapters 1 through 9 make a solid, useful course; at the graduate level the whole book can be covered. With Chapters 1, 3, and 4 as remedial reading, the course takes about one quarter. The remainder of a semester is easily filled with these chapters or other material we omitted.

The most important omitted material consists of methods not related to Newton's method for solving unconstrained minimization and nonlinear equation problems. Most of them are important only in special cases. The Nelder-Meade simplex algorithm [see, e.g., Avriel (1976)], an effective algorithm for problems with less than five variables, can be covered in an hour. Conjugate direction methods [see, e.g., Gill, Murray, and Wright (1981)] properly belong in a numerical linear algebra course, but because of their low storage requirements they are useful for optimization problems with very large numbers of variables. They can be covered usefully in two hours and completely in two weeks.

The omission we struggled most with is that of the Brown-Brent methods. These methods are conceptually elegant and startlingly effective for partly linear problems with good starting points. In their current form they are not competitive for general-purpose use, but unlike the simplex or conjugate-direction algorithms, they would not be covered elsewhere. This omission can be remedied in one or two lectures, if proofs are left out [see, e.g., Dennis (1977)]. The final important omission is that of the continuation or homotopy-based methods, which enjoyed a revival during the seventies. These elegant ideas can be effective as a last resort for the very hardest problems but are not yet competitive for most problems. The excellent survey by Allgower and Georg (1980) requires at least two weeks.

We have provided many exercises; many of them further develop ideas that are alluded to briefly in the text. The large appendix (by Schnabel) is intended to provide both a mechanism for class projects and an important reference for readers who wish to understand the details of the algorithms and perhaps to develop their own versions. The reader is encouraged to read the preface to the appendix at an early stage.

Several problems of terminology and notation were particularly troublesome. We have already mentioned the confusion over the terms "quasi-Newton" and "secant methods." In addition, we use the term "unconstrained optimization" in the title but "unconstrained minimization" in the text, since technically we consider only minimization. For maximization, turn the problems upside-down. The important term "global" has several interpretations, and we try to explain ours clearly in Section 1.1. Finally, a major notational problem was how to differentiate between the ith component of an n-vector x, a *scalar* usually denoted by x_i, and the ith iteration in a sequence of such x's, a *vector* also usually denoted x_i. After several false starts, we decided to allow this conflicting notation, since the intended meaning is always clear from the context; in fact, the notation is rarely used in both ways in any single section of the text.

We wanted to keep this book as short and inexpensive as possible without slighting the exposition. Thus, we have edited some proofs and topics in a merciless fashion. We have tried to use a notion of rigor consistent with good taste but subservient to insight, and to include proofs that give insight while omitting those that merely substantiate results. We expect more criticism for omissions than for inclusions, but as every teacher knows, the most difficult but important part in planning a course is deciding what to leave out.

We sincerely thank Idalia Cuellar, Arlene Hunter, and Dolores Pendel for typing the numerous drafts, and our students for their specific identification of unclear passages. David Gay, Virginia Klema, Homer Walker, Pete Stewart, and Layne Watson used drafts of the book in courses at MIT, Lawrence Livermore Laboratory, University of Houston, University of New Mexico, University of Maryland, and VPI, and made helpful suggestions. Trond Steihaug and Mike Todd read and commented helpfully on portions of the text.

Rice University — *J. E. Dennis, Jr.*

University of Colorado at Boulder — *Robert B. Schnabel*

Contents

||3|| NUMERICAL LINEAR ALGEBRA BACKGROUND 40

||4|| MULTIVARIABLE CALCULUS BACKGROUND 69

||5|| NEWTON'S METHOD FOR NONLINEAR EQUATIONS AND UNCONSTRAINED MINIMIZATION 86

||6|| GLOBALLY CONVERGENT MODIFICATIONS OF NEWTON'S METHOD 111

||7|| STOPPING, SCALING, AND TESTING *155*

||8|| SECANT METHODS FOR SYSTEMS OF NONLINEAR EQUATIONS *168*

||9|| SECANT METHODS FOR UNCONSTRAINED MINIMIZATION *194*

||10|| NONLINEAR LEAST SQUARES *218*

||11|| METHODS FOR PROBLEMS WITH SPECIAL STRUCTURE *239*

Before we begin, a program note

The first four chapters of this book contain the background material and motivation for the study of multivariable nonlinear problems. In Chapter 1 we introduce the problems we will be considering. Chapter 2 then develops some algorithms for nonlinear problems in just one variable. By developing these algorithms in a way that introduces the basic philosophy of all the nonlinear algorithms to be considered in this book, we hope to provide an accessible and solid foundation for the study of multivariable nonlinear problems. Chapters 3 and 4 contain the background material in numerical linear algebra and multivariable calculus required to extend our consideration to problems in more than one variable.

