

APPENDIX

A

# Background Material

## **A.1 ELEMENTS OF LINEAR ALGEBRA**

### **VECTORS AND MATRICES**

In this book we work exclusively with vectors and matrices whose components are real numbers. Vectors are usually denoted by lowercase roman characters, and matrices by uppercase roman characters. The space of real vectors of length  $n$  is denoted by  $\mathbb{R}^n$ , while the space of real  $m \times n$  matrices is denoted by  $\mathbb{R}^{m \times n}$ .

Given a vector  $x \in \mathbb{R}^n$ , we use  $x_i$  to denote its  $i$ th component. We invariably assume that  $x$  is a *column* vector, that is,

$$x = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}.$$

The transpose of  $x$ , denoted by  $x^T$  is the row vector

$$x^T = [ x_1 \quad x_2 \quad \cdots \quad x_n ],$$

and is often also written with parentheses as  $x = (x_1, x_2, \dots, x_n)$ . We write  $x \geq 0$  to indicate componentwise nonnegativity, that is,  $x_i \geq 0$  for all  $i = 1, 2, \dots, n$ , while  $x > 0$  indicates that  $x_i > 0$  for all  $i = 1, 2, \dots, n$ .

Given  $x \in \mathbb{R}^n$  and  $y \in \mathbb{R}^n$ , the standard inner product is  $x^T y = \sum_{i=1}^n x_i y_i$ .

Given a matrix  $A \in \mathbb{R}^{m \times n}$ , we specify its components by double subscripts as  $A_{ij}$ , for  $i = 1, 2, \dots, m$  and  $j = 1, 2, \dots, n$ . The transpose of  $A$ , denoted by  $A^T$ , is the  $n \times m$  matrix whose components are  $A_{ji}$ . The matrix  $A$  is said to be *square* if  $m = n$ . A square matrix is *symmetric* if  $A = A^T$ .

A square matrix  $A$  is *positive definite* if there is a positive scalar  $\alpha$  such that

$$x^T A x \geq \alpha x^T x, \quad \text{for all } x \in \mathbb{R}^n. \quad (\text{A.1})$$

It is *positive semidefinite* if

$$x^T A x \geq 0, \quad \text{for all } x \in \mathbb{R}^n.$$

We can recognize that a symmetric matrix is positive definite by computing its eigenvalues and verifying that they are all positive, or by performing a Cholesky factorization. Both techniques are discussed further in later sections.

The diagonal of the matrix  $A \in \mathbb{R}^{m \times n}$  consists of the elements  $A_{ii}$ , for  $i = 1, 2, \dots, \min(m, n)$ . The matrix  $A \in \mathbb{R}^{m \times n}$  is *lower triangular* if  $A_{ij} = 0$  whenever  $i < j$ ; that is, all elements above the diagonal are zero. It is *upper triangular* if  $A_{ij} = 0$  whenever  $i > j$ ; that is, all elements below the diagonal are zero.  $A$  is *diagonal* if  $A_{ij} = 0$  whenever  $i \neq j$ .

The identity matrix, denoted by  $I$ , is the square diagonal matrix whose diagonal elements are all 1.

A square  $n \times n$  matrix  $A$  is *nonsingular* if for any vector  $b \in \mathbb{R}^n$ , there exists  $x \in \mathbb{R}^n$  such that  $Ax = b$ . For nonsingular matrices  $A$ , there exists a unique  $n \times n$  matrix  $B$  such that  $AB = BA = I$ . We denote  $B$  by  $A^{-1}$  and call it the *inverse* of  $A$ . It is not hard to show that the inverse of  $A^T$  is the transpose of  $A^{-1}$ .

A square matrix  $Q$  is *orthogonal* if it has the property that  $QQ^T = Q^T Q = I$ . In other words, the inverse of an orthogonal matrix is its transpose.

## NORMS

For a vector  $x \in \mathbb{R}^n$ , we define the following norms:

$$\|x\|_1 \stackrel{\text{def}}{=} \sum_{i=1}^n |x_i|, \quad (\text{A.2a})$$

$$\|x\|_2 \stackrel{\text{def}}{=} \left( \sum_{i=1}^n x_i^2 \right)^{1/2} = (x^T x)^{1/2}, \quad (\text{A.2b})$$

$$\|x\|_\infty \stackrel{\text{def}}{=} \max_{i=1, \dots, n} |x_i|. \quad (\text{A.2c})$$

The norm  $\|\cdot\|_2$  is often called the *Euclidean norm*. We sometimes refer to  $\|\cdot\|_1$  as the  $\ell_1$  norm and to  $\|\cdot\|_\infty$  as the  $\ell_\infty$  norm. All these norms measure the length of the vector in some sense, and they are equivalent in the sense that each one is bounded above and below by a multiple of the other. To be precise, we have for all  $x \in \mathbb{R}^n$  that

$$\|x\|_\infty \leq \|x\|_2 \leq \sqrt{n}\|x\|_\infty, \quad \|x\|_\infty \leq \|x\|_1 \leq n\|x\|_\infty, \quad (\text{A.3})$$

and so on. In general, a norm is any mapping  $\|\cdot\|$  from  $\mathbb{R}^n$  to the nonnegative real numbers that satisfies the following properties:

$$\|x + z\| \leq \|x\| + \|z\|, \quad \text{for all } x, z \in \mathbb{R}^n; \quad (\text{A.4a})$$

$$\|x\| = 0 \Rightarrow x = 0; \quad (\text{A.4b})$$

$$\|\alpha x\| = |\alpha| \|x\|, \quad \text{for all } \alpha \in \mathbb{R} \text{ and } x \in \mathbb{R}^n. \quad (\text{A.4c})$$

Equality holds in (A.4a) if and only if one of the vectors  $x$  and  $z$  is a nonnegative scalar multiple of the other.

Another interesting property that holds for the Euclidean norm  $\|\cdot\| = \|\cdot\|_2$  is the Cauchy–Schwarz inequality, which states that

$$|x^T z| \leq \|x\| \|z\|, \quad (\text{A.5})$$

with equality if and only if one of these vectors is a nonnegative multiple of the other. We can prove this result as follows:

$$0 \leq \|\alpha x + z\|^2 = \alpha^2 \|x\|^2 + 2\alpha x^T z + \|z\|^2.$$

The right-hand-side is a convex function of  $\alpha$ , and it satisfies the required nonnegativity property only if there exist fewer than 2 distinct real roots, that is,

$$(2x^T z)^2 \leq 4\|x\|^2 \|z\|^2,$$

proving (A.5). Equality occurs when the quadratic  $\alpha$  has exactly one real root (that is,  $|x^T z| = \|x\| \|z\|$ ) and when  $\alpha x + z = 0$  for some  $\alpha$ , as claimed.

Any norm  $\|\cdot\|$  has a *dual norm*  $\|\cdot\|_D$  defined by

$$\|x\|_D = \max_{\|y\|=1} x^T y. \tag{A.6}$$

It is easy to show that the norms  $\|\cdot\|_1$  and  $\|\cdot\|_\infty$  are duals of each other, and that the Euclidean norm is its own dual.

We can derive definitions for certain matrix norms from these vector norm definitions. If we let  $\|\cdot\|$  be generic notation for the three norms listed in (A.2), we define the corresponding matrix norm as

$$\|A\| \stackrel{\text{def}}{=} \sup_{x \neq 0} \frac{\|Ax\|}{\|x\|}. \tag{A.7}$$

The matrix norms defined in this way are said to be *consistent* with the vector norms (A.2). Explicit formulae for these norms are as follows:

$$\|A\|_1 = \max_{j=1,\dots,n} \sum_{i=1}^m |A_{ij}|, \tag{A.8a}$$

$$\|A\|_2 = \text{largest eigenvalue of } (A^T A)^{1/2}, \tag{A.8b}$$

$$\|A\|_\infty = \max_{i=1,\dots,m} \sum_{j=1}^n |A_{ij}|. \tag{A.8c}$$

The Frobenius norm  $\|A\|_F$  of the matrix  $A$  is defined by

$$\|A\|_F = \left( \sum_{i=1}^m \sum_{j=1}^n A_{ij}^2 \right)^{1/2}. \tag{A.9}$$

This norm is useful for many purposes, but it is not consistent with any vector norm. Once again, these various matrix norms are equivalent with each other in a sense similar to (A.3).

For the Euclidean norm  $\|\cdot\| = \|\cdot\|_2$ , the following property holds:

$$\|AB\| \leq \|A\| \|B\|, \tag{A.10}$$

for all matrices  $A$  and  $B$  with consistent dimensions.

The *condition number* of a nonsingular matrix is defined as

$$\kappa(A) = \|A\| \|A^{-1}\|, \tag{A.11}$$

where any matrix norm can be used in the definition. Different norms can be the use of a subscript— $\kappa_1(\cdot)$ ,  $\kappa_2(\cdot)$ , and  $\kappa_\infty(\cdot)$ , respectively—with  $\kappa$  denoting  $\kappa_2$  by default.

Norms also have a meaning for scalar, vector, and matrix-valued functions that are defined on a particular domain. In these cases, we can define Hilbert spaces of functions for which the inner product and norm are defined in terms of an integral over the domain. We omit details, since all the development of this book takes place in the space  $\mathbb{R}^n$ , though many of the algorithms can be extended to more general Hilbert spaces. However, we mention for purposes of the analysis of Newton-like methods that the following inequality holds for functions of the type that we consider in this book:

$$\left\| \int_a^b F(t) dt \right\| \leq \int_a^b \|F(t)\| dt, \quad (\text{A.12})$$

where  $F$  is a continuous scalar-, vector-, or matrix-valued function on the interval  $[a, b]$ .

### SUBSPACES

Given the Euclidean space  $\mathbb{R}^n$ , the subset  $\mathcal{S} \subset \mathbb{R}^n$  is a *subspace of  $\mathbb{R}^n$*  if the following property holds: If  $x$  and  $y$  are any two elements of  $\mathcal{S}$ , then

$$\alpha x + \beta y \in \mathcal{S}, \quad \text{for all } \alpha, \beta \in \mathbb{R}.$$

For instance,  $\mathcal{S}$  is a subspace of  $\mathbb{R}^2$  if it consists of (i) the whole space  $\mathbb{R}^2$ ; (ii) any line passing through the origin; (iii) the origin alone; or (iv) the empty set.

Given any set of vectors  $a_i \in \mathbb{R}^n$ ,  $i = 1, 2, \dots, m$ , the set

$$\mathcal{S} = \{w \in \mathbb{R}^n \mid a_i^T w = 0, i = 1, 2, \dots, m\} \quad (\text{A.13})$$

is a subspace. However, the set

$$\{w \in \mathbb{R}^n \mid a_i^T w \geq 0, i = 1, 2, \dots, m\} \quad (\text{A.14})$$

is not in general a subspace. For example, if we have  $n = 2$ ,  $m = 1$ , and  $a_1 = (1, 0)^T$ , this set would consist of all vectors  $(w_1, w_2)^T$  with  $w_1 \geq 0$ , but then given two vectors  $x = (1, 0)^T$  and  $y = (2, 3)$  in this set, it is easy to choose multiples  $\alpha$  and  $\beta$  such that  $\alpha x + \beta y$  has a negative first component, and so lies outside the set.

Sets of the forms (A.13) and (A.14) arise in the discussion of second-order optimality conditions for constrained optimization.

