NEWTON'S METHOD

by

Jorge J. Moré and D. C. Sorensen

ARGONNE NATIONAL LABORATORY

ARGONNE NATIONAL LABORATORY, ARGONNE, ILLINOIS

Prepared for the U. S. DEPARTMENT OF ENERGY
under Contract W-31-109-Eng-38

DISCLAIMER

Printed in the United States of America
Available from
National Technical Information Service
U. S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

ANL--82-8

DE82 013692

Distribution Category:
Mathematics and Computers
(UC-32)

ANL-82-8

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

NEWTON'S METHOD

Jorge J. Moré and D. C. Sorensen

Applied Mathematics Division

February 1982

ABSTRACT

Newton's method plays a central role in the development of numerical techniques for optimization. In fact, most of the current practical methods for optimization can be viewed as variations on Newton's method. It is therefore important to understand Newton's method as an algorithm in its own right and as a key introduction to the most recent ideas in this area. One of the aims of this expository paper is to present and analyze two main approaches to Newton's method for unconstrained minimization: the line search approach and the trust region approach. The other aim is to present some of the recent developments in the optimization field which are related to Newton's method. In particular, we explore several variations on Newton's method which are appropriate for large scale problems, and we also show how quasi-Newton methods can be derived quite naturally from Newton's method.

Newton's Method

Jorge J. Moré and D. C. Sorensen

Applied Mathematics Division
Argonne National Laboratory
Argonne Illinois 60439

1. Introduction.

Many fundamental problems in science, engineering, and economics can be phrased in terms of minimizing a scalar valued function of several variables. Problems that arise in these practical settings usually have constraints placed upon the variables. Special techniques are required to handle these constraints but eventually the numerical techniques used must rely upon the efficient solution of unconstrained minimization problems.

Newton's method plays a central role in the development of numerical techniques for optimization. One of the reasons for its importance is that it arises very naturally from considering a Taylor approximation to the function. Because of its simplicity and wide applicability, Newton's method remains an important tool for solving many optimization problems. In fact, most of the current practical methods for optimization (e.g. quasi-Newton methods) can be viewed as variations on Newton's method. It is therefore important to understand Newton's method as an algorithm in its own right and as a key introduction to the most recent ideas in this area.

One of the aims of this paper is to present and analyze two main approaches to Newton's method for unconstrained minimization: the line search approach and the trust region approach. The other aim is to present some of the recent developments in the optimization field which are related to Newton's method. In particular, we explore several variations on Newton's method which are appropriate for large scale problems, and we also show how quasi-Newton methods can be derived quite naturally from Newton's method.

We assume familiarity with some of the basic notions from computational linear algebra (see, for example, Stewart [1973]), and the calculus of functions of several variables (see, for example, Chapter 7 of Bartle [1976] or Chapter 3 of

Work supported in part by the Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under Contract W-31-109-Eng-38.

Ortega and Rheinboldt [1970]), but otherwise the background necessary for this paper is minimal. We begin our development by reviewing some standard definitions and results.

Given a function $f : R^n \to R$ defined in an open set D, the unconstrained minimization problem is to find $x^* \in D$ such that

(1.1) $\quad f(x^*) \le f(x), \quad x \in N(x^*),$

for some open neighborhood $N(x^*)$ of the *local minimizer* x^*. If x^* is the only minimizer of f in $N(x^*)$ then x^* is an *isolated* minimizer of f. If $N(x^*)$ is all of D then x^* is a *global* minimizer of f in D.

The properties of local minimizers are better understood if we focus our attention on a reasonable class of functions. For our purposes, it is reasonable to assume that f is twice continuously differentiable. Under this assumption, the properties of local minimizers can be expressed in terms of the quadratic function

$$\psi(w) = \nabla f(x)^T w + \tfrac{1}{2} w^T \nabla^2 f(x) w$$

where $\nabla f(x)$ is the *gradient* of f at x and $\nabla^2 f(x)$ is the *Hessian matrix* of f at x. Recall that the i-th component of the gradient is $\partial_i f(x)$ and that the (i,j) element of the Hessian matrix is $\partial_{i,j} f(x)$. Since

(1.2) $\quad f(x + w) = f(x) + \psi(w) + o(\|w\|^2),$

the quadratic ψ is the *local quadratic model* at x of the possible reduction in f. Unless otherwise stated, in this paper $\|\cdot\|$ is the Euclidean norm on R^n, or the induced operator norm.

Theorem (1.3). Let $f : R^n \to R$ be twice continuously differentiable in an open set D. If $x^* \in D$ is a local minimizer of f then $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ is positive semidefinite. If $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ is positive definite for some $x^* \in D$, then x^* is an isolated local minimizer of f.

Proof: Let ψ be the local quadratic model at x^* of the possible reduction in f. If x^* is a local minimizer for f then (1.2) shows that

$$0 \le \psi(\alpha p) + o(\alpha^2) = \alpha \nabla f(x^*)^T p + \alpha^2 p^T \nabla^2 f(x^*) p + o(\alpha^2)$$

for each $p \in R^n$ and all α sufficiently small. This implies that $\nabla f(x^*)^T p = 0$ and that $p^T \nabla^2 f(x^*) p \ge 0$. Since p is arbitrary, we must conclude that $\nabla f(x^*) = 0$ and that $\nabla^2 f(x^*)$ is positive semidefinite. On the other hand, if $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ is positive definite then

$$\psi(w) = \tfrac{1}{2} w^T \nabla^2 f(x^*) w \ge \tfrac{1}{2} \lambda \|w\|^2,$$

where $\lambda > 0$ is the smallest eigenvalue of $\nabla^2 f(x^*)$. Now it follows from (1.2) that

x^* must be an isolated local minimizer for f. ∎

A point $x^* \in R^n$ such that $\nabla f(x^*) = 0$ is a *critical point* of f. Critical points can be divided into local minimizers, local maximizers, and *saddle points*. Theorem (1.3) shows, in particular, that if x^* is a critical point of f and $\nabla^2 f(x^*)$ is indefinite then x^* is a saddle point of f. If, however, $\nabla^2 f(x^*)$ is semidefinite and singular then Theorem (1.3) does not provide any information on the nature of the critical point. This gap between the necessary and sufficient conditions of Theorem (1.3) is illustrated by the 2-dimensional function

$$f(\xi_1,\xi_2) = \xi_1^3 + \xi_2^2.$$

Note that (0,0) is a critical point of f and that the Hessian matrix at (0,0) is positive semidefinite. However, (0,0) is a saddle point of f and not a local minimizer.

Algorithms for the unconstrained minimization of a function $f : R^n \to R$ are usually descent methods. Given an initial starting point x_0, a *descent method* generates a sequence of approximations $\{x_k\}$ to a local minimizer with the property that

$$(1.4) \qquad f(x_{k+1}) < f(x_k), \quad k \geq 0.$$

This descent condition alone is not sufficient to guarantee that the iterates $\{x_k\}$ approach a local minimizer. Stronger conditions are required to actually force the sequence into a neighborhood of a local minimizer. Once the iterates are in such a neighborhood, descent methods usually allow a rapidly convergent local method to determine the iterates. In this paper, the local method is Newton's iteration

$$x_{k+1} = x_k - \nabla^2 f(x_k)^{-1} \nabla f(x_k), \quad k \geq 0,$$

and our concern here is with modifications to this local method that will provide a general purpose algorithm.

An algorithm that is designed for general use should be analyzed as thoroughly as possible. The purpose of a convergence analysis is to predict the behavior of the sequences produced by the algorithm. This involves establishing properties of limit points and rates of convergence. These features, together with requirements of storage and computational effort, aid in the selection of an algorithm for a specific application. At the very least, we expect an unconstrained minimization algorithm to produce sequences which satisfy

$$(1.5) \qquad \lim_{k \to +\infty} \nabla f(x_k) = 0.$$

This condition guarantees that any limit point x^* of $\{x_k\}$ is a critical point of f. For algorithms which only use gradient information this is all that can be

expected. If an algorithm requires Hessian information, then it is reasonable to expect that the second order necessary conditions of Theorem (1.3) will be satisfied. This can be done by ensuring that

$$(1.6) \qquad \lim_{k \to +\infty} \inf \lambda_1(\nabla^2 f(x_k)) \geq 0,$$

where $\lambda_1(A)$ is the smallest eigenvalue of a symmetric matrix A. If (1.5) and (1.6) hold then any limit point x^* of $\{x_k\}$ satisfies the necessary conditions of Theorem (1.3).

In the remainder of this paper we shall derive Newton's method in its basic form and then introduce various modifications which have been devised to ensure that (1.4), (1.5), and (1.6) are satisfied by the sequences $\{x_k\}$ produced by the method. Techniques for forcing convergence from poor starting points is the subject of two sections. We discuss line search methods and trust region methods in Sections 4 and 5, respectively. Both approaches are important and can be applied to other optimization problems. Variations on Newton's method are discussed in Sections 6 and 7. Since the techniques for forcing strategies are all designed to bring the sequence into a neighborhood of a local minimizer and then switch automatically to Newton's method, it is most appropriate to begin, in Section 2, with a discussion of the unmodified local algorithm.

It will be worthwhile to have a specific problem in mind in order to appreciate some of the concerns we express with respect to implementation of the methods. The problem we consider is the simplest problem in the calculus of variations. An excellent introduction to this problem may be found in Fleming and Rishel [1975]. The problem is to minimize the functional

$$(1.7) \qquad J(u) \equiv \int_0^1 L(\tau, u, \dot{u})\, d\tau$$

over the set W of piecewise continuously differentiable functions u on the interval $[0, 1]$ with specified endpoints $u(0)$ and $u(1)$. We assume that L is twice continuously differentiable. Two classical problems of this form are the brachistochrone for which $L(\tau, u, v) = (u - \alpha)^{-1/2}(1 + v)^{1/2}$ for some constant α, and the minimal surface of revolution for which $L(\tau, u, v) = u\,(1 + v)^{1/2}$. An accessible introduction to the many applications related to the minimization of J may be found in Smith [1974].

The solution techniques available for minimizing J directly are very limited. In most practical settings one would almost surely need to resort to numerical techniques. One such technique is to discretize the continuous problem and then construct an approximate solution by solving the discrete problem. To see how this might be accomplished, consider a family of n-dimensional subspaces $W_n \subset W$. If $\{\varphi_k\}$ is a basis for W_n, then we can determine the minimum of J on

W_n by setting

$$f(x) = f(\xi_1, \xi_2, \cdots, \xi_n) \equiv J\left[\sum_{k=1}^{n} \xi_k \varphi_k\right],$$

and minimizing f. If Newton's method is used to minimize f, then we must be able to compute ∇f and $\nabla^2 f$. These derivatives can be computed by noting that the i-th component of ∇f is

$$\partial_i f(x) = \int_0^1 \left[(\partial_2 L)\varphi_i + (\partial_3 L)\dot{\varphi}_i\right] d\tau,$$

and that the (i,j) element of the Hessian $\nabla^2 f(x)$ is

$$\partial_{i,j} f(x) = \int_0^1 \left[(\partial_{2,2} L)\varphi_i\varphi_j + (\partial_{2,3} L)(\varphi_i\dot{\varphi}_j + \dot{\varphi}_i\varphi_j) + (\partial_{3,3} L)\dot{\varphi}_i\dot{\varphi}_j\right] d\tau,$$

where the partial derivatives of L are all evaluated at $(\tau, u(\tau), \dot{u}(\tau))$ and

$$(1.8) \qquad u(\tau) \equiv \sum_{k=1}^{n} \xi_k \varphi_k(\tau).$$

Once a solution x^* is found, the components of x^* can be used in (1.8) to construct an approximate minimizer u_n^* of J. Some analysis must be carried out to ensure that u_n^* is near a minimizer of J. An introduction to the type of analysis that is necessary may be found in Daniel [1971]. We shall only be concerned with the finite dimensional minimization process that occurs once n and a particular basis is selected.

In principle this is all that is required to apply Newton's method to f. However, some important practical considerations remain. First of all, a reasonable starting point x_0 may be difficult to provide so it becomes important to have an algorithm which converges from arbitrary starting points. Also, since the dimension of the approximating subspace W_n must increase in order to guarantee that u_n^* is close to a minimizer of J, it may be necessary to solve large dimensional problems. This means that the storage of the Hessian and that solutions of linear systems involving the Hessian matrix may become prohibitively costly. There is also the matter of evaluating the integrals to sufficient accuracy as well as the concern over how inaccuracies might affect the performance of the optimization algorithm. We shall take up some of these problems in this paper.

2. The Local Algorithm.

Newton's method can be studied from a local point of view in which we assume that the starting point x_0 is close to a local minimizer. This point of view is helpful because it provides information about the ultimate behavior of Newton's method. In Sections 4 and 5 we shall study Newton's method from a

global viewpoint.