1

Introduction

This book discusses the methods, algorithms, and analysis involved in the computational solution of three important nonlinear problems: solving systems of nonlinear equations, unconstrained minimization of a nonlinear functional, and parameter selection by nonlinear least squares. Section 1.1 introduces these problems and the assumptions we will make about them. Section 1.2 gives some examples of nonlinear problems and discusses some typical characteristics of problems encountered in practice; the reader already familiar with the problem area may wish to skip it. Section 1.3 summarizes the features of finite-precision computer arithmetic that the reader will need to know in order to understand the computer-dependent considerations of the algorithms in the text.

1.1 PROBLEMS TO BE CONSIDERED

This book discusses three nonlinear problems in real variables that arise often in practice. They are mathematically equivalent under fairly reasonable hypotheses, but we will not treat them all with the same algorithm. Instead we will show how the best current algorithms seek to exploit the structure of each problem.

The simultaneous nonlinear equations problem (henceforth called "nonlinear equations") is the most basic of the three and has the least exploitable structure. It is

$$\text{Given} \quad F: \mathbb{R}^n \longrightarrow \mathbb{R}^n,$$
$$\text{find} \quad x_* \in \mathbb{R}^n \quad \text{for which} \quad F(x_*) = 0 \in \mathbb{R}^n, \tag{1.1.1}$$

where $\mathbb{R}^n$ denotes n-dimensional Euclidean space. Of course, (1.1.1) is just the standard way of denoting a system of n nonlinear equations in n unknowns, with the convention that right-hand side of each equation is zero. An example is

$$F(x_1, x_2) = \begin{pmatrix} x_1^2 + x_2^3 + 7 \\ x_1 + x_2 + 1 \end{pmatrix},$$

which has $F(x_*) = 0$ for $x_* = (1, -2)^T$.

Certainly the x_* that solves (1.1.1) would be a minimizer of

$$\sum_{i=1}^{n} (f_i(x))^2,$$

where $f_i(x)$ denotes the ith component function of F. This is a special case of the *unconstrained minimization* problem

$$\text{Given} \quad f: \mathbb{R}^n \longrightarrow \mathbb{R}$$
$$\text{find } x_* \in \mathbb{R}^n \quad \text{for which} \quad f(x_*) \le f(x) \text{ for every } x \in \mathbb{R}^n, \tag{1.1.2}$$

which is the second problem we will consider. Usually (1.1.2) is abbreviated to

$$\min_{x \in \mathbb{R}^n} f: \mathbb{R}^n \longrightarrow \mathbb{R}. \tag{1.1.3}$$

An example is

$$\min_{x \in \mathbb{R}^n} f(x_1, x_2, x_3) = (x_1 - 3)^2 + (x_2 + 5)^4 + (x_3 - 8)^2,$$

which has the solution $x_* = (3, -5, 8)^T$.

In some applications, one is interested in solving a constrained version of (1.1.3),

$$\min_{x \in \Omega \subset \mathbb{R}^n} f: \mathbb{R}^n \longrightarrow \mathbb{R}, \tag{1.1.4}$$

where Ω is a closed connected region. If the solution to (1.1.4) lies in the interior of Ω, then (1.1.4) can still be viewed as an unconstrained minimization problem. However, if x_* is a boundary point of Ω, then the minimization of f over Ω becomes a constrained minimization problem. We will not consider the constrained problem because less is known about how it should be solved, and there is plenty to occupy us in considering unconstrained problems. Furthermore, the techniques for solving unconstrained problems are the foundation for constrained-problem algorithms. In fact, many attempts to solve constrained problems boil down to either solving a related unconstrained minimization problem whose solution $\hat{x}$ is at least very near the solution x_* of the constrained problem, or to finding a nonlinear system of equations whose simultaneous solution is the same x_*. Finally, a large percentage of the problems that we have met in practice are either unconstrained or else constrained in a very trivial way—for example, every component of x might have to be nonnegative.