A set of vectors  $\{s_1, s_2, \dots, s_m\}$  in  $\mathbb{R}^n$  is called a *linearly independent set* if there are no real numbers  $\alpha_1, \alpha_2, \dots, \alpha_m$  such that

$$\alpha_1 s_1 + \alpha_2 s_2 + \dots + \alpha_m s_m = 0,$$

unless we make the trivial choice  $\alpha_1 = \alpha_2 = \cdots = \alpha_m = 0$ . Another way to define linear independence is to say that none of the vectors  $s_1, s_2, \dots, s_m$  can be written as a linear combination of the other vectors in this set. If in fact we have  $s_i \in \mathcal{S}$  for all  $i = 1, 2, \dots, m$ , we say that  $\{s_1, s_2, \dots, s_m\}$  is a *spanning set* for  $\mathcal{S}$  if *any* vector  $s \in \mathcal{S}$  can be written as

$$s = \alpha_1 s_1 + \alpha_2 s_2 + \cdots + \alpha_m s_m,$$

for some particular choice of the coefficients  $\alpha_1, \alpha_2, \dots, \alpha_m$ .

If the vectors  $s_1, s_2, \dots, s_m$  are both linearly independent and a spanning set for  $\mathcal{S}$ , we call them a *basis* of  $\mathcal{S}$ . In this case,  $m$  (the number of elements in the basis) is referred to as the *dimension* of  $\mathcal{S}$ , and denoted by  $\dim(\mathcal{S})$ . Note that there are many ways to choose a basis of  $\mathcal{S}$  in general, but that all bases contain the same number of vectors.

If  $A$  is any real matrix, the *null space* is the subspace

$$\text{Null}(A) = \{w \mid Aw = 0\},$$

while the *range space* is

$$\text{Range}(A) = \{w \mid w = Av \text{ for some vector } v\}.$$

The *fundamental theorem of linear algebra* states that

$$\text{Null}(A) \oplus \text{Range}(A^T) = \mathbb{R}^n,$$

where  $n$  is the number of columns in  $A$ . (Here, “ $\oplus$ ” denotes the direct sum of two sets:  $\mathcal{A} \oplus \mathcal{B} = \{x + y \mid x \in \mathcal{A}, y \in \mathcal{B}\}$ .)

When  $A$  is square ( $n \times n$ ) and nonsingular, we have  $\text{Null}A = \text{Null}A^T = \{0\}$  and  $\text{Range}A = \text{Range}A^T = \mathbb{R}^n$ . In this case, the columns of  $A$  form a basis of  $\mathbb{R}^n$ , as do the columns of  $A^T$ .

## EIGENVALUES, EIGENVECTORS, AND THE SINGULAR-VALUE DECOMPOSITION

A scalar value  $\lambda$  is an *eigenvalue* of the  $n \times n$  matrix  $A$  if there is a nonzero vector  $q$  such that

$$Aq = \lambda q.$$

The vector  $q$  is called an *eigenvector* of  $A$ . The matrix  $A$  is nonsingular if none of its eigenvalues are zero. The eigenvalues of symmetric matrices are all real numbers, while nonsymmetric matrices may have imaginary eigenvalues. If the matrix is positive definite as well as symmetric, its eigenvalues are all *positive* real numbers.

All matrices  $A$  (not necessarily square) can be decomposed as a product of three matrices with special properties. When  $A \in \mathbb{R}^{m \times n}$  with  $m > n$ , (that is,  $A$  has more rows than columns), this *singular-value decomposition* (SVD) has the form

$$A = U \begin{bmatrix} S \\ 0 \end{bmatrix} V^T, \quad (\text{A.15})$$

where  $U$  and  $V$  are orthogonal matrices of dimension  $m \times m$  and  $n \times n$ , respectively, and  $S$  is an  $n \times n$  diagonal matrix with diagonal elements  $\sigma_i$ ,  $i = 1, 2, \dots, n$ , that satisfy

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0.$$

These diagonal values are called the singular values of  $A$ . We can define the condition number (A.11) of the  $m \times n$  (possibly nonsquare) matrix  $A$  to be  $\sigma_1/\sigma_n$ . (This definition is identical to  $\kappa_2(A)$  when  $A$  happens to be square and nonsingular.)

When  $m \leq n$  (the number of columns is at least equal to the number of rows), the SVD has the form

$$A = U \begin{bmatrix} S & 0 \end{bmatrix} V^T,$$

where again  $U$  and  $V$  are orthogonal of dimension  $m \times m$  and  $n \times n$ , respectively, while  $S$  is  $m \times m$  diagonal with nonnegative diagonal elements  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m$ .

When  $A$  is symmetric, its  $n$  real eigenvalues  $\lambda_1, \lambda_2, \dots, \lambda_n$  and their associated eigenvectors  $q_1, q_2, \dots, q_n$  can be used to write a *spectral decomposition* of  $A$  as follows:

$$A = \sum_{i=1}^n \lambda_i q_i q_i^T.$$

This decomposition can be restated in matrix form by defining

$$\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n), \quad Q = [q_1 \mid q_2 \mid \dots \mid q_n],$$

and writing

$$A = Q\Lambda Q^T. \quad (\text{A.16})$$

In fact, when  $A$  is positive definite as well as symmetric, this decomposition is identical to the singular-value decomposition (A.15), where we define  $U = V = Q$  and  $S = \Lambda$ . Note that the singular values  $\sigma_i$  and the eigenvalues  $\lambda_i$  coincide in this case.

In the case of the Euclidean norm (A.8b), we have for symmetric positive definite matrices  $A$  that the singular values and eigenvalues of  $A$  coincide, and that

$$\begin{aligned}\|A\| &= \sigma_1(A) = \text{largest eigenvalue of } A, \\ \|A^{-1}\| &= \sigma_n(A)^{-1} = \text{inverse of smallest eigenvalue of } A.\end{aligned}$$

Hence, we have for all  $x \in \mathbb{R}^n$  that

$$\sigma_n(A)\|x\|^2 = \|x\|^2/\|A^{-1}\| \leq x^T A x \leq \|A\|\|x\|^2 = \sigma_1(A)\|x\|^2.$$

For an orthogonal matrix  $Q$ , we have for the Euclidean norm that

$$\|Qx\| = \|x\|,$$

and that all the singular values of this matrix are equal to 1.

### DETERMINANT AND TRACE

The trace of an  $n \times n$  matrix  $A$  is defined by

$$\text{trace}(A) = \sum_{i=1}^n A_{ii}. \quad (\text{A.17})$$

If the eigenvalues of  $A$  are denoted by  $\lambda_1, \lambda_2, \dots, \lambda_n$ , it can be shown that

$$\text{trace}(A) = \sum_{i=1}^n \lambda_i, \quad (\text{A.18})$$

that is, the trace of the matrix is the sum of its eigenvalues.

The determinant of an  $n \times n$  matrix  $A$ , denoted by  $\det A$ , is the product of its eigenvalues; that is,

$$\det A = \prod_{i=1}^n \lambda_i. \quad (\text{A.19})$$

The determinant has several appealing (and revealing) properties. For instance,

$\det A = 0$  if and only if  $A$  is singular;

$\det AB = (\det A)(\det B)$ ;

$\det A^{-1} = 1/\det A$ .

Recall that any orthogonal matrix  $A$  has the property that  $QQ^T = Q^TQ = I$ , so that  $Q^{-1} = Q^T$ . It follows from the property of the determinant that  $\det Q = \det Q^T = \pm 1$ .

The properties above are used in the analysis of Chapter 6.

### MATRIX FACTORIZATIONS: CHOLESKY, LU, QR

Matrix factorizations are important both in the design of algorithms and in their analysis. One such factorization is the singular-value decomposition defined above in (A.15). Here we define the other important factorizations.

All the factorization algorithms described below make use of *permutation matrices*. Suppose that we wish to exchange the first and fourth rows of a matrix  $A$ . We can perform this operation by premultiplying  $A$  by a permutation matrix  $P$ , which is constructed by interchanging the first and fourth rows of an identity matrix that contains the same number of rows as  $A$ . Suppose, for example, that  $A$  is a  $5 \times 5$  matrix. The appropriate choice of  $P$  would be

$$P = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

A similar technique is used to find a permutation matrix  $P$  that exchanges columns of a matrix.

The LU factorization of a matrix  $A \in \mathbb{R}^{n \times n}$  is defined as

$$PA = LU, \tag{A.20}$$

where

$P$  is an  $n \times n$  permutation matrix (that is, it is obtained by rearranging the rows of the  $n \times n$  identity matrix),

$L$  is unit lower triangular (that is, lower triangular with diagonal elements equal to 1, and

$U$  is upper triangular.

This factorization can be used to solve a linear system of the form  $Ax = b$  efficiently by the following three-step process:

form  $\tilde{b} = Pb$  by permuting the elements of  $b$ ;

solve  $Lz = \tilde{b}$  by performing triangular forward-substitution, to obtain the vector  $z$ ;

solve  $Ux = z$  by performing triangular back-substitution, to obtain the solution vector  $x$ .

The factorization (A.20) can be found by using Gaussian elimination with row partial pivoting, an algorithm that requires approximately  $2n^3/3$  floating-point operations when  $A$  is dense. Standard software that implements this algorithm (notably, LAPACK [7]) is readily available. The method can be stated as follows.

**Algorithm A.1** (Gaussian Elimination with Row Partial Pivoting).

```

Given  $A \in \mathbb{R}^{n \times n}$ ;
Set  $P \leftarrow I, L \leftarrow 0$ ;
for  $i = 1, 2, \dots, n$ 
    find the index  $j \in \{i, i + 1, \dots, n\}$  such that  $|A_{ji}| = \max_{k=i, i+1, \dots, n} |A_{ki}|$ ;
    if  $A_{ij} = 0$ 
        stop; (* matrix  $A$  is singular *)
    if  $i \neq j$ 
        swap rows  $i$  and  $j$  of matrices  $A$  and  $L$ ;
    (* elimination step *)
     $L_{ii} \leftarrow 1$ ;
    for  $k = i + 1, i + 2, \dots, n$ 
         $L_{ki} \leftarrow A_{ki}/A_{ii}$ ;
        for  $l = i + 1, i + 2, \dots, n$ 
             $A_{kl} \leftarrow A_{kl} - L_{ki}A_{il}$ ;
        end (for)
    end (if)
end (for)
 $U \leftarrow$  upper triangular part of  $A$ .

```

Variants of the basic algorithm allow for rearrangement of the columns as well as the rows during the factorization, but these do not add to the practical stability properties of the algorithm. Column pivoting may, however, improve the performance of Gaussian elimination when the matrix  $A$  is sparse. by ensuring that the factors  $L$  and  $U$  are also reasonably sparse.

Gaussian elimination can be applied also to the case in which  $A$  is not square. When  $A$  is  $m \times n$ , with  $m > n$ , the standard row pivoting algorithm produces a factorization of the form (A.20), where  $L \in \mathbb{R}^{m \times n}$  is unit lower triangular and  $U \in \mathbb{R}^{n \times n}$  is upper triangular. When  $m < n$ , we can find an LU factorization of  $A^T$  rather than  $A$ , that is, we obtain

$$PA^T = \begin{bmatrix} L_1 \\ L_2 \end{bmatrix} U, \quad (\text{A.21})$$

where  $L_1$  is  $m \times m$  (square) unit lower triangular,  $U$  is  $m \times m$  upper triangular, and  $L_2$  is a general  $(n - m) \times m$  matrix. If  $A$  has full row rank, we can use this factorization to calculate

its null space explicitly as the space spanned by the columns of the matrix

$$M = P^T \begin{bmatrix} L_1^{-T} L_2^T \\ -I \end{bmatrix} U^{-T}. \quad (\text{A.22})$$

It is easy to check that  $M$  has dimensions  $n \times (n - m)$  and that  $AM = 0$ .

When  $A \in \mathbb{R}^{n \times n}$  is symmetric positive definite, it is possible to compute a similar but more specialized factorization at about half the cost—about  $n^3/3$  operations. This factorization, known as the Cholesky factorization, produces a matrix  $L$  such that

$$A = LL^T. \quad (\text{A.23})$$

(If we require  $L$  to have positive diagonal elements, it is uniquely defined by this formula.) The algorithm can be specified as follows.