Let $f : R^n \to R$ be a twice continuously differentiable function. Newton's method for the unconstrained minimization problem can be derived by assuming that we have an approximation x_k to a local minimizer of f and that in a neighborhood of x_k the approximation

$$f(x_k + w) \approx f(x_k) + \psi_k(w),$$

is appropriate, where

$$\psi_k(w) = \nabla f(x_k)^T w + \tfrac{1}{2} w^T \nabla^2 f(x_k) w$$

is the local quadratic model at x_k of the possible reduction in f. If this approximation is appropriate, then a presumably better approximation $x_{k+1} = x_k + s_k$ can be found by requiring that the step s_k be a minimizer of ψ_k. Theorem (1.3) shows that s_k must then satisfy

$$\nabla \psi_k(s_k) = \nabla f(x_k) + \nabla^2 f(x_k) s_k = 0.$$

Thus Newton's method takes an approximation x_0 and attempts to improve it through the iteration

$$x_{k+1} = x_k - \nabla^2 f(x_k)^{-1} \nabla f(x_k), \quad k \geq 0. \tag{2.1}$$

Note that in this derivation the only restriction on the step s_k is that it satisfy the system of linear equations $\nabla \psi_k(w) = 0$. In other words, we only require that s_k be a critical point of ψ_k. As a consequence, the Newton iteration (2.1) has the same behavior in the neighborhood of any critical point of f regardless of its type. This seems undesirable since we would like our algorithms to have a predilection towards local minimizers.

As it turns out, however, this behavior is just a consequence of the fact that iteration (2.1) is Newton's method for the solution of the system of nonlinear equations $\nabla f(x) = 0$. Since the local properties of iteration (2.1) only depend on the mapping $F(x) = \nabla f(x)$, let us consider Newton's method in this more general setting.

Let $F : R^n \to R^n$ be a mapping with range and domain in R^n and consider the problem of finding the solution to the system of n equations in n unknowns $F(x) = 0$, or equivalently,

$$f_i(\xi_1, \cdots, \xi_n) = 0, \quad 1 \leq i \leq n,$$

where f_i is the i-th component function of F. Newton's method for this problem can be derived by assuming that we have an approximation x_k to the solution of the system of nonlinear equations $F(x) = 0$, and that in a neighborhood of x_k the approximation

$$F(x_k + w) \approx L_k(w) \equiv F(x_k) + F'(x_k)w$$

is appropriate where $F'(x)$ is the *Jacobian matrix* of the mapping F at x. The next approximation $x_{k+1} = x_k + s_k$ can then be obtained by requiring that the step s_k satisfies the system of linear equations $L_k(w) = 0$. Thus Newton's method attempts to improve x_0 by the iteration

$$x_{k+1} = x_k - F'(x_k)^{-1}F(x_k), \quad k \geq 0. \tag{2.2}$$

In comparing iterations (2.1) and (2.2), note that (2.1) is a special case of the Newton iteration (2.2) applied to mapping $F(x) = \nabla f(x)$. Because of this relationship, it suffices to study the local behavior of iteration (2.2). The most important aspects of this local behavior are summarized in the following two theorems.

Theorem (2.3). Let $F: R^n \to R^n$ be a continuously differentiable mapping defined in an open set D, and assume that $F(x^*) = 0$ for some x^* in D and that $F'(x^*)$ is nonsingular. Then there is an open set S such that for any x_0 in S the Newton iterates (2.2) are well defined, remain in S, and converge to x^*.

Proof: Let α be a fixed constant in (0,1). Since F' is continuous at x^* and $F'(x^*)$ is nonsingular, there is an open ball $S \equiv \{x : \|x - x^*\| < \varepsilon\}$ and a positive constant μ such that

$$\|F'(x)^{-1}\| \leq \mu, \quad \|F'(y) - F'(x)\| \leq \frac{\alpha}{\mu},$$

for every x and y in S. Suppose that $x_k \in S$. Since x_{k+1} satisfies (2.2) and $F(x^*) = 0$ we have that

$$x_{k+1} - x^* = -F'(x_k)^{-1}\left[F(x_k) - F(x^*) - F'(x_k)(x_k - x^*)\right],$$

and hence

$$\|x_{k+1} - x^*\| \leq \mu \|F(x_k) - F(x^*) - F'(x_k)(x_k - x^*)\|.$$

Now, the fundamental theorem of integral calculus implies that

$$F(x_k) - F(x^*) - F'(x_k)(x_k - x^*) = \int_0^1 [F'(x^* + \xi(x_k - x^*)) - F'(x_k)](x_k - x^*)\,d\xi,$$

and hence,

$$\|x_{k+1} - x^*\| \leq \mu \left\{ \max_{0 \leq \xi \leq 1} \|F'(x^* + \xi(x_k - x^*)) - F'(x_k)\| \right\} \|x_k - x^*\|. \tag{2.4}$$

Thus,

$$\|x_{k+1} - x^*\| \leq \alpha \|x_k - x^*\|$$

as long as $x_k \in S$. Since $\alpha < 1$, this last inequality implies that if $x_0 \in S$ then $x_k \in S$ for $k = 1,2,\ldots$, and that $\{x_k\}$ converges to x^*. ∎

Theorem (2.3) states that Newton's method is *locally convergent* in the sense that if the starting point x_0 is sufficiently close to a solution x^* then Newton's method converges to x^*. Unfortunately, for many important problems the *domain of attraction* S guaranteed by Theorem (2.3) is quite small, and much research has gone into developing techniques to overcome this weakness of Newton's method. For systems of nonlinear equations this is a particularly active field, and much interest has been generated by the recent global Newton methods. For an account of some of this work see, for example, Keller [1978]. For the unconstrained minimization problem, the situation is in much better shape, and we examine two main approaches to globalizing Newton's method in Sections 4 and 5.

Although Theorem (2.3) is undeniably important, it does not tell the full story. It is not enough to know that a sequence converges if the rate is so slow that that we could not afford to see it converge. Generally, when analyzing an iterative method we are also interested in saying as much as possible about the expected rate of convergence of a sequence produced by the method. A reasonable optimization algorithm should be able to generate *linearly convergent* sequences $\{x_k\}$ in the sense that

$$\|x_{k+1} - x^*\| \leq \alpha \|x_k - x^*\|, \quad k \geq 0, \tag{2.5}$$

for some constant α in (0,1). If α is small then (2.5) is adequate, but if α is close to unity, say $\alpha \geq 0.9$, then (2.5) is not reassuring.

For many optimization algorithms which use second order information, it is possible to establish a stronger result than (2.5). A sequence $\{x_k\}$ converges *quadratically* to x^* if

$$\|x_{k+1} - x^*\| \leq \beta \|x_k - x^*\|^2, \quad k \geq 0, \tag{2.6}$$

for some constant $\beta > 0$. Since

$$\frac{\|x_{k+1} - x^*\|}{\|x^*\|} \leq \left(\beta \|x^*\|\right)\left(\frac{\|x_k - x^*\|}{\|x^*\|}\right)^2,$$

quadratic convergence implies that the number of significant digits of x_k as an approximation to x^* double at each iteration. Typically, as soon as two significant digits are obtained, the next three iterations will produce roughly sixteen significant digits.

There is a middle ground between (2.5) and (2.6). A sequence $\{x_k\}$ converges *superlinearly* to x^* if

$$\|x_{k+1} - x^*\| \leq \beta_k \|x_k - x^*\|, \quad k \geq 0, \tag{2.7}$$

for some sequence $\{\beta_k\}$ which converges to zero. It should be clear that a superlinearly convergent sequence is linearly convergent, and that a quadratically convergent sequence is superlinearly convergent. Also note that since

$$| \|x_{k+1}-x_k\| - \|x_k-x^*\| | \leq \|x_k-x^*\|,$$

it follows that

$$\lim_{k\to+\infty} \frac{\|x_{k+1}-x_k\|}{\|x_k-x^*\|} = 1,$$

when $\{x_k\}$ converges superlinearly to x^*. This is an important property because it implies that $\|x_{k+1}-x_k\|$ can be used to estimate the distance of x_k from x^*.

An iterative method is assigned a rate of convergence if it is possible to show that every convergent sequence produced will have at least this rate. Usually reasonable restrictions are imposed upon the domain of application for a method in order to obtain a useful assessment of this rate. Newton's method is usually quadratically convergent as we now demonstrate.

Theorem (2.8). Let $F: R^n \to R^n$ satisfy the assumptions of Theorem (2.3). Then the sequence $\{x_k\}$ produced by iteration (2.2) converges superlinearly to x^*. Moreover, if

(2.9) $$\|F'(x)-F'(x^*)\| \leq \kappa\|x-x^*\|, \quad x \in D,$$

for some constant $\kappa > 0$ then the sequence converges quadratically to x^*.

Proof: Convergence of the sequence $\{x_k\}$ was established in Theorem (2.3), so it only remains to establish the rate of convergence. To this end define

$$\beta_k \equiv \mu\left\{\max_{0\leq\xi\leq 1}\|F'(x^*+\xi(x_k-x^*)) - F'(x_k)\|\right\},$$

and assume that $x_0 \in S$ with μ and S defined in the proof of Theorem (2.3). The hypothesis on F' at x^* and the convergence of the sequence to x^* implies that $\{\beta_k\}$ converges to zero. Since inequality (2.4) shows that

$$\|x_{k+1}-x^*\| \leq \beta_k\|x_k-x^*\|,$$

this proves that $\{x_k\}$ converges superlinearly to x^*. Moreover, if (2.9) holds then

$$\beta_k \leq 2\mu\kappa\|x_k-x^*\|,$$

and hence $\{x_k\}$ converges quadratically to x^*. ∎

Note that the Lipschitz condition (2.9) is necessary in order to guarantee that Newton's method is quadratically convergent. For example, Newton's method applied to the 1-dimensional problem defined by

$$f(\xi) = \xi[1 + \log(|\xi|)], \quad \xi \neq 0, \quad f(0) = 0,$$

is precisely superlinearly convergent at $\xi^* = 0$; that is, if $\{\xi_k\}$ is a sequence generated by Newton's method, then the ratio

$$\frac{|\xi_{k+1}|}{|\xi_k|^{1+\rho}}$$

is unbounded for any $\rho > 0$.

Rate of convergence results are sometimes used to compare algorithms by claiming that the superior algorithm is the one with the highest rate of convergence. Claims of this type should be made with care because these results are asymptotic and thus it is usually not possible to establish the magnitude of the constants that appear in expressions like (2.5), (2.6) and (2.7). Moreover, rate of convergence results do not measure the work necessary to compute x_{k+1} from x_k, and in many cases this information is decisive in the choice of algorithm. For example, consider the class of quasi-Newton methods as described in Section 7. Sequences generated by these methods are known to be superlinearly convergent, and usually not quadratically convergent. However, since they do not require the computation of the Hessian matrix, quasi-Newton methods are often regarded as being superior to the quadratically convergent Newton methods.

3. Properties of Quadratic Functions.

Quadratic functions play an important role in the development of algorithms for optimization problems. For example, we have seen in Section 2 that in a neighborhood of a local minimizer of a function $f : R^n \to R$, Newton's method can be derived by requiring that the step be the minimizer of the local quadratic model

$$(3.1) \qquad \psi_k(w) = \nabla f(x_k)^T w + \tfrac{1}{2} w^T \nabla^2 f(x_k) w$$

of the expected reduction in f. It is therefore important to understand the properties of quadratic functions and to provide numerically stable algorithms for minimizing them.

Our first result completely describes the unconstrained minimization of quadratic functions.

Lemma (3.2). Let $\psi : R^n \to R$ be the quadratic function

$$(3.3) \qquad \psi(w) = g^T w + \tfrac{1}{2} w^T B w$$

where $g \in R^n$ and $B \in R^{n \times n}$ is a symmetric matrix.

a) *The quadratic ψ has a minimum if and only if B is positive semidefinite and g is in the range of B.*

b) *The quadratic ψ has a unique minimizer if and only if B is positive definite.*

c) *If B is positive semidefinite then every solution to the equation $Bp = -g$ is a global minimizer.*

Proof: Suppose B is positive semidefinite with g in the range of B. Then $Bp = -g$ has a solution, and thus

$$(3.4) \qquad \psi(p+w) = \psi(p) + (Bp + g)^T w + \tfrac{1}{2} w^T B w = \psi(p) + \tfrac{1}{2} w^T B w \geq \psi(p)$$

for every $w \in R^n$. On the other hand, if p is a minimizer of ψ then Theorem (1.3) implies that $Bp + g = \nabla\psi(p) = 0$, and that $B = \nabla^2\psi(p)$ is positive semidefinite. To establish b) and c) note that (3.4) holds whenever $Bp = -g$ and B is positive semidefinite, and that strict inequality holds for $w \neq 0$ if and only if B is positive definite. ▪

Given the quadratic ψ, there is an excellent numerical procedure for finding its minimizer. First an attempt is made at computing the Cholesky factorization of B. This factorization exists if and only if B is positive semidefinite, and in this case it leads to an upper triangular matrix R such that

$$B = R^T R .$$

If a negative diagonal is encountered during the factorization process, then B is not positive semidefinite and hence Lemma (3.2) shows that the quadratic ψ has no minimum. If the factorization is successful and R is nonsingular, then the minimizer is computed by solving the system $Bp = -g$, or equivalently,

$$R^T v = -g , \quad Rp = v .$$

If the factorization is successful but R is singular then B is positive semidefinite and singular. It still may be possible to compute a solution p, but from a numerical point of view, this computation is unstable because arbitrarily small perturbations can transform B into a positive definite matrix or an indefinite matrix.