The third problem that we consider is also a special case of unconstrained minimization, but owing to its importance and its special structure it is a research area all by itself. This is the *nonlinear least-squares* problem:

$$\text{Given} \quad R : \mathbb{R}^n \longrightarrow \mathbb{R}^m, \qquad m \geq n,$$

$$\text{find} \quad x_* \in \mathbb{R}^n \quad \text{for which} \quad \sum_{i=1}^{m} (r_i(x))^2 \text{ is minimized,} \tag{1.1.5}$$

where $r_i(x)$ denotes the ith component function of R. Problem (1.1.5) is most frequently met within the context of curve fitting, but it can arise whenever a nonlinear system has more nonlinear requirements than degrees of freedom.

We are concerned exclusively with the very common case when the nonlinear functions F, f, or R are at least once, twice, or twice continuously differentiable, respectively. We do not necessarily assume that the derivatives are analytically available, only that the functions are sufficiently smooth. For further comments on the typical size and other characteristics of nonlinear problems being solved today, see Section 1.2.

The typical scenario in the numerical solution of a nonlinear problem is that the user is asked to provide a subroutine to evaluate the problem function(s), and a starting point x_0 that is a crude approximation to the solution x_*. If they are readily available, the user is asked to provide first and perhaps second derivatives. Our emphasis in this book is on the most common difficulties encountered in solving problems in this framework: (1) what to do if the starting guess x_0 is not close to the solution x_* ("global method") and how to combine this effectively with a method that is used in the vicinity of the answer ("local method"); (2) what to do if analytic derivatives are not available; and (3) the construction of algorithms that will be efficient if evaluation of the problem function(s) is expensive. (It often is, sometimes dramatically so.) We discuss the basic methods and supply details of the algorithms that are currently considered the best ones for solving such problems. We also give the analysis that we believe is relevant to understanding these methods and extending or improving upon them in the future. In particular, we try to identify and emphasize the ideas and techniques that have evolved as the central ones in this field. We feel that the field has jelled to a point where these techniques are identifiable, and while some improvement is still likely, one no longer expects new algorithms to result in quantum jumps over the best being used today.

The techniques for solving the nonlinear equations and unconstrained minimization problems are closely related. Most of the book is concerned with these two problems. The nonlinear least-squares problem is just a special case of unconstrained minimization, but one can modify unconstrained minimization techniques to take special advantage of the structure of the nonlinear least-squares problem and produce better algorithms for it. Thus Chapter 10 is really an extensive worked-out example that illustrates how to apply and extend the preceding portion of the book.

One problem that we do not address in this book is finding the "global minimizer" of a nonlinear functional—that is, the absolute lowest point of $f(x)$ in the case when there are many distinct local minimizers, solutions to (1.1.2) in open connected regions of $\mathbb{R}^n$. This is a very difficult problem that is not nearly as extensively studied or as successfully solved as the problems we consider; two collections of papers on the subject are Dixon and Szegö (1975, 1978). Throughout this book we will use the word "global," as in "global method" or "globally convergent algorithm" to denote a method that is designed to converge to a *local* minimizer of a nonlinear functional or *some* solution of a system of nonlinear equations, *from almost any starting point*. It might be appropriate to call such methods *local* or *locally convergent*, but these descriptions are already reserved by tradition for another usage. Any method that is guaranteed to converge from every starting point is probably too inefficient for general use [see Allgower and Georg (1980)].

1.2 CHARACTERISTICS OF "REAL-WORLD" PROBLEMS

In this section we attempt to provide some feeling for nonlinear problems encountered in practice. First we give three real examples of nonlinear problems and some considerations involved in setting them up as numerical problems. Then we make some remarks on the size, expense, and other characteristics of nonlinear problems encountered in general.

One difficulty with discussing sample problems is that the background and algebraic description of problems in this field is rarely simple. Although this makes consulting work interesting, it is of no help in the introductory chapter of a numerical analysis book. Therefore we will simplify our examples when possible.

The simplest nonlinear problems are those in one variable. For example, a scientist may wish to determine the molecular configuration of a certain compound. The researcher derives an equation $f(x)$ giving the potential energy of a possible configuration as a function of the tangent x of the angle between its two components. Then, since nature will cause the molecule to assume the configuration with the minimum potential energy, it is desirable to find the x for which $f(x)$ is minimized. This is a minimization problem in the single variable x. It is likely to be highly nonlinear, owing to the physics of the function f. It truly is unconstrained, since x can take any real value. Since the