**Algorithm A.2** (Cholesky Factorization).

Given  $A \in \mathbb{R}^{n \times n}$  symmetric positive definite;

**for**  $i = 1, 2, \dots, n$ ;

$L_{ii} \leftarrow \sqrt{A_{ii}}$ ;

**for**  $j = i + 1, i + 2, \dots, n$

$L_{ji} \leftarrow A_{ji}/L_{ii}$ ;

**for**  $k = i + 1, i + 2, \dots, j$

$A_{jk} \leftarrow A_{jk} - L_{ji}L_{ki}$ ;

**end (for)**

**end (for)**

**end (for)**

Note that this algorithm references only the lower triangular elements of  $A$ ; in fact, it is only necessary to store these elements in any case, since by symmetry they are simply duplicated in the upper triangular positions.

Unlike the case of Gaussian elimination, the Cholesky algorithm can produce a valid factorization of a symmetric positive definite matrix without swapping any rows or columns. However, symmetric permutation (that is, reordering the rows and columns in the same way) can be used to improve the sparsity of the factor  $L$ . In this case, the algorithm produces a permutation of the form

$$P^T A P = LL^T$$

for some permutation matrix  $P$ .

The Cholesky factorization can be used to compute solutions of the system  $Ax = b$  by performing triangular forward- and back-substitutions with  $L$  and  $L^T$ , respectively, as in the case of  $L$  and  $U$  factors produced by Gaussian elimination.

The Cholesky factorization can also be used to verify positive definiteness of a symmetric matrix  $A$ . If Algorithm A.2 runs to completion with all  $L_{ii}$  values well defined and positive, then  $A$  is positive definite.

Another useful factorization of rectangular matrices  $A \in \mathbb{R}^{m \times n}$  has the form

$$AP = QR, \quad (\text{A.24})$$

where

$P$  is an  $n \times n$  permutation matrix,

$Q$  is  $m \times m$  orthogonal, and

$R$  is  $m \times n$  upper triangular.

In the case of a square matrix  $m = n$ , this factorization can be used to compute solutions of linear systems of the form  $Ax = b$  via the following procedure:

set  $\tilde{b} = Q^T b$ ;

solve  $Rz = \tilde{b}$  for  $z$  by performing back-substitution;

set  $x = P^T z$  by rearranging the elements of  $x$ .

For a dense matrix  $A$ , the cost of computing the QR factorization is about  $4m^2n/3$  operations. In the case of a square matrix, the operation count is about twice as high as for an LU factorization via Gaussian elimination. Moreover, it is more difficult to maintain sparsity in a QR factorization than in an LU factorization.

Algorithms to perform QR factorization are almost as simple as algorithms for Gaussian elimination and for Cholesky factorization. The most widely used algorithms work by applying a sequence of special orthogonal matrices to  $A$ , known either as Householder transformations or Givens rotations, depending on the algorithm. We omit the details, and refer instead to Golub and Van Loan [136, Chapter 5] for a complete description.

In the case of a rectangular matrix  $A$  with  $m < n$ , we can use the QR factorization of  $A^T$  to find a matrix whose columns span the null space of  $A$ . To be specific, we write

$$A^T P = QR = [ Q_1 \quad Q_2 ] R,$$

where  $Q_1$  consists of the first  $m$  columns of  $Q$ , and  $Q_2$  contains the last  $n - m$  columns. It is easy to show that columns of the matrix  $Q_2$  span the null space of  $A$ . This procedure yields a more satisfactory basis matrix for the null space than the Gaussian elimination procedure (A.22), because the columns of  $Q_2$  are orthogonal to each other and have unit length. It may be more expensive to compute, however, particularly in the case in which  $A$  is sparse.

When  $A$  has full column rank, we can make an identification between the  $R$  factor in (A.24) and the Cholesky factorization. By multiplying the formula (A.24) by its transpose,

we obtain

$$P^T A^T A P = R^T Q^T Q R = R^T R,$$

and by comparison with (A.23), we see that  $R^T$  is simply the Cholesky factor of the symmetric positive definite matrix  $P^T A^T A P$ . Recalling that  $L$  is uniquely defined when we restrict its diagonal elements to be positive, this observation implies that  $R$  is also uniquely defined for a given choice of permutation matrix  $P$ , provided that we enforce positiveness of the diagonals of  $R$ . Note, too, that since we can rearrange (A.24) to read  $A P R^{-1} = Q$ , we can conclude that  $Q$  is also uniquely defined under these conditions.

Note that by definition of the Euclidean norm and the property (A.10), and the fact that the Euclidean norms of the matrices  $P$  and  $Q$  in (A.24) are both 1, we have that

$$\|A\| = \|Q R P^T\| \leq \|Q\| \|R\| \|P^T\| = \|R\|,$$

while

$$\|R\| = \|Q^T A P\| \leq \|Q^T\| \|A\| \|P\| = \|A\|.$$

We conclude from these two inequalities that  $\|A\| = \|R\|$ . When  $A$  is square, we have by a similar argument that  $\|A^{-1}\| = \|R^{-1}\|$ . Hence the Euclidean-norm condition number of  $A$  can be estimated by substituting  $R$  for  $A$  in the expression (A.11). This observation is significant because various techniques are available for estimating the condition number of triangular matrices  $R$ ; see Golub and Van Loan [136, pp. 128–130] for a discussion.

### SYMMETRIC INDEFINITE FACTORIZATION

When matrix  $A$  is symmetric but indefinite, Algorithm A.2 will break down by trying to take the square root of a negative number. We can however produce a factorization, similar to the Cholesky factorization, of the form

$$P A P^T = L B L^T, \tag{A.25}$$

where  $L$  is unit lower triangular,  $B$  is a block diagonal matrix with blocks of dimension 1 or 2, and  $P$  is a permutation matrix. The first step of this symmetric indefinite factorization proceeds as follows. We identify a submatrix  $E$  of  $A$  that is suitable to be used as a pivot block. The precise criteria that can be used to choose  $E$  are described below, but we note here that  $E$  is either a single diagonal element of  $A$  (a  $1 \times 1$  pivot block), or else the  $2 \times 2$  block consisting of two diagonal elements of  $A$  (say,  $a_{ii}$  and  $a_{jj}$ ) along with the corresponding off-diagonal elements (that is,  $a_{ij}$  and  $a_{ji}$ ). In either case,  $E$  must be nonsingular. We then

find a permutation matrix  $P_1$  that makes  $E$  a leading principal submatrix of  $A$ , that is,

$$P_1 A P_1 = \begin{bmatrix} E & C^T \\ C & H \end{bmatrix}, \quad (\text{A.26})$$

and then perform a block factorization on this rearranged matrix, using  $E$  as the pivot block, to obtain

$$P_1 A P_1^T = \begin{bmatrix} I & 0 \\ C E^{-1} & I \end{bmatrix} \begin{bmatrix} E & 0 \\ 0 & H - C E^{-1} C^T \end{bmatrix} \begin{bmatrix} I & E^{-1} C^T \\ 0 & I \end{bmatrix}.$$

The next step of the factorization consists in applying exactly the same process to  $H - C E^{-1} C^T$ , known as the *remaining matrix* or the *Schur complement*, which has dimension either  $(n-1) \times (n-1)$  or  $(n-2) \times (n-2)$ . We now apply the same procedure recursively, terminating with the factorization (A.25). Here  $P$  is defined as a product of the permutation matrices from each step of the factorization, and  $B$  contains the pivot blocks  $E$  on its diagonal.

The symmetric indefinite factorization requires approximately  $n^3/3$  floating-point operations—the same as the cost of the Cholesky factorization of a positive definite matrix—but to this count we must add the cost of identifying suitable pivot blocks  $E$  and of performing the permutations, which can be considerable. There are various strategies for determining the pivot blocks, which have an important effect on both the cost of the factorization and its numerical properties. Ideally, our strategy for choosing  $E$  at each step of the factorization procedure should be inexpensive, should lead to at most modest growth in the elements of the remaining matrix at each step of the factorization, and should avoid excessive fill-in (that is,  $L$  should not be too much more dense than  $A$ ).

A well-known strategy, due to Bunch and Parlett [43], searches the whole remaining matrix and identifies the largest-magnitude diagonal and largest-magnitude off-diagonal elements, denoting their respective magnitudes by  $\xi_{\text{dia}}$  and  $\xi_{\text{off}}$ . If the diagonal element whose magnitude is  $\xi_{\text{dia}}$  is selected to be a  $1 \times 1$  pivot block, the element growth in the remaining matrix is bounded by the ratio  $\xi_{\text{dia}}/\xi_{\text{off}}$ . If this growth rate is acceptable, we choose this diagonal element to be the pivot block. Otherwise, we select the off-diagonal element whose magnitude is  $\xi_{\text{off}}$  ( $a_{ij}$ , say), and choose  $E$  to be the  $2 \times 2$  submatrix that includes this element, that is,

$$E = \begin{bmatrix} a_{ii} & a_{ij} \\ a_{ij} & a_{jj} \end{bmatrix}.$$

This pivoting strategy of Bunch and Parlett is numerically stable and guarantees to yield a matrix  $L$  whose maximum element is bounded by 2.781. Its drawback is that the evaluation of  $\xi_{\text{dia}}$  and  $\xi_{\text{off}}$  at each iteration requires many comparisons between floating-point numbers

to be performed:  $O(n^3)$  in total during the overall factorization. Since each comparison costs roughly the same as an arithmetic operation, this overhead is not insignificant.

The more economical pivoting strategy of Bunch and Kaufman [42] searches at most two columns of the working matrix at each stage and requires just  $O(n^2)$  comparisons in total. Its rationale and details are somewhat tricky, and we refer the interested reader to the original paper [42] or to Golub and Van Loan [136, Section 4.4] for details. Unfortunately, this algorithm can give rise to arbitrarily large elements in the lower triangular factor  $L$ , making it unsuitable for use with a modified Cholesky strategy.

The bounded Bunch–Kaufman strategy is essentially a compromise between the Bunch–Parlett and Bunch–Kaufman strategies. It monitors the sizes of elements in  $L$ , accepting the (inexpensive) Bunch–Kaufman choice of pivot block when it yields only modest element growth, but searching further for an acceptable pivot when this growth is excessive. Its total cost is usually similar to that of Bunch–Kaufman, but in the worst case it can approach the cost of Bunch–Parlett.

So far, we have ignored the effect of the choice of pivot block  $E$  on the sparsity of the final  $L$  factor. This consideration is important when the matrix to be factored is large and sparse, since it greatly affects both the CPU time and the amount of storage required by the algorithm. Algorithms that modify the strategies above to take account of sparsity have been proposed by Duff et al. [97], Duff and Reid [95], and Fourer and Mehrotra [113].

### SHERMAN–MORRISON–WOODBURY FORMULA

If the square nonsingular matrix  $A$  undergoes a rank-one update to become

$$\bar{A} = A + ab^T,$$

where  $a, b \in \mathbb{R}^n$ , then if  $\bar{A}$  is nonsingular, we have

$$\bar{A}^{-1} = A^{-1} - \frac{A^{-1}ab^T A^{-1}}{1 + b^T A^{-1}a}. \quad (\text{A.27})$$

It is easy to verify this formula: Simply multiply the definitions of  $\bar{A}$  and  $\bar{A}^{-1}$  together and check that they produce the identity.