Theorem (1.3) shows that in a neighborhood of a local minimizer of f, we can expect the Hessian matrix to be positive definite and then Lemma (3.2) shows that the local quadratic model (3.1) has a unique minimizer. Thus, in this case, the minimizer of the local quadratic model is a reasonable step for a minimization algorithm. However, away from a local minimizer the Hessian matrix $\nabla^2 f(x)$ may have a negative eigenvalue and then Lemma (3.2) tells us that this local quadratic model does not have a minimum. In fact, the model is not even bounded below. There are several remedies to this difficulty. One possibility is to modify the quadratic model by adding a positive semidefinite matrix

$E(x)$ so that

$$\nabla^2 f(x) + E(x)$$

is positive definite. When B is replaced by this matrix the step calculation can proceed as described above. An algorithm based upon this step calculation is described in Section 4.

Another possible remedy is to restrict the region in which we assume that the local quadratic model is appropriate. Locally the model still provides an excellent approximation to the expected reduction in f, so it is reasonable to restrict ψ to a ball $\{w : \|w\| \le \Delta\}$ for some $\Delta > 0$, and to compute a step as the minimizer of ψ on this ball. An algorithm based upon this step calculation is described in Section 5. The following result characterizes the solutions to the problem of minimizing a quadratic function on this restricted region.

Lemma (3.5). Let $\psi: R^n \to R$ be the quadratic function (3.3) and let $\Delta > 0$ be given. A point $p \in R^n$ solves the problem

(3.6) $\quad \min\{\psi(w) : \|w\| \le \Delta\}$

if and only if there is $\lambda \ge 0$ such that

(3.7) $\quad (B + \lambda I)p = -g\,, \quad \lambda(\Delta - \|p\|) = 0,$

with $B + \lambda I$ positive semidefinite.

Proof: Suppose that λ and p satisfy (3.7) with $B + \lambda I$ positive semidefinite. Then Lemma (3.2) implies that p minimizes the quadratic function

$$\hat{\psi}(w) = g^T w + \tfrac{1}{2} w^T (B + \lambda I) w\,.$$

Thus $\hat{\psi}(w) \ge \hat{\psi}(p)$ which implies that

(3.8) $\quad g^T w + \tfrac{1}{2} w^T B w \ge g^T p + \tfrac{1}{2} p^T B p + \frac{\lambda}{2}(p^T p - w^T w)$

for all $w \in R^n$. Since $\lambda p^T p = \lambda \Delta^2$ and $\lambda \ge 0$, it follows from (3.8) that $\psi(w) \ge \psi(p)$ whenever $\|w\| \le \Delta$, so p must solve (3.6).

Now suppose that p solves (3.6). If $\|p\| < \Delta$ then p is an unconstrained minimizer of ψ, so Lemma (3.2) implies that (3.7) holds with $\lambda = 0$, and that B is positive semidefinite. If $\|p\| = \Delta$ then p must also solve the equality constrained problem $\min\{\psi(w) : \|w\| = \Delta\}$. Therefore, the method of Lagrange ensures the existence of λ such that

$$\nabla L(p) = 0, \quad \text{where } L(w) \equiv \psi(w) + \frac{\lambda}{2}(w^T w - \Delta^2).$$

This implies that (3.7) holds for this λ and p. Moreover, since p solves (3.6) we have that (3.8) is valid for this λ and p whenever $\|w\| = \|p\|$. Using (3.7) to

replace g and then rearranging terms in (3.8) shows that

$$\tfrac{1}{2}(w-p)^T(B+\lambda I)(w-p) \geq 0$$

for every w with norm $\|p\|$. It follows readily from this inequality that $B + \lambda I$ is positive semidefinite. To show that $\lambda \geq 0$, note that Lemma (3.1) implies that (3.8) is valid for every $w \in R^n$. Now, if λ is not positive then (3.8) implies that $\psi(w) \geq \psi(p)$ whenever $\|w\| \geq \|p\|$. Since p solves (3.6) we must have that p is an unconstrained minimizer of ψ and then Lemma (3.1) implies that $\lambda = 0$. Hence, $\lambda \geq 0$ as claimed. ▪

Those familiar with various multiplier rules associated with mathematical programming will of course recognize that most of Lemma (3.5) could be obtained through direct application of these rules. At the very least one could invoke the Karush-Kuhn-Tucker conditions to avoid having to argue that the λ obtained from the Lagrange theory for equality constraints must be non-negative. Unfortunately, the multiplier theory for inequality constraints is not treated in most advanced calculus texts. McShane [1973] has an elementary treatment of the standard results, and Pourciau [1980] surveys the most recent results. It is interesting, however, that Lemma (3.5) cannot be obtained through direct application of the standard second order multiplier rules. The gap between necessary and sufficient second order conditions precludes this possibility since there is no such gap in this result.

Computing a numerical approximation to a solution of (3.6) requires some care. One immediate complication is that, due to the nonlinear constraint, there cannot be any general direct method for solving (3.6). In fact, when $g = 0$ a solution p to (3.6) must be an eigenvector of norm Δ corresponding to the smallest eigenvalue of B. Therefore,a general method for solving (3.6) must solve a symmetric eigenvalue problem in this special case.

The solution of (3.6) is straightforward if there are no solutions on the boundary of $\{w : \|w\| \leq \Delta\}$. In fact, it is not difficult to prove that (3.6) has no solution with $\|p\| = \Delta$ if and only if B is positive definite and $\|B^{-1}g\| < \Delta$.

If (3.6) has a solution on the boundary of $\{w : \|w\| \leq \Delta\}$, then Lemma (3.5) shows that it is reasonable to expect that the nonlinear equation

(3.9) $$\|p_\alpha\| = \Delta$$

where

$$p_\alpha = -(B + \alpha I)^{-1}g ,$$

has a solution $\lambda \geq 0$ in $(-\lambda_1, \infty)$ where λ_1 is the smallest eigenvalue of B. Note that (3.9) is a 1-dimensional zero finding problem in α that can be solved, for example, by Newton's method. However, since each evaluation of p_α requires the

solution of a system of linear equations, it is important to solve (3.9) with very few evaluations of p_α.

To solve (3.9), Reinsch [1967,1971] and Hebden [1973] observed independently that great advantage could be taken of the fact that the function $\|p_\alpha\|^2$ is a rational function in α with second order poles on a subset of the negatives of the eigenvalues of the symmetric matrix B. To see this consider the decomposition

$$B = Q\Lambda Q^T \text{ with } \Lambda = diag\,(\lambda_1, \lambda_2, \cdots, \lambda_n) \text{ and } Q^TQ = I,$$

and observe that

$$(3.10) \qquad \|p_\alpha\|^2 = \|Q(\Lambda + \alpha I)^{-1}Q^Tg\|^2 = \sum_{j=1}^{n} \frac{\gamma_j^2}{(\lambda_j + \alpha)^2}$$

where γ_i is the i-th component of Q^Tg. Knowledge of the functional form (3.10) shows that Newton's method may not be very efficient if it is applied to the function

$$\varphi_1(\alpha) = \|p_\alpha\| - \Delta.$$

A reason for this is that φ_1 has a pole at $-\lambda_1$, and thus Newton's method tends to perform poorly when the solution of (3.9) is near $-\lambda_1$. This point is clear from Figure 3.1 which shows a typical sketch of $\|p_\alpha\|$ with $\lambda_1 = -7.5$.

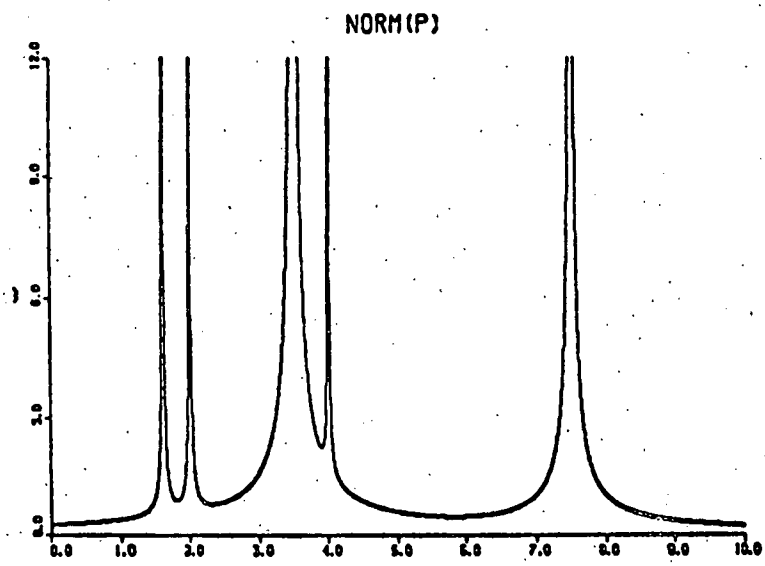

Figure 3.1

Reinsch and Hebden suggested that it is more efficient to apply Newton's method to the function

$$\varphi_2(\alpha) = \frac{1}{\Delta} - \frac{1}{\|p_\alpha\|}.$$

This function has no poles, and is almost linear near a solution of (3.9). This is illustrated quite well in Figure 3.2. This graph shows $(\|p_\alpha\|)^{-1}$ where $\|p_\alpha\|$ appears in Figure 3.1. Note that $(\|p_\alpha\|)^{-1}$ is almost linear for $\lambda > -\lambda_1$.

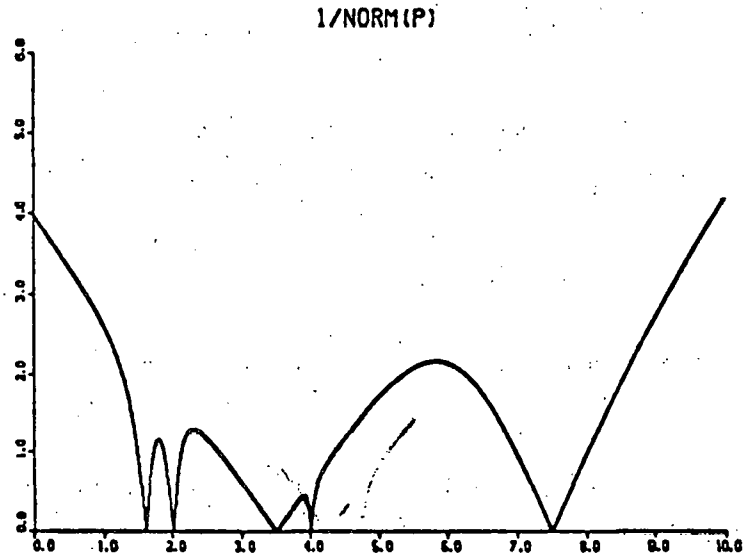

Figure 3.2

It is clear from Figure 3.2 that Newton's method is bound to perform well on φ_2. The Newton iteration applied to finding a zero of φ_2 takes the following form.

Algorithm (3.11).

1) Let λ_0 and $\Delta > 0$ be given.

2) For $k = 0,1,\ldots$ until "convergence"

a) Factor $B + \lambda_k I = R_k^T R_k$;

b) Solve $R_k^T R_k p_k = -g$;

c) Solve $R_k^T q_k = p_k$;

d) $\lambda_{k+1} = \lambda_k + \left(\frac{\|p_k\|}{\|q_k\|}\right)^2 \left(\frac{\|p_k\| - \Delta}{\Delta}\right)$;

If certain precautions are taken, then this basic iteration can be used to solve (3.6) in most cases. However, when B is indefinite there are cases in which the equation (3.9) has no solutions in $(-\lambda_1,\infty)$, and then Algorithm (3.11) fails. This happens, for example, when $g = 0$ and B is indefinite. It may also happen

when $g \neq 0$, as illustrated by the following simple example. If

$$B = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}, \quad g = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

then $\lambda_1 = -1$, and if $\alpha > 1$ then $\|p_\alpha\|^2 < \frac{1}{2}$. In our examples g is orthogonal to the eigenspace of B corresponding to the smallest eigenvalue. This is typical; g must be orthogonal to the eigenspace

$$S_1 \equiv \{z : Bz = \lambda_1 z, \ z \neq 0\}$$

corresponding to the smallest eigenvalue of B whenever (3.9) has no solutions in $(-\lambda_1, \infty)$. To see this it suffices to note that if g is not orthogonal to S_1, then $\gamma_1 \neq 0$ in (3.10), and hence

$$\lim_{\alpha \to -\lambda_1} \|p_\alpha\| = \infty, \quad \lim_{\alpha \to +\infty} \|p_\alpha\| = 0.$$

Since small perturbations of g lead to a nonzero γ_1, it is tempting to ignore the case when g is orthogonal to S_1. However, in many cases g is almost orthogonal to S_1, and in these cases an algorithm based completely on Newton's method would require a large number of iterations. This is not acceptable since a matrix factorization is required for each of these iterations.

Several algorithms have been proposed for the numerical solution of (3.6), but Gay [1981] was the first to show that his algorithm produced a nearly optimal solution. Gay's algorithm, however, may require a large number of iterations when g is orthogonal to S_1, and fails when $g = 0$ and B is indefinite. Moré and Sorensen [1981] have improved on Gay's algorithm, and their numerical results show that it is possible to produce a nearly optimal solution to (3.6) in all cases and with only a few iterations.

We have dealt with problem (3.6) at length because it arises in a variety of applications. For example, the solution of ill-posed problems in linear algebra usually requires the solution of (3.6) for a positive definite B. The literature on just this problem is extensive; for more information consult Eldén [1977], Gander [1978], and Varah [1979].

4. Line Search Methods.

In Section 2 we mentioned the difficulty of providing a starting point x_0 for Newton's method which is sufficiently close to a local minimizer. Overcoming this difficulty has been the subject of a considerable amount of recent research in numerical optimization, and in this section we discuss the line search approach to this problem. In discussing this approach it is best to consider general line search methods first, and then specialize to Newton's method.

Given an iterate x_k, the basic idea of a line search method is to compute a direction p_k and a parameter $\alpha_k > 0$ such that the next iterate $x_{k+1} = x_k + \alpha_k p_k$ has a lower function value. Convergence of the iterates to a minimizer depends on the choice of p_k and α_k.

A direction $p \in R^n$ is a *descent direction* for a function $f : R^n \to R$ at a point $x \in R^n$ if there is a constant $\bar{\alpha} > 0$ such that

(4.1) $\quad f(x + \alpha p) < f(x), \quad \alpha \in (0, \bar{\alpha}]$.

For differentiable functions, the easiest way to guarantee that (4.1) holds is to require that

(4.2) $\quad \nabla f(x)^T p < 0$.

In particular, the *steepest descent* choice $p = -\nabla f(x)$ satisfies (4.2). Condition (4.2) requires that the angle between $-\nabla f(x)$ and p be acute, and is equivalent to requiring that there is a positive definite matrix B such that

(4.3) $\quad p = -B^{-1} \nabla f(x)$.

This is not difficult to prove. If (4.3) holds then certainly (4.2) follows. Conversely, if (4.2) holds then

$$B = I - \frac{pp^T}{p^T p} - \frac{gg^T}{g^T p}, \quad g = \nabla f(x),$$

is positive definite and satisfies (4.3). Thus descent directions differ only in the choice of the positive definite matrix B in (4.3). The steepest descent method chooses B as the identity matrix and Newton's method chooses B as the Hessian matrix; the choice of B made in a quasi-Newton method (described in Section 7) is a compromise between these two choices.

Line search methods for differentiable functions assume that (4.2) holds. Note that if (4.2) does not hold then (4.1) can fail and then it may not be possible to make further reductions in f. Later on in this section we shall see that for convergence purposes it is necessary to require that p is not even nearly orthogonal to $\nabla f(x)$. This can be achieved by imposing a bound on the condition number of B in (4.3).

A *line search* algorithm examines points along the ray $\{x + \alpha p : \alpha \geq 0\}$ in search of a *steplength* α such that $f(x + \alpha p) < f(x)$. If p is a descent direction then such a point exists. In fact, the smallest positive local minimizer α^* of the univariate function

(4.4) $\quad \varphi(\alpha) \equiv f(x + \alpha p), \quad \alpha \geq 0$,

is such an α. However, it would not be practical to search for this point. Indeed, a line search algorithm is usually an iterative scheme for 1-dimensional

minimization, but the search process is usually terminated long before an accurate minimizer is found. Finding an accurate minimizer along a given ray usually does not yield a significantly larger reduction in f than a crude search, and better progress can often be made by making a reasonable reduction in the function f and then exploring other directions. These considerations have led to the development of stopping rules which terminate the line search process as soon as some minimal requirement is satisfied.

Given parameters $\mu \in (0, \frac{1}{2})$ and $\eta \in (\mu, 1)$, and a descent direction $p \in R^n$ which satisfies (4.2), the steplength $\alpha > 0$ belongs to $SR(\mu,\eta)$ if

$$f(x + \alpha p) \leq f(x) + \alpha\mu \nabla f(x)^T p ,$$

and

$$| \nabla f(x + \alpha p)^T p | \leq \eta | \nabla f(x)^T p | .$$

In other words, the set $SR(\mu,\eta)$ specifies the stopping rule. In terms of the function φ defined by (4.4), a steplength α belongs to $SR(\mu,\eta)$ if and only if

(4.5) $$\varphi(\alpha) \leq \varphi(0) + \alpha\mu\varphi'(0) ,$$

and

(4.6) $$| \varphi'(\alpha) | \leq \eta | \varphi'(0) | .$$

For a typical function φ, the set $SR(\mu,\eta)$ is shown in Figure 4.1.

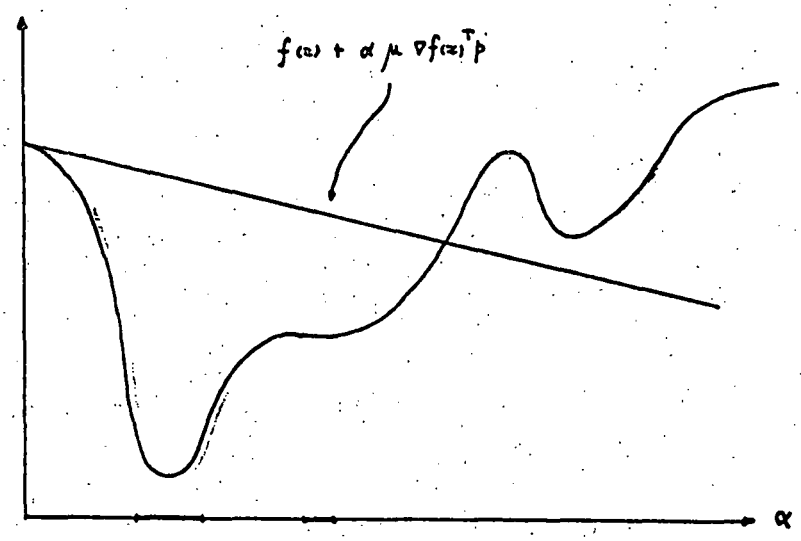

Figure 4.1

The intuitive nature of these rules should be clear. If α is not too small, then the first condition of the stopping rule SR(μ,η) forces a sufficient decrease in the function. However, (4.5) allows arbitrarily small choices of $\alpha > 0$, so this condition is not sufficient to guarantee convergence. The second condition rules out arbitrarily small choices of $\alpha > 0$ and usually implies that α is near a local minimizer of φ.

We assume that $\mu < \frac{1}{2}$ because if φ is a quadratic with $\varphi'(0) < 0$ and $\varphi''(0) > 0$ then the global minimizer α^* of φ satisfies

$$\varphi(\alpha^*) = \varphi(0) + \tfrac{1}{2}\alpha^*\varphi'(0),$$

and thus α^* satisfies (4.5) only if $\mu \leq \frac{1}{2}$. The restriction $\mu < \frac{1}{2}$ also allows $\alpha = 1$ to be ultimately acceptable to Newton and quasi-Newton methods; failure to take $\mu < \frac{1}{2}$ prevents these methods from converging superlinearly.

The restriction $\mu < \eta$ guarantees that under reasonable conditions SR(μ,η) contains a non-trivial interval. For example, suppose that φ is bounded below. Then there is a $\beta > 0$ such that

$$\varphi(\beta) \geq \varphi(0) + \mu\beta\varphi'(0). \tag{4.7}$$

Now let α^* be the smallest α in $(0,\beta]$ such that

$$\varphi(\alpha) = \varphi(0) + \mu\alpha\varphi'(0).$$

Then the mean value theorem shows that there is a τ such that

$$\varphi'(\tau) = \mu\varphi'(0) > \eta\varphi'(0), \quad 0 < \tau < \alpha^*.$$

In particular, since $\varphi'(0) < 0$ we must have that $\varphi'(\tau) < 0$. Hence, τ satisfies (4.6). Moreover, $\tau < \alpha^*$ implies that

$$\varphi(\tau) \leq \varphi(0) + \mu\tau\varphi'(0),$$

and thus τ also satisfies (4.5). Continuity of φ' now shows that SR(μ,η) contains a non-trivial interval.

The algorithms for selecting the steplength α are usually based upon minimizing a univariate quadratic or cubic model to φ defined by interpolation of function and first derivative at trial values of α. It is important to realize that it is possible to *safeguard* these algorithms so that they terminate in a finite number of steps.

Safeguarding a line search algorithm requires that we determine and update an *interval of uncertainty* I which contains points in SR(μ,η). The updating process must guarantee that the length of I tends to zero and that eventually I is contained in SR(μ,η).

To define the interval of uncertainty I, it is helpful to use an auxiliary function ψ defined by

$$\psi(\alpha) \equiv \varphi(\alpha) - \varphi(0) - \mu\alpha\varphi'(0),$$

and require that I be a closed interval with endpoints α_l and α_u such that

$$\psi'(\alpha_l)(\alpha_u - \alpha_l) < 0, \quad \psi(\alpha_l) \leq \psi(\alpha_u), \quad \psi(\alpha_l) \leq 0.$$

We now prove that $\psi'(\alpha^*) = 0$ and $\psi(\alpha^*) < 0$ for some α^* in I. As a consequence, if I is sufficiently small then

$$|\psi'(\alpha)| \leq (\eta - \mu)|\varphi'(0)|, \quad \psi(\alpha) \leq 0,$$

for all $\alpha \in I$, and thus I is contained in $SR(\mu,\eta)$. If we let α^* be a global minimizer of ψ on I then α^* cannot be an endpoint of I because this contradicts the above requirements on α_l and α_u. Hence, α^* is interior to I and thus $\psi'(\alpha^*) = 0$. Moreover, since $\psi(\alpha_l) \leq 0$ we must also have that $\psi(\alpha^*) < 0$.

We now show how to update I. Given a trial value α_t in I, we can determine a new interval I_+ with endpoints α_l^+ and α_u^+ as follows:

If $\psi(\alpha_t) \geq \psi(\alpha_l)$ then $\alpha_l^+ = \alpha_l$ and $\alpha_u^+ = \alpha_t$.

If $\psi(\alpha_t) < \psi(\alpha_l)$ and $\psi'(\alpha_t)(\alpha_t - \alpha_l) < 0$ then $\alpha_l^+ = \alpha_t$ and $\alpha_u^+ = \alpha_u$.

If $\psi(\alpha_t) < \psi(\alpha_l)$ and $\psi'(\alpha_t)(\alpha_t - \alpha_l) > 0$ then $\alpha_l^+ = \alpha_t$ and $\alpha_u^+ = \alpha_l$.

It is straightforward to show that α_l^+ and α_u^+ still define an interval of uncertainty unless $\psi'(\alpha_t) = 0$ and $\psi(\alpha_t) < \psi(\alpha_l)$. Of course, in this case α_t belongs to $SR(\mu,\eta)$ and there is no need to update I. Also note that these updating rules can be used to determine an initial interval of uncertainty. If we set $\alpha_l = 0$ then $\alpha_t > 0$ defines an interval of uncertainty if $\psi(\alpha_t) \geq 0$ or if $\psi'(\alpha_t) < 0$. For α_t sufficiently large, we must have that $\psi(\alpha_t) \geq 0$ unless φ is not bounded below.

There are many ways to compute the trial value of α_t; the only requirement on α_t is that the length of I tends to zero. This can be done by monitoring the length of I, and if say, the length of I is not reduced by a factor of 0.5 after two trials, then a bisection step can be used for the next trial α_t.

Theorem (4.8). Let $f : R^n \to R$ be continuously differentiable and bounded below on R^n, and assume that the starting point x_0 is such that ∇f is uniformly continuous on the level set

$$\Omega \equiv \{x \in R^n : f(x) \leq f(x_0)\}. \tag{4.9}$$

If the sequence $\{x_k\}$ is defined by $x_{k+1} = x_k + \alpha_k p_k$ where $\nabla f(x_k)^T p_k < 0$ and α_k is any steplength in $SR(\mu,\eta)$ then

$$\lim_{k \to +\infty} \left(\frac{\nabla f(x_k)^T p_k}{\|p_k\|} \right) = 0. \tag{4.10}$$

Proof: Since $\nabla f(x_k)^T p_k < 0$ and since f is bounded below, the sequence $\{x_k\}$ is well defined and lies in Ω. Moreover, $\{f(x_k)\}$ is decreasing and hence converges.

The proof is by contradiction. If (4.10) does not hold then there is an $\varepsilon > 0$ and a subsequence with index set K such that

$$-\frac{\nabla f(x_k)^T p_k}{\|p_k\|} \geq \varepsilon, \quad k \in K.$$

The first condition of the stopping rule SR(μ,η) shows that

$$f(x_k) - f(x_{k+1}) \geq \mu\alpha_k \|p_k\| \left[-\frac{\nabla f(x_k)^T p_k}{\|p_k\|}\right] \geq \mu\alpha_k \|p_k\| \varepsilon, \quad k \in K,$$

and since $\{f(x_k)\}$ is a convergent sequence, $\{\alpha_k p_k : k \in K\}$ converges to zero. Now, the second condition of the stopping rule SR(μ,η) yields the inequality

$$(1-\eta)(-\nabla f(x_k)^T p_k) \leq (\nabla f(x_k + \alpha_k p_k) - \nabla f(x_k))^T p_k, \quad k \geq 0,$$

and hence

$$\varepsilon \leq -\frac{\nabla f(x_k)^T p_k}{\|p_k\|} \leq \left[\frac{1}{1-\eta}\right] \|\nabla f(x_k + \alpha_k p_k) - \nabla f(x_k)\|, \quad k \in K.$$

However, since we have already shown that $\{\alpha_k p_k : k \in K\}$ converges to zero, this contradicts the uniform continuity of ∇f on Ω. ∎

Wolfe [1969] proved Theorem (4.8) under various choices of steplength rules, while Gill and Murray [1973] obtained a variation of Theorem (4.10) with the steplength chosen by a safeguarded algorithm designed for 1-dimensional minimization. Note, however, that 1-dimensional minimization algorithms must be modified in order to find points in SR(μ,η) because the first condition of this rule may exclude all 1-dimensional local minimizers. Also note that Theorem (4.8) holds under the hypothesis that $f : R^n \to R$ is continuously differentiable in an open set D and that the level set

$$\Omega = \{x \in D : f(x) \leq f(x_0)\} \tag{4.11}$$

is compact. The proof is almost identical to that of Theorem (4.8); the only difference occurs in proving that SR(μ,η) is not empty.