This formula can be extended to higher-rank updates. Let  $U$  and  $V$  be matrices in  $\mathbb{R}^{n \times p}$  for some  $p$  between 1 and  $n$ . If we define

$$\hat{A} = A + UV^T,$$

then  $\hat{A}$  is nonsingular if and only if  $(I + V^T A^{-1}U)$  is nonsingular, and in this case we have

$$\hat{A}^{-1} = A^{-1} - A^{-1}U(I + V^T A^{-1}U)^{-1}V^T A^{-1}. \quad (\text{A.28})$$

We can use this formula to solve linear systems of the form  $\bar{A}x = d$ . Since

$$x = \hat{A}^{-1}d = A^{-1}d - A^{-1}U(I + V^T A^{-1}U)^{-1}V^T A^{-1}d,$$

we see that  $x$  can be found by solving  $p + 1$  linear systems with the matrix  $A$  (to obtain  $A^{-1}d$  and  $A^{-1}U$ ), inverting the  $p \times p$  matrix  $I + V^T A^{-1}U$ , and performing some elementary matrix algebra. Inversion of the  $p \times p$  matrix  $I + V^T A^{-1}U$  is inexpensive when  $p \ll n$ .

### INTERLACING EIGENVALUE THEOREM

The following result is proved for example in Golub and Van Loan [136, Theorem 8.1.8].

**Theorem A.1** (Interlacing Eigenvalue Theorem).

Let  $A \in \mathbb{R}^{n \times n}$  be a symmetric matrix with eigenvalues  $\lambda_1, \lambda_2, \dots, \lambda_n$  satisfying

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n,$$

and let  $z \in \mathbb{R}^n$  be a vector with  $\|z\| = 1$ , and  $\alpha \in \mathbb{R}$  be a scalar. Then if we denote the eigenvalues of  $A + \alpha z z^T$  by  $\xi_1, \xi_2, \dots, \xi_n$  (in decreasing order), we have for  $\alpha > 0$  that

$$\xi_1 \geq \lambda_1 \geq \xi_2 \geq \lambda_2 \geq \xi_3 \geq \dots \geq \xi_n \geq \lambda_n,$$

with

$$\sum_{i=1}^n \xi_i - \lambda_i = \alpha. \quad (\text{A.29})$$

If  $\alpha < 0$ , we have that

$$\lambda_1 \geq \xi_1 \geq \lambda_2 \geq \xi_2 \geq \lambda_3 \geq \dots \geq \lambda_n \geq \xi_n,$$

where the relationship (A.29) is again satisfied.

Informally stated, the eigenvalues of the modified matrix “interlace” the eigenvalues of the original matrix, with nonnegative adjustments if the coefficient  $\alpha$  is positive, and nonpositive adjustments if  $\alpha$  is negative. The total magnitude of the adjustments equals  $\alpha$ , whose magnitude is identical to the Euclidean norm  $\|\alpha z z^T\|_2$  of the modification.

### ERROR ANALYSIS AND FLOATING-POINT ARITHMETIC

In most of this book our algorithms and analysis deal with real numbers. Modern digital computers, however, cannot store or compute with general real numbers. Instead,

they work with a subset known as *floating-point numbers*. Any quantities that are stored on the computer, whether they are read directly from a file or program or arise as the intermediate result of a computation, must be approximated by a floating-point number. In general, then, the numbers that are produced by practical computation differ from those that would be produced if the arithmetic were exact. Of course, we try to perform our computations in such a way that these differences are as tiny as possible.

Discussion of errors requires us to distinguish between *absolute error* and *relative error*. If  $x$  is some exact quantity (scalar, vector, matrix) and  $\tilde{x}$  is its approximate value, the absolute error is the norm of the difference, namely,  $\|x - \tilde{x}\|$ . (In general, any of the norms (A.2a), (A.2b), and (A.2c) can be used in this definition.) The relative error is the ratio of the absolute error to the size of the exact quantity, that is,

$$\frac{\|x - \tilde{x}\|}{\|x\|}.$$

When this ratio is significantly less than one, we can replace the denominator by the size of the approximate quantity—that is,  $\|\tilde{x}\|$ —without affecting its value very much.

Most computations associated with optimization algorithms are performed in double-precision arithmetic. Double-precision numbers are stored in words of length 64 bits. Most of these bits (say  $t$ ) are devoted to storing the *fractional part*, while the remainder encode the *exponent*  $e$  and other information, such as the sign of the number, or an indication of whether it is zero or “undefined.” Typically, the fractional part has the form

$$.d_1d_2 \dots d_t,$$

where each  $d_i, i = 1, 2, \dots, t$ , is either zero or one. (In some systems  $d_1$  is implicitly assumed to be 1 and is not stored.) The value of the floating-point number is then

$$\sum_{i=1}^t d_i 2^{-i} \times 2^e.$$

The value  $2^{-t-1}$  is known as *unit roundoff* and is denoted by  $\mathbf{u}$ . Any real number whose absolute value lies in the range  $[2^L, 2^U]$  (where  $L$  and  $U$  are lower and upper bounds on the value of the exponent  $e$ ) can be approximated to within a relative accuracy of  $\mathbf{u}$  by a floating-point number, that is,

$$\text{fl}(x) = x(1 + \epsilon), \quad \text{where } |\epsilon| \leq \mathbf{u}, \quad (\text{A.30})$$

where  $\text{fl}(\cdot)$  denotes floating-point approximation. The value of  $\mathbf{u}$  for double-precision IEEE arithmetic is about  $1.1 \times 10^{-16}$ . In other words, if the real number  $x$  and its floating-point approximation are both written as base-10 numbers (the usual fashion), they agree to at least 15 digits.

For further information on floating-point computations, see Overton [233], Golub and Van Loan [136, Section 2.4], and Higham [169].

When an arithmetic operation is performed with one or two floating-point numbers, the result must also be stored as a floating-point number. This process introduces a small *roundoff error*, whose size can be quantified in terms of the size of the arguments. If  $x$  and  $y$  are two floating-point numbers, we have that

$$|\text{fl}(x * y) - x * y| \leq \mathbf{u}|x * y|, \quad (\text{A.31})$$

where  $*$  denotes any of the operations  $+$ ,  $-$ ,  $\times$ ,  $\div$ .

Although the error in a single floating-point operation appears benign, more significant errors may occur when the arguments  $x$  and  $y$  are floating-point approximations of two *real* numbers, or when a sequence of computations are performed in succession. Suppose, for instance, that  $x$  and  $y$  are large real numbers whose values are very similar. When we store them in a computer, we approximate them with floating-point numbers  $\text{fl}(x)$  and  $\text{fl}(y)$  that satisfy

$$\text{fl}(x) = x + \epsilon_x, \quad \text{fl}(y) = y + \epsilon_y, \quad \text{where } |\epsilon_x| \leq \mathbf{u}|x|, |\epsilon_y| \leq \mathbf{u}|y|.$$

If we take the difference of the two stored numbers, we obtain a final result  $\text{fl}(\text{fl}(x) - \text{fl}(y))$  that satisfies

$$\text{fl}(\text{fl}(x) - \text{fl}(y)) = (\text{fl}(x) - \text{fl}(y))(1 + \epsilon_{xy}), \quad \text{where } |\epsilon_{xy}| \leq \mathbf{u}.$$

By combining these expressions, we find that the difference between this result and the true value  $x - y$  may be as large as

$$\epsilon_x + \epsilon_y + \epsilon_{xy},$$

which is bounded by  $\mathbf{u}(|x| + |y| + |x - y|)$ . Hence, since  $x$  and  $y$  are large and close together, the relative error is approximately  $2\mathbf{u}|x|/|x - y|$ , which may be quite large, since  $|x| \gg |x - y|$ .

This phenomenon is known as *cancellation*. It can also be explained (less formally) by noting that if both  $x$  and  $y$  are accurate to  $k$  digits, and if they agree in the first  $\bar{k}$  digits, then their difference will contain only about  $k - \bar{k}$  significant digits—the first  $\bar{k}$  digits cancel each other out. This observation is the reason for the well-known adage of numerical computing—that one should avoid taking the difference of two similar numbers if at all possible.

## CONDITIONING AND STABILITY

*Conditioning* and *stability* are two terms that are used frequently in connection with numerical computations. Unfortunately, their meaning sometimes varies from author to author, but the general definitions below are widely accepted, and we adhere to them in this book.

Conditioning is a property of the numerical problem at hand (whether it is a linear algebra problem, an optimization problem, a differential equations problem, or whatever). A problem is said to be *well conditioned* if its solution is not affected greatly by small perturbations to the data that define the problem. Otherwise, it is said to be *ill conditioned*.

A simple example is given by the following  $2 \times 2$  system of linear equations:

$$\begin{bmatrix} 1 & 2 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \end{bmatrix}.$$

By computing the inverse of the coefficient matrix, we find that the solution is simply

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} -1 & 2 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 3 \\ 2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

If we replace the first right-hand-side element by 3.00001, the solution becomes  $(x_1, x_2)^T = (0.99999, 1.00001)^T$ , which is only slightly different from its exact value  $(1, 1)^T$ . We would note similar insensitivity if we were to perturb the other elements of the right-hand-side or elements of the coefficient matrix. We conclude that this problem is well conditioned. On the other hand, the problem

$$\begin{bmatrix} 1.00001 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 2.00001 \\ 2 \end{bmatrix}$$

is ill conditioned. Its exact solution is  $x = (1, 1)^T$ , but if we change the first element of the right-hand-side from 2.00001 to 2, the solution would change drastically to  $x = (0, 2)^T$ .

For general square linear systems  $Ax = b$  where  $A \in \mathbb{R}^{n \times n}$ , the condition number of the matrix (defined in (A.11)) can be used to quantify the conditioning. Specifically, if we perturb  $A$  to  $\tilde{A}$  and  $b$  to  $\tilde{b}$  and take  $\tilde{x}$  to be the solution of the perturbed system  $\tilde{A}\tilde{x} = \tilde{b}$ , it can be shown that

$$\frac{\|x - \tilde{x}\|}{\|x\|} \approx \kappa(A) \left[ \frac{\|A - \tilde{A}\|}{\|A\|} + \frac{\|b - \tilde{b}\|}{\|b\|} \right]$$

(see, for instance, Golub and Van Loan [136, Section 2.7]). Hence, a large condition number  $\kappa(A)$  indicates that the problem  $Ax = b$  is ill conditioned, while a modest value indicates well conditioning.

Note that the concept of conditioning has nothing to do with the particular algorithm that is used to solve the problem, only with the numerical problem itself.

*Stability*, on the other hand, is a property of the algorithm. An algorithm is stable if it is guaranteed to produce accurate answers to all well-conditioned problems in its class, even when floating-point arithmetic is used.

As an example, consider again the linear equations  $Ax = b$ . We can show that Algorithm A.1, in combination with triangular substitution, yields a computed solution  $\tilde{x}$  whose relative error is approximately

$$\frac{\|x - \tilde{x}\|}{\|x\|} \approx \kappa(A) \frac{\text{growth}(A)}{\|A\|} \mathbf{u}, \quad (\text{A.32})$$

where  $\text{growth}(A)$  is the size of the largest element that arises in  $A$  during execution of Algorithm A.1. In the worst case, we can show that  $\text{growth}(A)/\|A\|$  may be around  $2^{n-1}$ , which indicates that Algorithm A.1 is an unstable algorithm, since even for modest  $n$  (say,  $n = 200$ ), the right-hand-side of (A.32) may be large even when  $\kappa(A)$  is modest. In practice, however, large growth factors are rarely observed, so we conclude that Algorithm A.1 is stable for all practical purposes.

Gaussian elimination without pivoting, on the other hand, is definitely unstable. If we omit the possible exchange of rows in Algorithm A.1, the algorithm will fail to produce a factorization even of some well-conditioned matrices, such as

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}.$$

For systems  $Ax = b$  in which  $A$  is symmetric positive definite, the Cholesky factorization in combination with triangular substitution constitutes a stable algorithm for producing a solution  $x$ .

## **A.2 ELEMENTS OF ANALYSIS, GEOMETRY, TOPOLOGY**

### **SEQUENCES**

Suppose that  $\{x_k\}$  is a sequence of points belonging to  $\mathbb{R}^n$ . We say that a sequence  $\{x_k\}$  converges to some point  $x$ , written  $\lim_{k \rightarrow \infty} x_k = x$ , if for any  $\epsilon > 0$ , there is an index  $K$  such that

$$\|x_k - x\| \leq \epsilon, \quad \text{for all } k \geq K.$$

For example, the sequence  $\{x_k\}$  defined by  $x_k = (1 - 2^{-k}, 1/k^2)^T$  converges to  $(1, 0)^T$ .

Given an index set  $\mathcal{S} \subset \{1, 2, 3, \dots\}$ , we can define a *subsequence* of  $\{t_k\}$  corresponding to  $\mathcal{S}$ , and denote it by  $\{t_k\}_{k \in \mathcal{S}}$ .

We say that  $\hat{x} \in \mathbb{R}^n$  is an *accumulation point* or *limit point* for  $\{x_k\}$  if there is an infinite set of indices  $k_1, k_2, k_3, \dots$  such that the subsequence  $\{x_{k_i}\}_{i=1,2,3,\dots}$  converges to  $\hat{x}$ ; that is,

$$\lim_{i \rightarrow \infty} x_{k_i} = \hat{x}.$$

Alternatively, we say that for any  $\epsilon > 0$  and all positive integers  $K$ , we have

$$\|x_k - x\| \leq \epsilon, \quad \text{for some } k \geq K.$$

An example is given by the sequence

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1/4 \\ 1/4 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1/8 \\ 1/8 \end{bmatrix}, \dots, \quad (\text{A.33})$$

which has exactly two limit points:  $\hat{x} = (0, 0)^T$  and  $\hat{x} = (1, 1)^T$ . A sequence can even have an infinite number of limit points. An example is the sequence  $x_k = \sin k$ , for which every point in the interval  $[-1, 1]$  is a limit point. A sequence converges if and only if it has exactly one limit point.

A sequence is said to be a Cauchy sequence if for any  $\epsilon > 0$ , there exists an integer  $K$  such that  $\|x_k - x_l\| \leq \epsilon$  for all indices  $k \geq K$  and  $l \geq K$ . A sequence converges if and only if it is a Cauchy sequence.

We now consider *scalar sequences*  $\{t_k\}$ , that is,  $t_k \in \mathbb{R}$  for all  $k$ . This sequence is said to be *bounded above* if there exists a scalar  $u$  such that  $t_k \leq u$  for all  $k$ , and *bounded below* if there is a scalar  $v$  with  $t_k \geq v$  for all  $k$ . The sequence  $\{t_k\}$  is said to be *nondecreasing* if  $t_{k+1} \geq t_k$  for all  $k$ , and *nonincreasing* if  $t_{k+1} \leq t_k$  for all  $k$ . If  $\{t_k\}$  is *nondecreasing and bounded above*, then it converges, that is,  $\lim_{k \rightarrow \infty} t_k = t$  for some scalar  $t$ . Similarly, if  $\{t_k\}$  is *nonincreasing and bounded below*, it converges.

We define the *supremum* of the scalar sequence  $\{t_k\}$  as the smallest real number  $u$  such that  $t_k \leq u$  for all  $k = 1, 2, 3, \dots$ , and denote it by  $\sup\{t_k\}$ . The *infimum*, denoted by  $\inf\{t_k\}$ , is the largest real number  $v$  such that  $v \leq t_k$  for all  $k = 1, 2, 3, \dots$ . We can now define the sequence of suprema as  $\{u_i\}$ , where

$$u_i \stackrel{\text{def}}{=} \sup\{t_k \mid k \geq i\}.$$

Clearly,  $\{u_i\}$  is a nonincreasing sequence. If bounded below, it converges to a finite number  $\bar{u}$ , which we call the “lim sup” of  $\{t_k\}$ , denoted by  $\limsup t_k$ . Similarly, we can denote the sequence of infima by  $\{v_i\}$ , where

$$v_i \stackrel{\text{def}}{=} \inf\{t_k \mid k \geq i\},$$

which is nondecreasing. If  $\{v_i\}$  is bounded above, it converges to a point  $\bar{v}$  which we call the “lim inf” of  $\{t_k\}$ , denoted by  $\liminf t_k$ . As an example, the sequence  $1, \frac{1}{2}, 1, \frac{1}{4}, 1, \frac{1}{8}, \dots$  has a lim inf of 0 and a lim sup of 1.

### RATES OF CONVERGENCE

One of the key measures of performance of an algorithm is its rate of convergence. Here, we define the terminology associated with different types of convergence.

Let  $\{x_k\}$  be a sequence in  $\mathbb{R}^n$  that converges to  $x^*$ . We say that the convergence is *Q-linear* if there is a constant  $r \in (0, 1)$  such that

$$\frac{\|x_{k+1} - x^*\|}{\|x_k - x^*\|} \leq r, \quad \text{for all } k \text{ sufficiently large.} \quad (\text{A.34})$$

This means that the distance to the solution  $x^*$  decreases at each iteration by at least a constant factor bounded away from 1. For example, the sequence  $1 + (0.5)^k$  converges Q-linearly to 1, with rate  $r = 0.5$ . The prefix “Q” stands for “quotient,” because this type of convergence is defined in terms of the quotient of successive errors.

The convergence is said to be *Q-superlinear* if

$$\lim_{k \rightarrow \infty} \frac{\|x_{k+1} - x^*\|}{\|x_k - x^*\|} = 0.$$

For example, the sequence  $1 + k^{-k}$  converges superlinearly to 1. (Prove this statement!) *Q-quadratic* convergence, an even more rapid convergence rate, is obtained if

$$\frac{\|x_{k+1} - x^*\|}{\|x_k - x^*\|^2} \leq M, \quad \text{for all } k \text{ sufficiently large,}$$

where  $M$  is a positive constant, not necessarily less than 1. An example is the sequence  $1 + (0.5)^{2^k}$ .

The speed of convergence depends on  $r$  and (more weakly) on  $M$ , whose values depend not only on the algorithm but also on the properties of the particular problem. Regardless of these values, however, a quadratically convergent sequence will always eventually converge faster than a linearly convergent sequence.

Obviously, any sequence that converges Q-quadratically also converges Q-super-linearly, and any sequence that converges Q-superlinearly also converges Q-linearly. We can also define higher rates of convergence (cubic, quartic, and so on), but these are less interesting in practical terms. In general, we say that the Q-order of convergence is  $p$  (with  $p > 1$ ) if there is a positive constant  $M$  such that

$$\frac{\|x_{k+1} - x^*\|}{\|x_k - x^*\|^p} \leq M, \quad \text{for all } k \text{ sufficiently large.}$$

Quasi-Newton methods for unconstrained optimization typically converge Q-superlinearly, whereas Newton's method converges Q-quadratically under appropriate assumptions. In contrast, steepest descent algorithms converge only at a Q-linear rate, and when the problem is ill-conditioned the convergence constant  $r$  in (A.34) is close to 1.

In the book, we omit the letter  $Q$  and simply talk about superlinear convergence, quadratic convergence, and so on.

A slightly weaker form of convergence, characterized by the prefix "R" (for "root"), is concerned with the overall rate of decrease in the error, rather than the decrease over each individual step of the algorithm. We say that convergence is *R-linear* if there is a sequence of nonnegative scalars  $\{v_k\}$  such that

$$\|x_k - x^*\| \leq v_k \text{ for all } k, \text{ and } \{v_k\} \text{ converges Q-linearly to zero.}$$

The sequence  $\{\|x_k - x^*\|\}$  is said to be *dominated* by  $\{v_k\}$ . For instance, the sequence

$$x_k = \begin{cases} 1 + (0.5)^k, & k \text{ even,} \\ 1, & k \text{ odd,} \end{cases} \quad (\text{A.35})$$

(the first few iterates are 2, 1, 1.25, 1, 1.03125, 1, ...) converges R-linearly to 1, because we have  $(1 + (0.5)^k) - 1 = (0.5)^k$ , and the sequence  $\{(0.5)^k\}$  converges Q-linearly to zero. Likewise, we say that  $\{x_k\}$  converges R-superlinearly to  $x^*$  if  $\{\|x_k - x^*\|\}$  is dominated by a sequence of scalars converging Q-superlinearly to zero, and  $\{x_k\}$  converges R-quadratically to  $x^*$  if  $\{\|x_k - x^*\|\}$  is dominated by a sequence converging Q-quadratically to zero.

Note that in the R-linear sequence (A.35), the error actually increases at every second iteration! Such behavior occurs even in sequences whose R-rate of convergence is arbitrarily high, but it cannot occur for Q-linear sequences, which insist on a decrease at every step  $k$ , for  $k$  sufficiently large.

For an extensive discussion of convergence rates see Ortega and Rheinboldt [230].

## TOPOLOGY OF THE EUCLIDEAN SPACE $\mathbb{R}^n$

The set  $\mathcal{F}$  is *bounded* if there is some real number  $M > 0$  such that

$$\|x\| \leq M, \quad \text{for all } x \in \mathcal{F}.$$

A subset  $\mathcal{F} \subset \mathbb{R}^n$  is *open* if for every  $x \in \mathcal{F}$ , we can find a positive number  $\epsilon > 0$  such that the ball of radius  $\epsilon$  around  $x$  is contained in  $\mathcal{F}$ ; that is,

$$\{y \in \mathbb{R}^n \mid \|y - x\| \leq \epsilon\} \subset \mathcal{F}.$$

The set  $\mathcal{F}$  is *closed* if for all possible sequences of points  $\{x_k\}$  in  $\mathcal{F}$ , all limit points of  $\{x_k\}$  are elements of  $\mathcal{F}$ . For instance, the set  $\mathcal{F} = (0, 1) \cup (2, 10)$  is an open subset of  $\mathbb{R}$ , while

$\mathcal{F} = [0, 1] \cup [2, 5]$  is a closed subset of  $\mathbb{R}$ . The set  $\mathcal{F} = (0, 1]$  is a subset of  $\mathbb{R}$  that is neither open nor closed.

The *interior* of a set  $\mathcal{F}$ , denoted by  $\text{int } \mathcal{F}$ , is the largest open set contained in  $\mathcal{F}$ . The *closure* of  $\mathcal{F}$ , denoted by  $\text{cl } \mathcal{F}$ , is the smallest closed set containing  $\mathcal{F}$ . In other words, we have

$$x \in \text{cl } \mathcal{F} \quad \text{if } \lim_{k \rightarrow \infty} x_k = x \text{ for some sequence } \{x_k\} \text{ of points in } \mathcal{F}.$$

If  $\mathcal{F} = (-1, 1] \cup [2, 4)$ , then

$$\text{cl } \mathcal{F} = [-1, 1] \cup [2, 4], \quad \text{int } \mathcal{F} = (-1, 1) \cup (2, 4).$$

Note that if  $\mathcal{F}$  is open, then  $\text{int } \mathcal{F} = \mathcal{F}$ , while if  $\mathcal{F}$  is closed, then  $\text{cl } \mathcal{F} = \mathcal{F}$ .

We note the following facts about open and closed sets. The union of finitely many closed sets is closed, while any intersection of closed sets is closed. The intersection of finitely many open sets is open, while any union of open sets is open.