The specialization of Theorem (4.8) to algorithms of the Newton class is almost immediate. In this case,

$$p_k = -B_k^{-1} \nabla f(x_k), \tag{4.12}$$

where $\{B_k\}$ is a sequence of positive definite matrices with uniformly bounded condition numbers; that is, there is a constant $\kappa > 0$ such that

$$\|B_k\| \|B_k^{-1}\| \leq \kappa, \quad k \geq 0. \tag{4.13}$$

Under this assumption we have that

$$-\frac{\nabla f(x_k)^T p_k}{\|p_k\|} \geq \left[\frac{1}{\kappa}\right] \|\nabla f(x_k)\|,$$

and thus (4.10) implies that $\{\nabla f(x_k)\}$ converges to zero. In particular, every limit point of $\{x_k\}$ is a critical point of f.

For a line search method, this is the strongest type of result possible. It is not possible to prove that the limit points of $\{x_k\}$ are local minimizers because, for example, if x_0 is any critical point of f then a line search method terminates at x_0.

The choice of B_k in (4.12) is guided by a desire to satisfy (4.13) and still guarantee a fast rate of convergence. In the steepest descent method B_k is the identity matrix. For this method (4.13) is satisfied but convergence can be quite slow. The convergence of Newton's method is quite rapid when it occurs, but since $B_k = \nabla^2 f(x_k)$ is not necessarily positive definite, there is no guarantee of convergence. Modifications to Newton's method have been designed to overcome this problem. They set

$$B_k = \nabla^2 f(x_k) + E_k \tag{4.14}$$

where E_k is chosen so that B_k is positive definite and satisfies (4.13). There are many ways to do this, but one of the most effective methods is due to Gill and Murray [1974b].

Given a symmetric matrix A and parameters $\varepsilon \geq 0$ and $\beta > 0$, Gill and Murray's method produces an upper triangular matrix R and a diagonal matrix $E = \text{diag}(\varepsilon_i) \geq 0$ such that $A + E = R^T R$. The i-th step of the algorithm sets

$$\gamma_{ij} = a_{ij} - \sum_{k=1}^{i-1} r_{ki} r_{kj}, \quad i \leq j \leq n,$$

$$\mu_i = \max\{\, |\gamma_{ij}| : i < j \leq n \,\},$$

$$r_{ii} = \max\{\varepsilon, |\gamma_{ii}|^{1/2}, \frac{\mu_i}{\beta}\},$$

$$r_{ij} = \frac{\gamma_{ij}}{r_{ii}}, \quad i < j \leq n,$$

$$\varepsilon_i = r_{ii}^2 - \gamma_{ii}.$$

Note that if $\varepsilon = 0$ then it is possible that $r_{ii} = 0$, but in this case set $r_{ij} = 0$.

The idea behind the Gill and Murray algorithm is to increase the diagonal elements of A so that $A + E$ has a Cholesky decomposition. The increase in r_{ii} is designed to ensure that r_{ij} is bounded relative to $\|A\|$, and that if A is sufficiently positive definite then A is not modified. Note that the increase in r_{ii}

due to the term μ_i / β forces $| r_{ij} | \leq \beta$ for $i < j$. Hence,

$$| \gamma_{ij} | \leq | a_{ij} | + \beta^2 n , \quad i \leq j ,$$

This shows that μ_i / β is bounded in terms of β and $| a_{ij} |$. It is sensible to choose β so that this bound on μ_i / β is as small as possible, and this leads to a choice of

$$\beta^2 = \frac{1}{n} \max\{ | a_{ij} | : i \neq j \} .$$

This choice, however, may conflict with the desire to leave A unmodified whenever A is sufficiently positive definite. The definition of r_{ii} shows that in order to accomplish this β cannot be too small. It is sufficient to require that

$$\beta^2 \geq \max\{ | a_{ii} | : 1 \leq i \leq n \} .$$

To establish this claim we first show that if A is positive definite then $\gamma_{jj} > 0$. The proof is easy. Given an index j, define $p \in R^n$ by letting $Rp = r_{jj} e_j$. Then $p_j = 1$ and

$$0 < p^T A p = r_{jj}^2 - p^T E p \leq r_{jj}^2 - \varepsilon_j = \gamma_{jj} .$$

Now, since $\gamma_{jj} > 0$, it follows that

$$\beta^2 \geq a_{jj} > \sum_{k=1}^{j-1} r_{kj}^2 \geq r_{ij}^2 , \quad i < j ,$$

and since $\gamma_{ij} = r_{ij} r_{ii}$, we have that $\mu_i^2 < \beta^2 r_{ii}^2$. Hence, if A is positive definite then $r_{ii} = \max\{ \varepsilon , \gamma_{ii}^{1/2} \}$. This shows that if A is sufficiently positive definite then $r_{ii} = \gamma_{ii}^{1/2}$, and thus $E = 0$.

A reasonable way to guarantee that r_{ij} is bounded and that $E = 0$ whenever A is sufficiently positive definite, is to choose

$$\beta^2 = \max\left\{ \frac{1}{n} \max\{ | a_{ij} | : i \neq j \} , \max\{ | a_{ii} | : 1 \leq i \leq n \} \right\} .$$

For this choice of β it is not difficult to prove that

$$\varepsilon \leq | r_{ii} | \leq \max\{ \varepsilon , 2n\beta \} , \quad | r_{ij} | \leq \beta , \quad i < j .$$

For $\varepsilon > 0$, these inequalities show that if the Gill and Murray algorithm is applied to a bounded sequence $\{A_k\}$ of symmetric matrices, then (4.13) is satisfied.

We now have all the ingredients for a modified Newton's method with a line search. In this method, we compute p_k via (4.12) and (4.14), and determine E_k by Gill and Murray's modification of the Cholesky factorization with some $\varepsilon > 0$.

Theorem (4.15). Let $f : R^n \to R$ be twice continuously differentiable on an open set D, and assume that the starting point x_0 is such that the level set (4.11) is compact. If the sequence $\{x_k\}$ is defined by $x_{k+1} = x_k + \alpha_k p_k$ where p_k is

computed by the modified Newton's method and α_k is any steplength in SR(μ,η) then

$$\lim_{k\to+\infty} \nabla f(x_k) = 0.$$

Proof: We have already noted that if $\varepsilon > 0$ and $\{A_k\}$ is bounded then (4.13) holds. In this case $A_k = \nabla^2 f(x_k)$, and since Ω is compact, $\{A_k\}$ is bounded. Thus our result is a consequence of Theorem (4.8). ▪

There is an interesting variation of Theorem (4.15) which shows that the iterates $\{x_k\}$ usually converge. In this variation the stopping rule SR(μ,η) is modified by the addition of an upper bound β on the steplength. We accept β as the steplength only if

$$f(x + \beta p) \geq f(x) + \beta\mu \nabla f(x)^T p .$$

Note that if this condition is not satisfied then (4.7) holds, and then there is an α in $(0,\beta)$ which satisfies SR(μ,η). It is not difficult to show that Theorem (4.15) holds with this variation, and moreover, that

$$\lim_{k\to+\infty} \| x_{k+1} - x_k \| = 0. \tag{4.16}$$

This shows that if $\{x_k\}$ has an isolated limit point x^* then $\{x_k\}$ converges to x^*. In particular, note that if $\nabla^2 f(x^*)$ is nonsingular at a limit point x^*, then x^* is an isolated solution of $\nabla f(x) = 0$ and hence x^* is also an isolated limit point of $\{x_k\}$. The structure of the set of limit points of $\{x_k\}$ is further restricted by a result of Ostrowski [1966], page 203, which states that if $\{x_k\}$ is a bounded sequence and (4.16) holds, then the set of limit points of $\{x_k\}$ is connected.

To investigate the rate of convergence of the modified Newton's method, assume that the sequence $\{x_k\}$ converges to a point x^* at which $\nabla^2 f(x^*)$ is sufficiently positive definite in the sense that $E_k = 0$ for all k sufficiently large. Then

$$p_k = p_k^N \equiv -\nabla^2 f(x_k)^{-1} \nabla f(x_k), \tag{4.17}$$

and it can then be shown that there is a k_0 such that the steplength $\alpha_k = 1$ is in SR(μ,η) for $k \geq k_0$. With this choice of α_k, the rate of convergence is given by Theorem (2.8).

The above argument relies on the fact that $\alpha_k = 1$ is eventually in SR(μ,η). To establish this result it is only necessary to assume that $\{p_k\}$ tends to the Newton step in both length and direction; that is,

$$\lim_{k\to+\infty} \left(\frac{\| p_k - p_k^N \|}{\| p_k \|} \right) = 0,$$

where p_k^N is the Newton step (4.17). For a proof of this result, and a discussion of its relationship to quasi-Newton methods, see Dennis and Moré [1977].

5. Trust Region Methods.

In Newton's method with a line search the Hessian is modified when it is not sufficiently positive definite. This modification to the quadratic model guarantees convergence but seems to ignore the role of the quadratic model as a local approximation to the objective function. We now consider an alternative approach in which the quadratic model is not modified but instead, the quadratic model is only considered in a restricted *trust* region. We mentioned this technique briefly in Section 3 as motivation for Lemma (3.5); its use for globalizing Newton's method has resulted in reliable algorithms with strong convergence properties. In this section we introduce the main ideas of this approach and establish some of the basic convergence properties.

Let $f : R^n \to R$ be a twice continuously differentiable function. In Newton's method with a trust region strategy, each iterate x_k has a bound Δ_k such that

$$f(x_k + w) \approx f(x_k) + \psi_k(w), \quad \|w\| \leq \Delta_k \ ,$$

where

$$\psi_k(w) = \nabla f(x_k)^T w + \tfrac{1}{2} w^T \nabla^2 f(x_k) w \ .$$

is the quadratic model of the possible reduction in f within a neighborhood of the iterate x_k. This suggests that it may be desirable to compute a step s_k which approximately solves the problem

(5.1) $\quad \min\{\psi_k(w) : \|w\| \leq \Delta_k\}.$

If the step is satisfactory in the sense that $x_k + s_k$ produces a sufficient reduction in f, then Δ_k can be increased; if the step is unsatisfactory then Δ_k should be decreased. The following algorithm expresses these ideas in more detail.

Algorithm (5.2). Let $0 < \mu < \eta < 1$ and $0 < \gamma_1 < \gamma_2 < 1 < \gamma_3$ be specified constants.

1) Let $x_0 \in R^n$ and $\Delta_0 > 0$ be given.

2) For $k = 0,1,\ldots$ until "convergence"

a) Compute $\nabla f(x_k)$ and $\nabla^2 f(x_k)$.

b) Determine an approximate solution s_k to problem (5.1).

c) Compute $\rho_k = \dfrac{f(x_k + s_k) - f(x_k)}{\psi_k(s_k)}$.

d) If $\rho_k \leq \mu$ then $\Delta_k := \Delta \in [\gamma_1 \Delta_k, \gamma_2 \Delta_k]$ and go to b).

e) $x_{k+1} = x_k + s_k$.

f) If $\rho_k \leq \eta$ then $\Delta_{k+1} \in [\gamma_2 \Delta_k, \Delta_k]$ else $\Delta_{k+1} \in [\Delta_k, \gamma_3 \Delta_k]$.

This is a basic form of a trust region Newton's method. An interesting variant of this algorithm includes a scaling matrix for the variables. In this variation subproblem (5.1) is replaced by

$$\min\{\psi_k(w) : \| D_k w \| \leq \Delta_k\}$$

where D_k is a nonsingular matrix. We shall not discuss this generalization here; however, it is important to note that all of the results presented here hold for this variant if $\{D_k\}$ has uniformly bounded condition numbers. Such a modification can be very important in practice when the units of the variables are on widely different scales. Another variation is to use the hypercube $\{w : \| w \|_\infty \leq \Delta_k\}$ as the trust region in (5.1). In this variation subproblem (5.1) is replaced by the quadratic programming problem

$$(5.3) \qquad \min\{\psi_k(w) : | w^T e_i | \leq \Delta_k \ , \ 1 \leq i \leq n\}.$$

A difficulty with the hypercube approach is that it is quite expensive to compute a reasonable approximation to the global minimizer of (5.3). This is not the case with (5.1), and thus we shall only consider (5.1) in this section.

Just as in the case of a line search we are not interested in solving the model problem (5.1) with great accuracy. Instead, we are interested in providing relaxed conditions for accepting an approximate solution s_k to problem (5.1) which are sufficient to force convergence of the sequence $\{x_k\}$ generated by Algorithm (5.2). In fact, there are conditions which guarantee much more than convergence of the method. If ψ_k^* is the optimal value of (5.1), and if the approximate solution s_k to (5.1) satisfies

$$(5.4) \qquad -\psi_k(s_k) \geq \beta_1 |\psi_k^*| \quad \text{with} \quad \| s_k \| \leq \beta_2 \Delta_k.$$

for specified constants $\beta_1 > 0$ and $\beta_2 > 0$, then it is possible to prove that under suitable conditions on f, the sequence $\{x_k\}$ is convergent to a point x^* with $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ positive semidefinite.

It is not difficult to obtain a vector s_k which satisfies (5.4), although as mentioned at the end of Section 3, this requires attention to a number of details. Given σ in (0,1), the algorithm of Moré and Sorensen [1981], for example, finds a vector s_k such that

$$\psi_k(s_k) - \psi_k^* \leq \sigma(2-\sigma) \mid \psi_k^* \mid , \quad \| s_k \| \leq (1+\sigma)\Delta_k ,$$

provided $\psi_k^* \neq 0$. Of course, if $\psi_k^* = 0$, then $\nabla f(x_k) = 0$ and $\nabla^2 f(x_k)$ is positive semidefinite, so Algorithm (5.2) terminates at x_k. It is also worthy of mention, that if $\sigma = 0.1$ then the cost of this algorithm is quite reasonable. On the average the approximate solution of each model problem requires less than two factorizations of a symmetric positive definite matrix of order n.

Condition (5.4) can be expressed in an alternate form which is more convenient for proofs of convergence. If $p_k \in R^n$ is a solution to problem (5.1) then Lemma (3.5) implies that there is a parameter $\lambda_k \geq 0$ such that

$$(\nabla^2 f(x_k) + \lambda_k I)p_k = -\nabla f(x_k), \quad \lambda_k(\Delta_k - \| p_k \|) = 0.$$

Now let $R_k^T R_k$ be the Cholesky factorization of $\nabla^2 f(x_k) + \lambda_k I$. Then

$$(5.5) \qquad \mid \psi_k^* \mid = \tfrac{1}{2}(\| R_k p_k \|^2 + \lambda_k \Delta_k^2).$$

This expression for ψ_k^* shows that if (5.4) holds then

$$(5.6) \qquad -\psi_k(s_k) \geq \tfrac{1}{2}\beta_1(\| R_k p_k \|^2 + \lambda_k \Delta_k^2),$$

and thus the iterates $\{x_k\}$ generated by Algorithm (5.2) satisfy

$$(5.7) \qquad f(x_k) - f(x_{k+1}) \geq \tfrac{1}{2}\mu\beta_1(\| R_k p_k \|^2 + \lambda_k \Delta_k^2).$$

These two inequalities are essential to the proof of our next result.

Theorem (5.8). Let $f : R^n \to R$ be twice continuously differentiable on an open set D, and assume that the starting point x_0 is such that the level set

$$\Omega = \{ x \in D : f(x) \leq f(x_0) \}$$

is compact. If the sequence $\{x_k\}$ is produced by Algorithm (5.2) where s_k satisfies (5.4), then either the algorithm terminates at $x_l \in \Omega$ because $\nabla f(x_l) = 0$ and $\nabla^2 f(x_l)$ is positive semidefinite , or $\{x_k\}$ has a limit point x^* in Ω with $\nabla f(x^*) = 0$ and $\nabla^2 f(x^*)$ positive semidefinite.

Proof: If $\nabla f(x_l) = 0$ and $\nabla^2 f(x_l)$ is positive semidefinite for some iterate x_l in Ω then the algorithm terminates; otherwise (5.4) implies that $\psi_k(s_k) < 0$ for $k \geq 0$ and thus $\{x_k\}$ is well defined and lies in Ω.

Let us now prove the result under the assumption that $\{\lambda_k\}$ is not bounded away from zero. If some subsequence of $\{\lambda_k\}$ converges to zero then since Ω is compact we can assume, without loss of generality, that the same subsequence of $\{x_k\}$ converges to some x^* in the level set Ω. Since $\nabla^2 f(x_k) + \lambda_k I$ is positive semidefinite, $\nabla^2 f(x^*)$ is also positive semidefinite , and $\nabla f(x^*) = 0$ follows by noting that

$$\| R_k p_k \|^2 \geq \frac{\| \nabla f(x_k) \|^2}{\| \nabla^2 f(x_k) \| + \lambda_k}$$

and that (5.7) implies that $\{\| R_k p_k \|\}$ converges to zero.

We can show that $\{\lambda_k\}$ is not bounded away from zero by contradiction. If $\lambda_k \geq \varepsilon > 0$ then (5.4) and (5.6) yield that

$$-\psi_k(s_k) \geq \tfrac{1}{2}\beta_1 \lambda_k \Delta_k^2 \geq \tfrac{1}{2}\left(\frac{\beta_1}{\beta_2^2}\right)\varepsilon \| s_k \|^2.$$

Now, a standard estimate is that

$$(5.9) \qquad | f(x_k + s_k) - f(x_k) - \psi_k(s_k) | \leq \tfrac{1}{2} \| s_k \|^2 \max_{0 \leq \xi \leq 1} \| \nabla^2 f(x_k + \xi s_k) - \nabla^2 f(x_k) \|,$$

and thus the last two inequalities show that

$$(5.10) \qquad \left| \rho_k - 1 \right| \leq \left(\frac{\beta_2^2}{\beta_1 \varepsilon}\right) \max_{0 \leq \xi \leq 1} \| \nabla^2 f(x_k + \xi s_k) - \nabla^2 f(x_k) \| .$$

Inequality (5.7) implies that $\{\Delta_k\}$ converges to zero and hence $\{\| s_k \|\}$ also converges to zero. Thus the uniform continuity of $\nabla^2 f$ on Ω together with (5.10) implies that $\rho_k > \eta$ for all k sufficiently large and then the updating rules for Δ_k yield that $\{\Delta_k\}$ is bounded away from zero. This is in contradiction of the fact that $\{\Delta_k\}$ converges to zero. ■

The result we have just established is only a sample of the available convergence results for Algorithm (5.2) under assumption (5.4) for s_k. This theorem extends results of Fletcher [1980] and Sorensen [1980] by admitting inexact solutions to the model problem (5.1). The following additional results are known.

a) The sequence $\{\nabla f(x_k)\}$ converges to zero.

b) If x^* is an isolated limit point of $\{x_k\}$ then $\nabla^2 f(x^*)$ is positive semidefinite.

c) If $\nabla^2 f(x^*)$ is nonsingular for some limit point x^* of $\{x_k\}$ then $\{x_k\}$ converges to x^*.

Thomas [1975] proved a), while Moré and Sorensen [1981] established b) and c) as extensions of results due to Sorensen [1980]. Of these results, b) is characteristic of the trust region approach, and is the only result that does not hold for Newton's method with a line search. This difference between the two approaches is of theoretical importance. From a practical viewpoint, however, it can be argued that a more important difference is that with a line search approach the search for a lower function value occurs in a 1-dimensional subspace, while with a trust region approach the search is not restricted to a lower dimensional subspace.

An additional result which is helpful in establishing rate of convergence results is that if $\{x_k\}$ converges to x^* and $\nabla^2 f(x^*)$ is positive definite, then the sequence $\{\Delta_k\}$ is bounded away from zero. To prove this first note that if $\varepsilon_0 > 0$ is a lower bound on the eigenvalues of $\nabla^2 f(x_k)$ then (5.5) shows that

$$|\psi_k^*| \geq \tfrac{1}{2}\varepsilon_0 \min\{\Delta_k^2, \|s_k^N\|^2\}$$

where

(5.11) $$s_k^N \equiv -\nabla^2 f(x_k)^{-1} \nabla f(x_k).$$

Now, since $\psi_k(s_k) \leq 0$, we have that

$$\tfrac{1}{2}\|s_k\| \leq \|\nabla^2 f(x_k)^{-1}\| \, \|\nabla f(x_k)\|,$$

and thus $\tfrac{1}{2}\|s_k\| \leq \kappa \|s_k^N\|$ where κ is an upper bound on the condition number of $\nabla^2 f(x_k)$. Hence, assumption (5.4) shows that there is a constant $\varepsilon_1 > 0$ with

$$-\psi_k(s_k) \geq \tfrac{1}{2}\varepsilon_1 \|s_k\|^2.$$

This estimate and (5.9) then yield that

$$\left| \rho_k - 1 \right| \leq \left(\frac{1}{\varepsilon_1} \right) \max_{0 \leq \xi \leq 1} \|\nabla^2 f(x_k + \xi s_k) - \nabla^2 f(x_k)\|,$$

and thus $\rho_k > \eta$ for all k sufficiently large. It follows that $\{\Delta_k\}$ is bounded away from zero as desired.

Rate of convergence results can be obtained with the additional - but mild - assumption that there is a constant $\beta_3 > 0$ such that if $\|s_k\| \leq \beta_3 \Delta_k$ then $\nabla^2 f(x_k)$ is positive definite and $s_k = s_k^N$ where s_k^N is the Newton step (5.11). With this assumption in mind, suppose that $\{x_k\}$ converges to x^* and that $\nabla^2 f(x^*)$ is positive definite. Then $\{s_k\}$ converges to zero, and hence $\|s_k\| \leq \beta_3 \Delta_k$ for all k sufficiently large. Thus there is a $k_0 \geq 0$ such that $s_k = s_k^N$ for $k \geq k_0$, and then the rate of convergence results are provided by Theorem 2.8.

6. Approximations to the Hessian Matrix.

The methods we have described in the previous sections all require the computation of the Hessian matrix. This can be a difficult and error prone task, and in some cases analytic expressions for the entries of the Hessian matrix may not even be available. What can be done in these cases?

An obvious way to overcome these difficulties is to approximate the Hessian matrix with differences of gradients. However, there are several things to consider. Which difference approximation should be used ? How large should the difference parameter be ? How is the performance of the minimization method affected when difference approximations are used ?

The two most common type of difference approximations use the forward difference and central difference formulas. The forward difference approximation is based on the Taylor's series expansion

(6.1) $$\left(\frac{1}{\alpha}\right)[\nabla f(x + \alpha p) - \nabla f(x)] = \nabla^2 f(x) p + O(\alpha).$$

while the central difference approximation is based on

$$(\frac{1}{2\alpha})[\nabla f(x+\alpha p) - 2\nabla f(x) + \nabla f(x-\alpha p)] = \nabla^2 f(x)p + O(\alpha^2).$$

In optimization work, forward differences are quite common because they require fewer gradient evaluations and usually provide the necessary accuracy. If forward differences are used, an approximation $A(x)$ to the Hessian matrix at some $x \in R^n$ can be obtained by setting

$$A(x)e_j = (\frac{1}{\alpha_j})[\nabla f(x+\alpha e_j) - \nabla f(x)], \quad 1 \leq j \leq n.$$

for some *difference parameter* $\alpha_j \neq 0$. Unfortunately, this approximation does not necessarily provide a symmetric matrix. This important feature of the Hessian can be obtained by using the symmetric matrix

$$\tfrac{1}{2}[A(x) + A(x)^T]$$

as the approximation to the Hessian matrix at x.

The choice of difference parameter presents a dilemma. In order to preserve the superlinear rate of convergence enjoyed by Newton's method it is necessary to force the difference parameter to zero. However, as the difference parameter α_j becomes small, the differences loose significance due to cancellation. To prevent this loss of significance, the difference parameter must stay above a certain threshold value. This dilemma can usually be resolved in practice because it is not necessary to provide a Hessian approximation of high accuracy. If the Hessian approximation has an accuracy comparable to the desired accuracy in the solution to the optimization problem, then convergence usually takes place at practically a quadratic rate. Less accurate Hessian approximations decrease the rate of convergence but do not prevent convergence. These remarks assume that the gradient is evaluated accurately; if this is not the case, we may not even be able to compute a descent direction.

Techniques for choosing the difference parameter in (6.1) require information about ∇f in a neighborhood of x which is obtained by evaluating ∇f at several points near x. For many practical problems it would be too expensive to acquire this information at each iterate. A sensible strategy for an optimization algorithm is to choose the difference parameter at a typical x (possibly the starting point x_0), and to use this choice until it is deemed unsuitable. The difference parameter is only recomputed when the quality of the difference approximation starts to degrade.