The set  $\mathcal{F}$  is *compact* if every sequence  $\{x^k\}$  of points in  $\mathcal{F}$  has at least one limit point, and all such limit points are in  $\mathcal{F}$ . (This definition is equivalent to the more formal one involving covers of  $\mathcal{F}$ .) The following is a central result in topology:

$$\mathcal{F} \in \mathbb{R}^n \text{ is closed and bounded} \Rightarrow \mathcal{F} \text{ is compact.}$$

Given a point  $x \in \mathbb{R}^n$ , we call  $\mathcal{N} \in \mathbb{R}^n$  a *neighborhood* of  $x$  if it is an open set containing  $x$ . An especially useful neighborhood is the *open ball of radius  $\epsilon$  around  $x$* , which is denoted by  $\mathbf{B}(x, \epsilon)$ ; that is,

$$\mathbf{B}(x, \epsilon) = \{y \mid \|y - x\| < \epsilon\}.$$

Given a set  $\mathcal{F} \subset \mathbb{R}^n$ , we say that  $\mathcal{N}$  is a *neighborhood* of  $\mathcal{F}$  if there is  $\epsilon > 0$  such that

$$\cup_{x \in \mathcal{F}} \mathbf{B}(x, \epsilon) \subset \mathcal{N}.$$

### CONVEX SETS IN $\mathbb{R}^n$

A *convex combination* of a finite set of vectors  $\{x_1, x_2, \dots, x_m\}$  in  $\mathbb{R}^m$  is any vector  $x$  of the form

$$x = \sum_{i=1}^m \alpha_i x_i, \quad \text{where } \sum_{i=1}^m \alpha_i = 1, \quad \text{and } \alpha_i \geq 0 \text{ for all } i = 1, 2, \dots, m.$$

The *convex hull* of  $\{x_1, x_2, \dots, x_m\}$  is the set of all convex combinations of these vectors.

A *cone* is a set  $\mathcal{F}$  with the property that for all  $x \in \mathcal{F}$  we have

$$x \in \mathcal{F} \Rightarrow \alpha x \in \mathcal{F}, \quad \text{for all } \alpha > 0. \tag{A.36}$$

For instance, the set  $\mathcal{F} \subset \mathbb{R}^2$  defined by

$$\{(x_1, x_2)^T \mid x_1 > 0, x_2 \geq 0\}$$

is a cone in  $\mathbb{R}^2$ . Note that cones are not necessarily convex. For example, the set  $\{(x_1, x_2)^T \mid x_1 \geq 0 \text{ or } x_2 \geq 0\}$ , which encompasses three quarters of the two-dimensional plane, is a cone.

The *cone generated by*  $\{x_1, x_2, \dots, x_m\}$  is the set of all vectors  $x$  of the form

$$x = \sum_{i=1}^m \alpha_i x_i, \quad \text{where } \alpha_i \geq 0 \text{ for all } i = 1, 2, \dots, m.$$

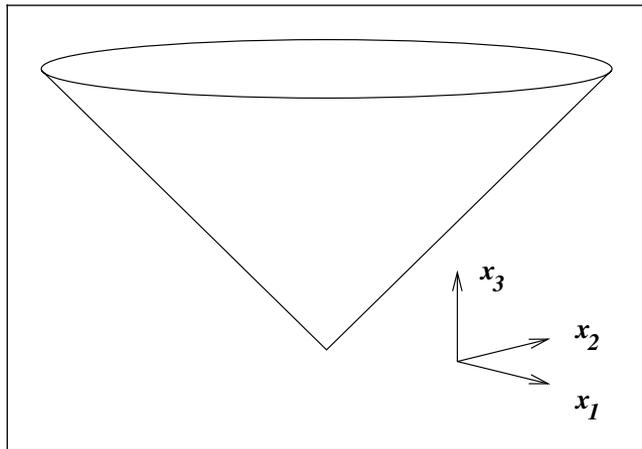
Note that all cones of this form are convex.

Finally, we define the affine hull and relative interior of a set. An *affine set* in  $\mathbb{R}^n$  is a the set of all vectors  $\{x\} \oplus \mathcal{S}$ , where  $x \in \mathbb{R}^n$  and  $\mathcal{S}$  is a subspace of  $\mathbb{R}^n$ . Given  $\mathcal{F} \subset \mathbb{R}^n$ , the *affine hull* of  $\mathcal{F}$  (denoted by  $\text{aff } \mathcal{F}$ ) is the smallest affine set containing  $\mathcal{F}$ . For instance, when  $\mathcal{F}$  is the “ice-cream cone” defined in three dimensions as

$$\Omega = \left\{ x \in \mathbb{R}^3 \mid x_3 \geq 2\sqrt{x_1^2 + x_2^2} \right\} \quad (\text{A.37})$$

(see Figure A.1), we have  $\text{aff } \mathcal{F} = \mathbb{R}^3$ . If  $\mathcal{F}$  is the set of two isolated points  $\mathcal{F} = \{(1, 0, 0)^T, (0, 2, 0)^T\}$ , we have

$$\text{aff } \mathcal{F} = \{(1, 0, 0)^T + \alpha(-1, 2, 0)^T \mid \text{for all } \alpha \in \mathbb{R}\}.$$



**Figure A.1** “Ice-cream cone” set.

The *relative interior*  $\text{ri } \mathcal{F}$  of the set  $\mathcal{F}$  is its *interior relative to*  $\text{aff } \mathcal{F}$ . If  $x \in \mathcal{F}$ , then  $x \in \text{ri } \mathcal{F}$  if there is an  $\epsilon > 0$  such that

$$(x + \epsilon \mathcal{B}) \cap \text{aff } \mathcal{F} \subset \mathcal{F}.$$

Referring again to the ice-cream cone (A.37), we have that

$$\text{ri } \mathcal{F} = \left\{ x \in \mathbb{R}^3 \mid x_3 > 2\sqrt{x_1^2 + x_2^2} \right\}.$$

For the set of two isolated points  $\mathcal{F} = \{(1, 0, 0)^T, (0, 2, 0)^T\}$ , we have  $\text{ri } \mathcal{F} = \emptyset$ . For the set  $\mathcal{F}$  defined by

$$\mathcal{F} \stackrel{\text{def}}{=} \{x \in \mathbb{R}^3 \mid x_1 \in [0, 1], x_2 \in [0, 1], x_3 = 0\},$$

we have that

$$\text{aff } \mathcal{F} = \mathbb{R} \times \mathbb{R} \times \{0\}, \quad \text{ri } \mathcal{F} = \{x \in \mathbb{R}^3 \mid x_1 \in (0, 1), x_2 \in (0, 1), x_3 = 0\}.$$

## CONTINUITY AND LIMITS

Let  $f$  be a function that maps some domain  $\mathcal{D} \subset \mathbb{R}^n$  to the space  $\mathbb{R}^m$ . For some point  $x_0 \in \text{cl } \mathcal{D}$ , we write

$$\lim_{x \rightarrow x_0} f(x) = f_0 \tag{A.38}$$

(spoken “the limit of  $f(x)$  as  $x$  approaches  $x_0$  is  $f_0$ ”) if for all  $\epsilon > 0$ , there is a value  $\delta > 0$  such that

$$\|x - x_0\| < \delta \text{ and } x \in \mathcal{D} \Rightarrow \|f(x) - f_0\| < \epsilon.$$

We say that  $f$  is *continuous* at  $x_0$  if  $x_0 \in \mathcal{D}$  and the expression (A.38) holds with  $f_0 = f(x_0)$ . We say that  $f$  is *continuous* on its domain  $\mathcal{D}$  if  $f$  is continuous for all  $x_0 \in \mathcal{D}$ .

An example is provided by the function

$$f(x) = \begin{cases} -x & \text{if } x \in [-1, 1], x \neq 0, \\ 5 & \text{for all other } x \in [-10, 10]. \end{cases} \tag{A.39}$$

This function is defined on the domain  $[-10, 10]$  and is continuous at all points of the domain except the points  $x = 0$ ,  $x = 1$ , and  $x = -1$ . At  $x = 0$ , the expression (A.38) holds with  $f_0 = 0$ , but the function is not continuous at this point because  $f_0 \neq f(0) = 5$ . At

$x = -1$ , the limit (A.38) is not defined, because the function values in the neighborhood of this point are close to both 5 and  $-1$ , depending on whether  $x$  is slightly smaller or slightly larger than  $-1$ . Hence, the function is certainly not continuous at this point. The same comments apply to the point  $x = 1$ .

In the special case of  $n = 1$  (that is, the argument of  $f$  is a real scalar), we can also define the *one-sided limit*. Given  $x_0 \in \text{cl}\mathcal{D}$ , We write

$$\lim_{x \downarrow x_0} f(x) = f_0 \quad (\text{A.40})$$

(spoken “the limit of  $f(x)$  as  $x$  approaches  $x_0$  from above is  $f_0$ ”) if for all  $\epsilon > 0$ , there is a value  $\delta > 0$  such that

$$x_0 < x < x_0 + \delta \text{ and } x \in \mathcal{D} \Rightarrow \|f(x) - f_0\| < \epsilon.$$

Similarly, we write

$$\lim_{x \uparrow x_0} f(x) = f_0 \quad (\text{A.41})$$

(spoken “the limit of  $f(x)$  as  $x$  approaches  $x_0$  from below is  $f_0$ ”) if for all  $\epsilon > 0$ , there is a value  $\delta > 0$  such that

$$x_0 - \delta < x < x_0 \text{ and } x \in \mathcal{D} \Rightarrow \|f(x) - f_0\| < \epsilon.$$

For the function defined in (A.39), we have that

$$\lim_{x \downarrow 1} f(x) = 5, \quad \lim_{x \uparrow 1} f(x) = 1.$$

Considering again the general case of  $f : \mathcal{D} \rightarrow \mathbb{R}^m$  where  $\mathcal{D} \subset \mathbb{R}^n$  for general  $m$  and  $n$ . The function  $f$  is said to be *Lipschitz continuous* on some set  $\mathcal{N} \subset \mathcal{D}$  if there is a constant  $L > 0$  such that

$$\|f(x_1) - f(x_0)\| \leq L\|x_1 - x_0\|, \quad \text{for all } x_0, x_1 \in \mathcal{N}. \quad (\text{A.42})$$

( $L$  is called the *Lipschitz constant*.) The function  $f$  is *locally Lipschitz continuous* at a point  $\bar{x} \in \text{int}\mathcal{D}$  if there is some neighborhood  $\mathcal{N}$  of  $\bar{x}$  with  $\mathcal{N} \subset \mathcal{D}$  such that the property (A.42) holds for some  $L > 0$ .

If  $g$  and  $h$  are two functions mapping  $\mathcal{D} \subset \mathbb{R}^n$  to  $\mathbb{R}^m$ , Lipschitz continuous on a set  $\mathcal{N} \subset \mathcal{D}$ , their sum  $g + h$  is also Lipschitz continuous, with Lipschitz constant equal to the sum of the Lipschitz constants for  $g$  and  $h$  individually. If  $g$  and  $h$  are two functions mapping  $\mathcal{D} \subset \mathbb{R}^n$  to  $\mathbb{R}$ , the product  $gh$  is Lipschitz continuous on a set  $\mathcal{N} \subset \mathcal{D}$  if both  $g$  and  $h$  are Lipschitz continuous on  $\mathcal{N}$  and both are bounded on  $\mathcal{N}$  (that is, there is  $M > 0$

such that  $|g(x)| \leq M$  and  $|h(x)| \leq M$  for all  $x \in \mathcal{N}$ . We prove this claim via a sequence of elementary inequalities, for arbitrary  $x_0, x_1 \in \mathcal{N}$ :

$$\begin{aligned} & |g(x_0)h(x_0) - g(x_1)h(x_1)| \\ & \leq |g(x_0)h(x_0) - g(x_1)h(x_0)| + |g(x_1)h(x_0) - g(x_1)h(x_1)| \\ & = |h(x_0)| |g(x_0) - g(x_1)| + |g(x_1)| |h(x_0) - h(x_1)| \\ & \leq 2ML \|x_0 - x_1\|, \end{aligned} \tag{A.43}$$

where  $L$  is an upper bound on the Lipschitz constant for both  $g$  and  $h$ .