There are several algorithms for choosing the difference parameter at a point. Discussing these algorithms in detail is not within the scope of this paper, but we want to mention some of the ideas behind these algorithms. The main

ideas are clear in the 1-dimensional case, so consider a differentiable function $\varphi: R \to R$, let $\varphi_c(\alpha)$ denote the computed value of $\varphi(\alpha)$, and let

$$\varepsilon(\alpha) = \varphi_c(\alpha) - \varphi(\alpha),$$

be the (absolute) error in the computed value. The smoothness of φ_c depends on the method used to evaluate φ on the computer, but in all cases φ_c is a step function. A reason for this is that a computer with l decimal digits of accuracy does not distinguish between numbers with the same first l digits. We mention this fact because it implies that φ_c is not differentiable. With these remarks in mind, note that our problem is to determine an α such that

$$(6.2) \qquad (\frac{1}{\alpha})[\varphi_c(\alpha) - \varphi_c(0)]$$

is close to $\varphi'(0)$. If we assume that we have an open neighborhood I of $\alpha = 0$, and a bound ε_0 such that

$$|\varepsilon(\alpha)| \leq \varepsilon_0, \quad \alpha \in I,$$

then it is not difficult to determine the difference parameter. Note that a Taylor's expansion of φ shows that

$$\varphi_c(\alpha) - \varphi_c(0) - \alpha\varphi'(0) = \tfrac{1}{2}\varphi''(\xi)\alpha^2 + [\varepsilon(\alpha) - \varepsilon(0)],$$

for some ξ with $|\xi| < |\alpha|$, and hence

$$|(\frac{1}{\alpha})[\varphi_c(\alpha) - \varphi_c(0)] - \varphi'(0)| \leq \tfrac{1}{2}\eta_0|\alpha| + \frac{2\varepsilon_0}{|\alpha|},$$

where η_0 is a bound for φ'' on I. This bound on the error between (6.2) and $\varphi'(0)$ has the correct qualitative behavior. If α is too small then the error is dominated by ε_0, while if α is too large then the error is determined by the curvature of φ. It is reasonable to choose α so that this bound is minimized, and this leads to a choice of

$$(6.3) \qquad \alpha = 2\left(\frac{\varepsilon_0}{\eta_0}\right)^{\frac{1}{2}}.$$

An algorithm for determining ε_0 and η_0 can be based on the work of Hamming [1971], pages 163-173. The basic idea is that the 4th and 5th order differences of φ_c are a measure of ε_0 and that the 3nd order differences can be used to estimate η_0. It is necessary to take some precautions, but in general we have found that an algorithm based on these ideas and (6.3) is quite effective.

Another way to overcome the difficulties mentioned at the beginning of this section is to approximate the Hessian matrix directly. As an illustration, recall that in the example of Section 1, the (i,j) element of the Hessian $\nabla^2 f(x)$ is

$$\theta_{i,j} f(x) = \int_0^1 \left[(\partial_{2,2} L)\varphi_i \varphi_j + (\partial_{2,3} L)(\varphi_i \dot{\varphi}_j + \dot{\varphi}_i \varphi_j) + (\partial_{3,3} L)\dot{\varphi}_i \dot{\varphi}_j \right] d\tau,$$

where the partial derivatives of L are all evaluated at $(\tau, u(\tau), \dot{u}(\tau))$ and

$$(6.4) \qquad u(\tau) \equiv \sum_{k=1}^{n} \xi_k \varphi_k(\tau).$$

In principle, these integrals can be evaluated with an appropriate quadrature, and the results used to define an approximation to the Hessian matrix. This requires the storage of a symmetric matrix of order n and the evaluation of $n(n+1)/2$ integrals over $[0,1]$. Since the dimension n must increase in order to refine the accuracy of u as an approximation to the continuous problem, it is usually necessary to solve large dimensional problems, and clearly, the cost of these requirements can then be prohibitive even for moderate values of n.

The above problems can be greatly reduced if we choose a basis $\{\varphi_i\}$ whose elements vanish on most of the interval $[0,1]$. For example, we could choose a B-spline basis. To illustrate this possibility, let $\tau_j = jh \equiv j/n$, and define

$$\varphi_j(\tau) = 1 - \frac{|\tau - \tau_j|}{h}, \quad \tau \in [\tau_{j-1}, \tau_{j+1}], \quad \varphi_j(\tau) = 0 \quad \text{otherwise}.$$

These functions are smooth B-splines of order 2. It is a simple matter to verify that

$$\varphi_i \varphi_j = \varphi_i \dot{\varphi}_j = \dot{\varphi}_i \dot{\varphi}_j \equiv 0, \quad |i-j| > 1,$$

and therefore the Hessian matrix is tridiagonal. Thus the storage is now of order n, and it is only necessary to evaluate $2n$ integrals over intervals of length $1/n$. Similar remarks hold for the computation of the gradient. If we had chosen a basis of smooth B-splines of order k then the Hessian has bandwidth $2k-1$. The computation of the gradient and Hessian is now more expensive, but there is an increase in the accuracy of (6.4) as an approximation to the continuous problem. For more information, see de Boor [1978] on splines, and Gill and Murray [1973] on the numerical solution of problems in the calculus of variations.

Large scale optimization problems frequently exhibit special structure such as sparsity in the Hessian matrix. Approximation of sparse Hessians by differences is attractive because the number of gradient differences required is often small compared to n. For example, if the Hessian matrix is tridiagonal (as in the above example), then 3 gradient differences suffice to approximate the Hessian.

A technique for estimating general sparse Hessian matrices is based on the work of Curtis, Powell, and Reid [1974]. They pointed out that a group of columns of $\nabla^2 f(x)$ can be approximated with one gradient difference if columns

in this group do not have a nonzero in the same row position. To see this, let I be the indices of a group of columns with this property, and let p be a vector with component $\rho_j = 0$ if j belongs to I and $\rho_j \neq 0$ otherwise. Then

$$\nabla^2 f(x)p = \sum_{j \in I} \rho_j \nabla^2 f(x) e_j ,$$

and since the columns with indices in I do not have nonzeroes in the same row position, for each (i,j) element of $\nabla^2 f(x)$ with $j \in I$ we have that

$$(\nabla^2 f(x)p)_i = \rho_j \, e_i^T \nabla^2 f(x) e_j .$$

In view of (6.1), it follows that we can approximate all the columns with indices in I with just one gradient difference.

For a tridiagonal matrix, it is easy to see that columns with indices of the form $l \ mod \ 3$ can be placed in the l-th group. Hence, as noted above, a tridiagonal matrix can be estimated with 3 gradient differences.

For general sparsity patterns it is not straightforward to partition the columns of the matrix into the least number of groups so that columns in a group do not have a nonzero in the same row position. Curtis, Powell, and Reid [1974] suggested an algorithm but did not analyze the problem. Coleman and Moré [1981] have approached this partitioning problem through its equivalence to certain graph coloring problems, and have used this point of view to analyze the partitioning problem and to suggest improved algorithms. Their numerical results show that these improved algorithms are nearly optimal on practical problems.

The partitioning technique that we have described for estimating sparse Hessians does not make any use of the symmetry of the matrix. Powell and Toint [1979] have pointed out that it is often possible to use symmetry to reduce the number of required gradient differences. They proposed several ways of doing this, and with one of their methods it is possible to estimate a tridiagonal Hessian matrix with 2 gradient differences. It turns out that their methods can also be analyzed with graph theory techniques; a treatment from this point of view is given by Coleman and Moré [1982].

7. Quasi-Newton Methods.

For some problems the objective function and its gradient are so expensive to calculate that we are not willing to compute a difference approximation to the Hessian matrix. These are not necessarily large dimensional problems. For example, the problem might be to minimize the L_2-norm of the solution to a differential equation that depends on a few parameters. In this case each function evaluation required by the optimization method actually involves a numerical solution of a differential equation.

In an effort to reduce the computational requirements of Newton's method, Davidon [1959] introduced a revolutionary idea which provides a way to approximate the Hessian matrix using only the gradient information gathered at each iterate. This idea has led to a highly successful class of methods which today are usually called quasi-Newton methods. There is a huge literature on quasi-Newton methods; our purpose in this section is to provide a brief introduction to the two most powerful members of this class and to contrast quasi-Newton methods with methods in the Newton class. For a thorough discussion of various aspects of quasi-Newton methods, see the survey paper of Dennis and Moré [1977].

In very simplistic terms a quasi-Newton method might be termed as an "earn while you learn" method. It is to be contrasted with methods in the Newton class through the manner of maintaining an approximate Hessian. In quasi-Newton methods, the approximate Hessian must satisfy the quasi-Newton equation. To derive this equation, suppose that we have a positive definite approximation B_k to the Hessian of f at x_k. We can then compute a descent direction p_k via

$$(7.1) \qquad p_k = -B_k^{-1}\nabla f(x_k),$$

and a steplength α_k in $SR(\mu,\eta)$. This defines the step $s_k = \alpha_k p_k$ and the next iterate $x_{k+1} = x_k + s_k$. Since

$$\left[\int_0^1 \nabla^2 f(x_k + \xi s_k)\, d\xi\right] s_k = \nabla f(x_{k+1}) - \nabla f(x_k),$$

it might be reasonable to seek an update to the approximate Hessian which satisfies

$$(7.2) \qquad B_{k+1}s_k = y_k \equiv \nabla f(x_{k+1}) - \nabla f(x_k).$$

This is the quasi-Newton equation and a method for generating B_{k+1} from B_k so that (7.2) holds is a quasi-Newton update. The quasi-Newton equation is essentially a gradient difference along the most recent search direction. Thus, quasi-Newton methods only use the search direction to obtain curvature information, while methods in the Newton class use n directions.

Various specific formulas exist for updating the matrix B_k, and we shall concentrate on those updates which guarantee that B_{k+1} is symmetric and positive definite. Note that if B_{k+1} is positive definite then (7.2) implies that $y_k^T s_k$ is positive. This condition is satisfied whenever α_k is in $SR(\mu,\eta)$ because then

$$y_k^T s_k = \nabla f(x_{k+1})s_k - \nabla f(x_k)s_k \geq (1-\eta)\,|\nabla f(x_k)^T s_k|.$$

We shall show below that if $y_k^T s_k$ is positive, then there are symmetric and positive definite matrices which satisfy the quasi-Newton equation.

In a discussion of quasi-Newton updates, it is customary and convenient to drop subscripts. Assume, therefore, that we have an approximate Hessian B, and vectors s and y with $y^T s$ positive. We then want to obtain update formulas that produce matrices B_+ according to the quasi-Newton equation

(7.3) $\quad B_+ s = y .$

The simplest derivation of such a formula is to ask for the nearest matrix to B which satisfies (7.3). If $B_+ = B + E$, then our problem is to find a solution to

(7.4) $\quad \min\{ \| E \| : (B + E)s = y \},$

where $\| \cdot \|$ is a suitable matrix norm. It is natural to choose the Frobenius norm defined by

$$\| A \|_F^2 = \sum_{i=1}^{n} \| Av_i \|^2 = \mathit{trace}\,(A^T A),$$

for any set $v_1, \cdots, v_n$ of orthonormal vectors, because this is the Euclidean norm in the space of matrices. With the Frobenius norm it is a simple matter to verify that

$$E = \frac{(y - Bs)s^T}{s^T s}$$

solves problem (7.4). Just note that E satisfies $(B + E)s = y$ and that if $\hat{E}$ is any other matrix that satisfies this equation then

$$\| E \|_F = \| \frac{(y - Bs)s^T}{s^T s} \|_F = \| \frac{\hat{E} s s^T}{s^T s} \|_F \leq \| \hat{E} \|_F \| \frac{s s^T}{s^T s} \|_F = \| \hat{E} \|_F .$$

This E is the unique solution to (7.4) since $\| \cdot \|_F$ is convex and the constraint in (7.4) is linear. The explicit updating formula for B is therefore given by

$$B_+ = B + \frac{(y - Bs)s^T}{s^T s}.$$

This is Broyden's [1965] rank-1 update formula; it is a rank-1 update because $\mathit{rank}\,(E) \leq 1$. Broyden's update is the most powerful quasi-Newton update for the solution of systems of nonlinear equations. For minimization, however, there are more suitable updates. A reason for this is that B_+ is usually neither symmetric nor positive definite even though B might possess these properties. In Section 4 we saw the importance of these properties in obtaining convergence for descent methods. Therefore, we are very interested in updating formulas which maintain symmetry and positive definiteness in the matrices B. One way to obtain these properties is to require that $B = R^T R$ and obtain $R_+ = R + E$ such that $B_+ = R_+^T R_+$ satisfies (7.3). It would be quite natural to seek the correction E as a solution to the problem

(7.5) $\quad \min\{\|E\|_F : (R+E)^T(R+E)s = y\}.$

Variational techniques can be used to show that the correction E which solves this problem is a rank-2 update to R. To our knowledge this approach has not been tried in practice because there is a rank-1 correction E which meets our requirements and has proven to be extremely successful. We might motivate a derivation of this rank-1 update by considering the implications of the quasi-Newton equation on the factors of B_+. If $B_+ = R_+^T R_+$ satisfies (7.3) then E must satisfy

(7.6) $\quad (R+E)^T v = y\ ,\ \ (R+E)s = v\ ,\ \ v^T v = y^T s > 0.$

Conversely, if E satisfies (7.6) for a given vector v, then $B_+ = R_+^T R_+$ satisfies (7.3). Thus a reasonable alternative to solving problem (7.5) is to specify a vector v of norm $(y^T s)^{1/2}$ and then obtain a correction matrix E which satisfies (7.6) as well as

(7.7) $\quad \min\{\|E\|_F : (R+E)^T v = y\}.$

The solution to (7.7) for a given v is

$$E = \frac{v(y - R^T v)^T}{v^T v},$$

and it follows that this E satisfies (7.6) if and only if $v = \tau Rs$ for some τ. Since we must have $v^T v = y^T s$, this condition determines τ. We have thus shown that

(7.8) $\quad R_+ = R + \dfrac{v(y - R^T v)^T}{v^T v},\ \ v = (y^T s)^{1/2}\dfrac{Rs}{\|Rs\|},$

induces a symmetric positive definite quasi-Newton update.

The updating formula we have just derived was discovered independently by Broyden[1969,1970], Fletcher[1970], Goldfarb[1970], Shanno[1970], and is often referred to as the BFGS formula. We have concentrated upon this particular update because it appears to work best in practice. The derivation of the BFGS update which we have presented is due to Dennis and Schnabel [1981]. Davidon and Sorensen [1980] have provided another derivation of the BFGS update by obtaining the quasi-Newton equation (7.6) as interpolation conditions on the gradient of the local quadratic model, and then showing that $v = \tau Rs$ is a consistent choice. However, while numerous derivations of this update formula have been given, the superior performance of the BFGS update has not yet received a satisfactory explanation.

A quasi-Newton method based on the BFGS update would generate a sequence $\{x_k\}$ defined by $x_{k+1} = x_k + p_k$ where p_k is the direction (7.1) and $\{B_k\}$ is chosen by the BFGS formula. If B_0 is symmetric and positive definite, and if α_k satisfies the stopping rules SR(μ,η), then $y_k^T s_k$ is positive and thus B_k is well

defined and positive definite for $k > 0$. For this method Powell [1975] has shown that if $f : R^n \to R$ satisfies the assumptions of Theorem (4.15), and if f is convex on the level set Ω, then $\{\nabla f(x_k)\}$ converges to zero. Moreover, if $\{x_k\}$ converges to x^* and $\nabla^2 f(x^*)$ is positive definite, then $\{x_k\}$ converges superlinearly to x^*. In practice this method converges for general functions f, so there is a wide gap between this result and what is observed in practice.

At first sight, it would seem that Powell's result is a rather straightforward extension of the analysis of Newton's method. However, this would be the case only if we could show that the condition numbers of $\{B_k\}$ are uniformly bounded. Interestingly enough, Powell shows this, but only after convergence has been established.

The form of the update (7.8) is quite amenable to stable numerical computation. In particular, it is possible to maintain the matrices R in triangular form. This facilitates the solution of the system (7.1) and reduces the storage. The reduction to triangular form can be accomplished by standard (see, for example, Gill, Golub, Murray, and Saunders [1974]) matrix updating techniques. If R is upper triangular, then a product of elementary rotations $Q = Q_1 Q_2 \cdots Q_{2n}$ can be constructed in such a way that

$$\hat{R}_+ = QR_+ = Q\left[R + \frac{v(y - R^T v)^T}{v^T v}\right] \tag{7.9}$$

is also upper triangular. Since Q is orthogonal,

$$\hat{R}_+^T \hat{R}_+ = R_+^T Q^T Q R_+ = R_+^T R_+ ,$$

and thus $\hat{R}_+$ is the required factor of B_+. Since Q is the product of $2n$ elementary rotations, the arithmetic required in (7.9) is on the order of n^2 floating point operations. This is to be compared to the order of n^3 operations required to form B_+ and then factor. Another advantage of keeping R in triangular form is that the condition number of triangular matrices can easily be monitored. This provides the opportunity to alter these matrices when extreme ill-conditioning occurs.

8. Current Research.

It is fairly safe to say that Newton's method and quasi-Newton methods are understood well enough to provide reliable software for general small to medium size unconstrained minimization problems. Several subroutines are available through software libraries and others are under development.

Currently, researchers are focusing much of their attention upon large scale problems. The ground rules for what constitutes an effective algorithm can change drastically when the number of variables becomes large. We have tacitly

assumed that the solution of a linear system of order n is, at worst, comparable in cost to the evaluation of the gradient and Hessian. This assumption may not be valid in large scale problems, and then it is necessary to take advantage of the special structure of the problem. With suitable modifications, Newton's method can still be an effective tool for large scale problems. We have already mentioned, in Section 6, one possible modification in connection with the estimation of sparse Hessian matrices by gradient differences. Modifications can also be made to the algorithms for determining the Newton direction. For example, since the Newton methods of this paper only require the Cholesky decomposition of a symmetric matrix, for sparse problems it is possible to reduce the amount of work and storage required by this decomposition. This is a well understood problem; see, for example, George and Liu [1979]. Another possibility is to only determine an approximation to the Newton direction. This possibility is explored, for example, by Dembo and Steihaug [1980].

So far modifications of quasi-Newton methods to account for sparsity have not had the resounding success that these methods have had in the dense case. This is despite intensive effort in this area. There may be fundamental reasons for this as noted by Sorensen [1981]. However, it would seem that this subject is just not fully understood at present, and thus this is still a very active research area. The interested reader should consult Steihaug [1980] and Toint [1981] for information and additional references.

The situation in greatest need of research at present arises when the Hessian matrix cannot be stored in fast memory. Currently the method of choice for this situation is a conjugate direction method. It would take a full article to describe these methods. Fletcher [1980] has a nice introduction to the basic ideas behind conjugate direction methods, while Buckley [1978a,1978b] and Gill and Murray [1979] describe some of the recent work in this area.

The development of methods for these very difficult and highly practical situations hinges upon a thorough understanding of Newton's method. It is our hope that this paper will provide a basis for work in these areas.

9. References

Bartle, R. G. [1976] The Elements of Real Analysis, John Wiley & Sons.

Broyden, C. G. [1965]. A class of methods for solving nonlinear simultaneous equations, Math. Comp. 19, 577-593.

Broyden, C. G. [1969]. A new double-rank minimization algorithm, Notices Amer. Math. Soc. 16, 670.

Broyden, C. G. [1970]. The convergence of single-rank quasi-Newton methods, Math. Comp. 24, 365-382.

Buckley, A. [1978a]. A combined conjugate-gradient quasi-Newton minimization algorithm, Math. Programming 15, 200-210.

Buckley, A. [1978b]. Extending the relationship between the conjugate gradient and BFGS algorithms, Math. Programming 15, 343-348.

Coleman, T. F. and Moré, J. J. [1981]. Estimation of sparse Jacobian matrices and graph coloring problems, Argonne National Laboratory, Report ANL-81-39, Argonne, Illinois.

Coleman, T. F. and Moré, J. J. [1982]. Estimation of sparse Hessian matrices and graph coloring problems, manuscript in preparation.

Curtis, A. R., Powell, M. J. D. and Reid, J. K. [1974]. On the estimation of sparse Jacobian matrices, J. Inst. Math. Appl. 13, 117-119.

Daniel, J. W. [1971]. The Approximate Minimization of Functionals, Prentice-Hall.

Davidon, W. C. [1959]. Variable metric method for minimization, Argonne National Laboratory, Report ANL-5990 (Rev.). Argonne, Illinois

Davidon, W. C. and Sorensen, D. C. [1980]. Numerical methods for unconstrained minimization, Proceedings COMPCON 80, IEEE Computer Society, Long Beach, California.

De Boor, C. [1978]. A Practical Guide to Splines, Springer-Verlag.

Dembo, R. S. and Steihaug, T. [1980]. Truncated Newton algorithms for large scale unconstrained optimization, Yale University, School of Organization and Management, Report 48, New Haven, Connecticut.

Dennis, J. E. and Moré, J. J. [1977]. Quasi-Newton methods, motivation and theory, SIAM Rev. 19, 46-89.

Dennis, J. E. and Schnabel, R. B. [1981]. A new derivation of symmetric positive definite secant updates, Nonlinear Programming 4, O. L. Mangasarian, R. R. Meyer, and S. M. Robinson, eds., Academic Press.

Eldén, L. [1977]. Algorithms for the regularization of ill-conditioned least squares problems, BIT 17, 134-145.

Fleming, W. H. and Rishel, R. W. [1975]. Deterministic and Stochastic Optimal Control, Springer-Verlag.

Fletcher, R. [1970]. A new approach to variable metric algorithms, Comput. J. 13, 317-322.

Fletcher, R. [1980]. Practical Methods of Optimization, Volume 1: Unconstrained Optimization, John Wiley & Sons.

Gander, W. [1978]. On the linear least squares problem with a quadratic constraint, Computer Science Department, Report 78-697, Stanford University, Stanford, California.

Gay, D. M. [1981]. Computing optimal locally constrained steps, SIAM J. Sci. Stat. Comput. 2, 186-197.

George, A. and Liu, J. [1981]. Computer Solution of Large Positive Definite Systems, Prentice-Hall.

Gill, P. E., Golub, G. H., Murray, W. and Saunders, M. A. [1974]. Methods for modifying matrix factorizations, Math. Comp. 28, 505-536.

Gill, P. E. and Murray, W. [1973]. The numerical solution of a problem in the calculus of variation, Recent Mathematical Developments in Control, D. J. Bell, ed., Academic Press.

Gill, P. E. and Murray, W. [1974a]. Safeguarded steplength algorithms for optimization using descent methods, National Physical Laboratory, Report NAC 37, Teddington, England.

Gill, P. E. and Murray, W. [1974b]. Newton-type methods for unconstrained and linearly constrained optimization, Math. Programming 7, 311-350.

Gill, P. E. and Murray W. [1979]. Conjugate-gradient methods for large-scale nonlinear optimization, Stanford University, Systems Optimization Laboratory, Report SOL 79-15, Stanford, California.

Goldfarb, D. [1970]. A family of variable-metric methods derived by variational means, Math. Comp. 24, 23-26.

Hamming, R. W. [1971]. Introduction to Applied Numerical Analysis, McGraw-Hill.

Hebden, M. D. [1973]. An algorithm for minimization using exact second derivatives, Atomic Energy Research Establishment, Report T.P. 515, Harwell, England.

Keller, H. B. [1978]. Global homotopies and Newton methods, Recent Advances in Numerical Analysis, C. de Boor and G. H. Golub, eds., Academic Press.

McShane, E. J. [1973]. The Lagrange multiplier rule, Amer. Math. Monthly 80, 922-925.

Moré, J. J. and Sorensen, D. C. [1981]. Computing a trust region step, Argonne National Laboratory, Report ANL-81-83, Argonne, Illinois.

Ortega, J. M. and Rheinboldt, W. C. [1970]. Iterative solution of nonlinear equations in several variables, Academic Press.

Ostrowski, A. [1966]. Solutions of Equations and Systems of Equations, Academic Press.

Pourciau, B. H. [1980]. Modern multiplier rules, Amer. Math. Monthly 87, 433-452.

Powell, M. J. D. [1975]. Convergence properties of a class of minimization algorithms, Nonlinear Programming 2, O. L. Mangasarian, R. R. Meyer, and S. M. Robinson, eds., Academic Press.

Powell, M. J. D. and Toint, Ph. L. [1979]. On the estimation of sparse Hessian matrices, SIAM J. Numer. Anal. 16, 1060-1074.

Reinsch, C. H. [1967]. Smoothing by spline functions, Numer. Math. 10, 177-183.

Reinsch, C. H. [1971]. Smoothing by spline functions II, Numer. Math. 16, 451-454.

Shanno, D. F. [1970]. Conditioning of quasi-Newton methods for function minimization, Math. Comp. 24, 647-656.

Sorensen, D. C. [1980]. Newton's method with a model trust-region modification, Argonne National Laboratory, Report ANL-80-106, Argonne, Illinois. SIAM J. Numer. Anal., to appear.

Sorensen, D. C. [1981]. An example concerning quasi-Newton estimation of a sparse Hessian, SIGNUM Newsletter 16, 8-10.

Smith, D. R. [1974]. Variational Methods in Optimization, Prentice-Hall.

Steihaug, T. [1980]. Quasi-Newton methods for large scale nonlinear problems, Yale University, School of Organization and Management, Report 49, New Haven, Connecticut.

Stewart, G. W. [1973]. Introduction to Matrix Computations, Academic Press.

Thomas, S. W. [1975]. Sequential estimation techniques for quasi-Newton algorithms, Ph. D. dissertation, Cornell University, Ithaca, New York.

Toint, Ph. L. [1981]. Towards an efficient sparsity exploiting Newton method for minimization, Sparse Matrices and their Uses, I. S. Duff, ed., Academic Press.

Varah, J. M. [1979]. A practical examination of some numerical methods for linear discrete ill-posed problems, SIAM Rev. 21, 100-111.

Wolfe, P. [1969]. Convergence conditions for ascent methods, SIAM Rev. 11, 226-235.

Distribution for ANL-82-8

Internal:

A. B. Krisciunas
P. C. Messina
J. J. More' (42)
D. M. Pahis
L. M. Phebus (2)
G. W. Pieper
R. J. Royston
D. C. Sorensen (41)
ANL Patent Department
ANL Contract File
ANL Libraries
TIS Files (6)

External:

DOE-TIC, for distribution per UC-32 (182)
Manager, Chicago Operations Office, DOE
President, Argonne Universities Association
Applied Mathematics Division Review Committee:
  J. C. Browne, U. Texas at Austin
  S. Gerhart, Software Research Associates
  W. C. Lynch, Xerox Corp.
  J. M. Ortega, U. Virginia
  S. Rosen, Purdue U.
  F. H. Westervelt, Wayne State U.
  M. F. Wheeler, Rice U.
D. Austin, BES-DOE