### DERIVATIVES

Let  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  be a real-valued function of a real variable (sometimes known as a *univariate* function). The first derivative  $\phi'(\alpha)$  is defined by

$$\frac{d\phi}{d\alpha} = \phi'(\alpha) \stackrel{\text{def}}{=} \lim_{\epsilon \rightarrow 0} \frac{\phi(\alpha + \epsilon) - \phi(\alpha)}{\epsilon}. \tag{A.44}$$

The second derivative is obtained by substituting  $\phi$  by  $\phi'$  in this same formula; that is,

$$\frac{d^2\phi}{d\alpha^2} = \phi''(\alpha) \stackrel{\text{def}}{=} \lim_{\epsilon \rightarrow 0} \frac{\phi'(\alpha + \epsilon) - \phi'(\alpha)}{\epsilon}. \tag{A.45}$$

Suppose now that  $\alpha$  in turn depends on another quantity  $\beta$  (we denote this dependence by writing  $\alpha = \alpha(\beta)$ ). We can use the *chain rule* to calculate the derivative of  $\phi$  with respect to  $\beta$ :

$$\frac{d\phi(\alpha(\beta))}{d\beta} = \frac{d\phi}{d\alpha} \frac{d\alpha}{d\beta} = \phi'(\alpha)\alpha'(\beta). \tag{A.46}$$

Consider now the function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ , which is a real-valued function of  $n$  independent variables. We typically gather the variables into a vector  $x = (x_1, x_2, \dots, x_n)^T$ . We say that  $f$  is differentiable at  $x$  if there exists a vector  $g \in \mathbb{R}^n$  such that

$$\lim_{y \rightarrow 0} \frac{f(x + y) - f(x) - g^T y}{\|y\|} = 0, \tag{A.47}$$

where  $\|\cdot\|$  is any vector norm of  $y$ . (This type of differentiability is known as *Frechet differentiability*.) If  $g$  satisfying (A.47) exists, we call it the *gradient* of  $f$  at  $x$ , and denote it

by  $\nabla f(x)$ , written componentwise as

$$\nabla f(x) = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix}. \quad (\text{A.48})$$

Here,  $\partial f/\partial x_i$  represents the partial derivative of  $f$  with respect to  $x_i$ . By setting  $y = \epsilon e_i$  in (A.47), where  $e_i$  is the vector in  $\mathbb{R}^n$  consisting all all zeros, except for a 1 in position  $i$ , we obtain

$$\begin{aligned} & \frac{\partial f}{\partial x_i} \\ & \stackrel{\text{def}}{=} \lim_{\epsilon \rightarrow 0} \frac{f(x_1, \dots, x_{i-1}, x_i + \epsilon, x_{i+1}, \dots, x_n) - f(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n)}{\epsilon} \\ & = \frac{f(x + \epsilon e_i) - f(x)}{\epsilon}. \end{aligned}$$

A gradient with respect to only a subset of the unknowns can be expressed by means of a subscript on the symbol  $\nabla$ . Thus for the function of two vector variables  $f(z, t)$ , we use  $\nabla_z f(z, t)$  to denote the gradient with respect to  $z$  (holding  $t$  constant).

The matrix of second partial derivatives of  $f$  is known as the *Hessian*, and is defined as

$$\nabla^2 f(x) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}.$$

We say that  $f$  is *differentiable* on a domain  $\mathcal{D}$  if  $\nabla f(x)$  exists for all  $x \in \mathcal{D}$ , and *continuously differentiable* if  $\nabla f(x)$  is a continuous functions of  $x$ . Similarly,  $f$  is *twice differentiable* on  $\mathcal{D}$  if  $\nabla^2 f(x)$  exists for all  $x \in \mathcal{D}$  and *twice continuously differentiable* if  $\nabla^2 f(x)$  is continuous on  $\mathcal{D}$ . Note that when  $f$  is twice continuously differentiable, the Hessian is a symmetric matrix, since

$$\frac{\partial^2 f}{\partial x_i \partial x_j} = \frac{\partial^2 f}{\partial x_j \partial x_i}, \quad \text{for all } i, j = 1, 2, \dots, n.$$

When  $f$  is a vector valued function that is  $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$  (See Chapters 10 and 11), we define  $\nabla f(x)$  to be the  $n \times m$  matrix whose  $i$ th column is  $\nabla f_i(x)$ , that is, the gradient of

$f_i$  with respect to  $x$ . Often, for notational convenience, we prefer to work with the transpose of his matrix, which has dimensions  $m \times n$ . This matrix is called the *Jacobian* and is often denoted by  $J(x)$ . Specifically, the  $(i, j)$  element of  $J(x)$  is  $\partial f_i(x)/\partial x_j$ .

When the vector  $x$  in turn depends on another vector  $t$  (that is,  $x = x(t)$ ), we can extend the chain rule (A.46) for the univariate function. Defining

$$h(t) = f(x(t)), \quad (\text{A.49})$$

we have

$$\nabla h(t) = \sum_{i=1}^n \frac{\partial f}{\partial x_i} \nabla x_i(t) = \nabla x(t) \nabla f(x(t)). \quad (\text{A.50})$$

□ **EXAMPLE A.1**

Let  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$  be defined by  $f(x_1, x_2) = x_1^2 + x_1 x_2$ , where  $x_1 = \sin t_1 + t_2^2$  and  $x_2 = (t_1 + t_2)^2$ . Defining  $h(t)$  as in (A.49), the chain rule (A.50) yields

$$\begin{aligned} \nabla h(t) &= \sum_{i=1}^n \frac{\partial f}{\partial x_i} \nabla x_i(t) \\ &= (2x_1 + x_2) \begin{bmatrix} \cos t_1 \\ 2t_2 \end{bmatrix} + x_1 \begin{bmatrix} 2(t_1 + t_2) \\ 2(t_1 + t_2) \end{bmatrix} \\ &= (2(\sin t_1 + t_2^2) + (t_1 + t_2)^2) \begin{bmatrix} \cos t_1 \\ 2t_2 \end{bmatrix} + (\sin t_1 + t_2^2) \begin{bmatrix} 2(t_1 + t_2) \\ 2(t_1 + t_2) \end{bmatrix}. \end{aligned}$$

If, on the other hand, we substitute directly for  $x$  into the definition of  $f$ , we obtain

$$h(t) = f(x(t)) = (\sin t_1 + t_2^2)^2 + (\sin t_1 + t_2^2) (t_1 + t_2)^2.$$

The reader should verify that the gradient of this expression is identical to the one obtained above by applying the chain rule. □

Special cases of the chain rule can be derived when  $x(t)$  in (A.50) is a linear function of  $t$ , say  $x(t) = Ct$ . We then have  $\nabla x(t) = C^T$ , so that

$$\nabla h(t) = C^T \nabla f(Ct).$$

In the case in which  $f$  is a scalar function, we can differentiate twice using the chain rule to obtain

$$\nabla^2 h(t) = C^T \nabla^2 f(Ct) C.$$

(The proof of this statement is left as an exercise.)

### DIRECTIONAL DERIVATIVES

The *directional derivative* of a function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  in the direction  $p$  is given by

$$D(f(x); p) \stackrel{\text{def}}{=} \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon p) - f(x)}{\epsilon}. \quad (\text{A.51})$$

The directional derivative may be well defined even when  $f$  is not continuously differentiable; in fact, it is most useful in such situations. Consider for instance the  $\ell_1$  norm function  $f(x) = \|x\|_1$ . We have from the definition (A.51) that

$$D(\|x\|_1; p) = \lim_{\epsilon \rightarrow 0} \frac{\|x + \epsilon p\|_1 - \|x\|_1}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{\sum_{i=1}^n |x_i + \epsilon p_i| - \sum_{i=1}^n |x_i|}{\epsilon}.$$

If  $x_i > 0$ , we have  $|x_i + \epsilon p_i| = |x_i| + \epsilon p_i$  for all  $\epsilon$  sufficiently small. If  $x_i < 0$ , we have  $|x_i + \epsilon p_i| = |x_i| - \epsilon p_i$ , while if  $x_i = 0$ , we have  $|x_i + \epsilon p_i| = \epsilon |p_i|$ . Therefore, we have

$$D(\|x\|_1; p) = \sum_{i|x_i < 0} -p_i + \sum_{i|x_i > 0} p_i + \sum_{i|x_i = 0} |p_i|,$$

so the directional derivative of this function exists for any  $x$  and  $p$ . The first derivative  $\nabla f(x)$  does *not* exist, however, whenever any of the components of  $x$  are zero.

When  $f$  is in fact continuously differentiable in a neighborhood of  $x$ , we have

$$D(f(x); p) = \nabla f(x)^T p.$$

To verify this formula, we define the function

$$\phi(\alpha) = f(x + \alpha p) = f(y(\alpha)), \quad (\text{A.52})$$

where  $y(\alpha) = x + \alpha p$ . Note that

$$\lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon p) - f(x)}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{\phi(\epsilon) - \phi(0)}{\epsilon} = \phi'(0).$$

By applying the chain rule (A.50) to  $f(y(\alpha))$ , we obtain

$$\begin{aligned}\phi'(\alpha) &= \sum_{i=1}^n \frac{\partial f(y(\alpha))}{\partial y_i} \nabla y_i(\alpha) \\ &= \sum_{i=1}^n \frac{\partial f(y(\alpha))}{\partial y_i} p_i = \nabla f(y(\alpha))^T p = \nabla f(x + \alpha p)^T p.\end{aligned}\tag{A.53}$$

We obtain (A.51) by setting  $\alpha = 0$  and comparing the last two expressions.

### MEAN VALUE THEOREM

We now recall the mean value theorem for univariate functions. Given a continuously differentiable function  $\phi : \mathbb{R} \rightarrow \mathbb{R}$  and two real numbers  $\alpha_0$  and  $\alpha_1$  that satisfy  $\alpha_1 > \alpha_0$ , we have that

$$\phi(\alpha_1) = \phi(\alpha_0) + \phi'(\xi)(\alpha_1 - \alpha_0)\tag{A.54}$$

for some  $\xi \in (\alpha_0, \alpha_1)$ . An extension of this result to a multivariate function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is that for any vector  $p$  we have

$$f(x + p) = f(x) + \nabla f(x + \alpha p)^T p,\tag{A.55}$$

for some  $\alpha \in (0, 1)$ . (This result can be proved by defining  $\phi(\alpha) = f(x + \alpha p)$ ,  $\alpha_0 = 0$ , and  $\alpha_1 = 1$  and applying the chain rule, as above.)

---

#### □ EXAMPLE A.2

Consider  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$  defined by  $f(x) = x_1^3 + 3x_1x_2^2$ , and let  $x = (0, 0)^T$  and  $p = (1, 2)^T$ . It is easy to verify that  $f(x) = 0$  and  $f(x + p) = 13$ . Since

$$\nabla f(x + \alpha p) = \begin{bmatrix} 3(x_1 + \alpha p_1)^2 + 3(x_2 + \alpha p_2)^2 \\ 6(x_1 + \alpha p_1)(x_2 + \alpha p_2) \end{bmatrix} = \begin{bmatrix} 15\alpha^2 \\ 12\alpha^2 \end{bmatrix},$$

we have that  $\nabla f(x + \alpha p)^T p = 39\alpha^2$ . Hence the relation (A.55) holds when we set  $\alpha = 1/\sqrt{13}$ , which lies in the open interval  $(0, 1)$ , as claimed. □

---

An alternative expression to (A.55) can be stated for twice differentiable functions: We have

$$f(x + p) = f(x) + \nabla f(x)^T p + \frac{1}{2} p^T \nabla^2 f(x + \alpha p)^T p, \quad (\text{A.56})$$

for some  $\alpha \in (0, 1)$ . In fact, this expression is one form of Taylor's theorem, Theorem 2.1 in Chapter 2, to which we refer throughout the book.

The extension of (A.55) to a vector-valued function  $r : \mathbb{R}^n \rightarrow \mathbb{R}^m$  for  $m > 1$  is not immediate. There is in general no scalar  $\alpha$  such that the natural extension of (A.55) is satisfied. However, the following result is often a useful analog. As in (10.3), we denote the Jacobian of  $r(x)$ , by  $J(x)$ , where  $J(x)$  is the  $m \times n$  matrix whose  $(j, i)$  entry is  $\partial r_j / \partial x_i$ , for  $j = 1, 2, \dots, m$  and  $i = 1, 2, \dots, n$ , and assume that  $J(x)$  is defined and continuous on the domain of interest. Given  $x$  and  $p$ , we then have

$$r(x + p) - r(x) = \int_0^1 J(x + \alpha p) p \, d\alpha. \quad (\text{A.57})$$

When  $p$  is sufficiently small in norm, we can approximate the right-hand side of this expression adequately by  $J(x)p$ , that is,

$$r(x + p) - r(x) \approx J(x)p.$$

If  $J$  is Lipschitz continuous in the vicinity of  $x$  and  $x + p$  with Lipschitz constant  $L$ , we can use (A.12) to estimate the error in this approximation as follows:

$$\begin{aligned} \|r(x + p) - r(x) - J(x)p\| &= \left\| \int_0^1 [J(x + \alpha p) - J(x)] p \, d\alpha \right\| \\ &\leq \int_0^1 \|J(x + \alpha p) - J(x)\| \|p\| \, d\alpha \\ &\leq \int_0^1 L\alpha \|p\|^2 \, d\alpha = \frac{1}{2} L \|p\|^2. \end{aligned}$$

### IMPLICIT FUNCTION THEOREM

The implicit function theorem lies behind a number of important results in local convergence theory of optimization algorithms and in the characterization of optimality (see Chapter 12). Our statement of this result is based on Lang [187, p. 131] and Bertsekas [19, Proposition A.25].

**Theorem A.2** (Implicit Function Theorem).

Let  $h : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$  be a function such that

- (i)  $h(z^*, 0) = 0$  for some  $z^* \in \mathbb{R}^n$ ,
- (ii) the function  $h(\cdot, \cdot)$  is continuously differentiable in some neighborhood of  $(z^*, 0)$ , and
- (iii)  $\nabla_z h(z, t)$  is nonsingular at the point  $(z, t) = (z^*, 0)$ .

Then there exist open sets  $\mathcal{N}_z \subset \mathbb{R}^n$  and  $\mathcal{N}_t \subset \mathbb{R}^m$  containing  $z^*$  and  $0$ , respectively, and a continuous function  $z : \mathcal{N}_t \rightarrow \mathcal{N}_z$  such that  $z^* = z(0)$  and  $h(z(t), t) = 0$  for all  $t \in \mathcal{N}_t$ . Further,  $z(t)$  is uniquely defined. Finally, if  $h$  is  $p$  times continuously differentiable with respect to both its arguments for some  $p > 0$ , then  $z(t)$  is also  $p$  times continuously differentiable with respect to  $t$ , and we have

$$\nabla z(t) = -\nabla_t h(z(t), t) [\nabla_z h(z(t), t)]^{-1}$$

for all  $t \in \mathcal{N}_t$ .

This theorem is frequently applied to parametrized systems of linear equations, in which  $z$  is obtained as the solution of

$$M(t)z = g(t),$$

where  $M(\cdot) \in \mathbb{R}^{n \times n}$  has  $M(0)$  nonsingular, and  $g(\cdot) \in \mathbb{R}^n$ . To apply the theorem, we define

$$h(z, t) = M(t)z - g(t).$$

If  $M(\cdot)$  and  $g(\cdot)$  are continuously differentiable in some neighborhood of  $0$ , the theorem implies that  $z(t) = M(t)^{-1}g(t)$  is a continuous function of  $t$  in some neighborhood of  $0$ .

## ORDER NOTATION

In much of our analysis we are concerned with how the members of a sequence behave *eventually*, that is, when we get far enough along in the sequence. For instance, we might ask whether the elements of the sequence are bounded, or whether they are similar in size to the elements of a corresponding sequence, or whether they are decreasing and, if so, how rapidly. *Order notation* is useful shorthand to use when questions like these are being examined. It saves us defining many constants that clutter up the argument and the analysis.

We will use three varieties of order notation:  $O(\cdot)$ ,  $o(\cdot)$ , and  $\Omega(\cdot)$ . Given two nonnegative infinite sequences of scalars  $\{\eta_k\}$  and  $\{\nu_k\}$ , we write

$$\eta_k = O(\nu_k)$$

if there is a positive constant  $C$  such that

$$|\eta_k| \leq C|v_k|$$

for all  $k$  sufficiently large. We write

$$\eta_k = o(v_k)$$

if the sequence of ratios  $\{\eta_k/v_k\}$  approaches zero, that is,

$$\lim_{k \rightarrow \infty} \frac{\eta_k}{v_k} = 0.$$

Finally, we write

$$\eta_k = \Omega(v_k)$$

if there are two constants  $C_0$  and  $C_1$  with  $0 < C_0 \leq C_1 < \infty$  such that

$$C_0|v_k| \leq |\eta_k| \leq C_1|v_k|,$$

that is, the corresponding elements of both sequences stay in the same ballpark for all  $k$ . This definition is equivalent to saying that  $\eta_k = O(v_k)$  and  $v_k = O(\eta_k)$ .

The same notation is often used in the context of quantities that depend continuously on each other as well. For instance, if  $\eta(\cdot)$  is a function that maps  $\mathbb{R}$  to  $\mathbb{R}$ , we write

$$\eta(v) = O(v)$$

if there is a constant  $C$  such that  $|\eta(v)| \leq C|v|$  for all  $v \in \mathbb{R}$ . (Typically, we are interested only in values of  $v$  that are either very large or very close to zero; this should be clear from the context. Similarly, we use

$$\eta(v) = o(v) \tag{A.58}$$

to indicate that the ratio  $\eta(v)/v$  approaches zero either as  $v \rightarrow 0$  or  $v \rightarrow \infty$ . (Again, the precise meaning should be clear from the context.)

As a slight variant on the definitions above, we write

$$\eta_k = O(1)$$

to indicate that there is a constant  $C$  such that  $|\eta_k| \leq C$  for all  $k$ , while

$$\eta_k = o(1)$$

indicates that  $\lim_{k \rightarrow \infty} \eta_k = 0$ . We sometimes use vector and matrix quantities as arguments, and in these cases the definitions above are intended to apply to the norms of these quantities. For instance, if  $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ , we write  $f(x) = O(\|x\|)$  if there is a constant  $C > 0$  such that  $\|f(x)\| \leq C\|x\|$  for all  $x$  in the domain of  $f$ . Typically, as above, we are interested only in some subdomain of  $f$ , usually a small neighborhood of 0. As before, the precise meaning should be clear from the context.

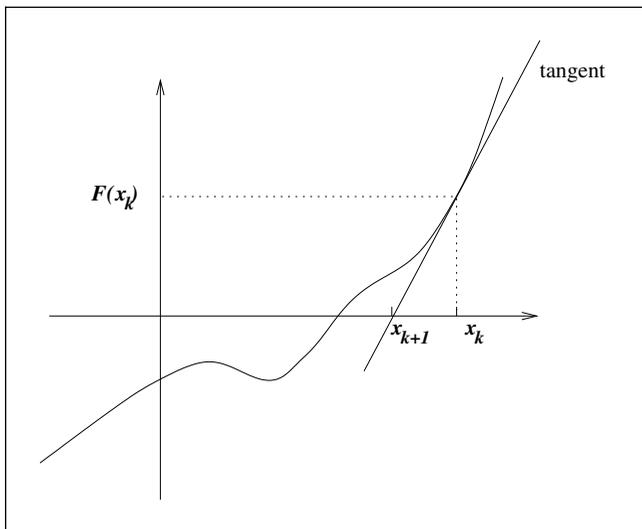
### ROOT-FINDING FOR SCALAR EQUATIONS

In Chapter 11 we discussed methods for finding solutions of nonlinear systems of equations  $F(x) = 0$ , where  $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ . Here we discuss briefly the case of scalar equations ( $n = 1$ ), for which the algorithm is easy to illustrate. Scalar root-finding is needed in the trust-region algorithms of Chapter 4, for instance. Of course, the general theorems of Chapter 11 can be applied to derive rigorous convergence results for this special case.

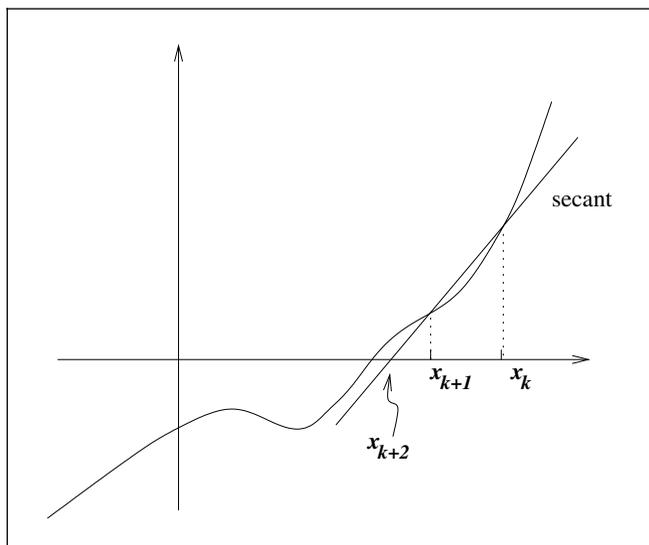
The basic step of Newton's method (Algorithm Newton of Chapter 11) in the scalar case is simply

$$p_k = -F(x_k)/F'(x_k), \quad x_{k+1} \leftarrow x_k + p_k \quad (\text{A.59})$$

(cf. (11.6)). Graphically, such a step involves taking the tangent to the graph of  $F$  at the point  $x_k$  and taking the next iterate to be the intersection of this tangent with the  $x$  axis (see Figure A.2). Clearly, if the function  $F$  is nearly linear, the tangent will be quite a good approximation to  $F$  itself, so the Newton iterate will be quite close to the true root of  $F$ .



**Figure A.2** One step of Newton's method for a scalar equation.



**Figure A.3** One step of the secant method for a scalar equation.

The secant method for scalar equations can be viewed as the specialization of Broyden's method to the case of  $n = 1$ . The issues are simpler in this case, however, since the secant equation (11.27) completely determines the value of the  $1 \times 1$  approximate Hessian  $B_k$ . That is, we do not need to apply extra conditions to ensure that  $B_k$  is fully determined. By combining (11.24) with (11.27), we find that the secant method for the case of  $n = 1$  is defined by

$$B_k = (F(x_k) - F(x_{k-1})) / (x_k - x_{k-1}), \quad (\text{A.60a})$$

$$p_k = -F(x_k) / B_k, \quad x_{k+1} = x_k + p_k. \quad (\text{A.60b})$$

By illustrating this algorithm, we see the origin of the term “secant.”  $B_k$  approximates the slope of the function at  $x_k$  by taking the secant through the points  $(x_{k-1}, F(x_{k-1}))$  and  $(x_k, F(x_k))$ , and  $x_{k+1}$  is obtained by finding the intersection of this secant with the  $x$  axis. The method is illustrated in Figure A.3.

# APPENDIX *B*

## A Regularization Procedure

The following algorithm chooses parameters  $\delta, \gamma$  that guarantee that the regularized primal-dual matrix (19.25) is nonsingular and satisfies the inertia condition (19.24). The algorithm assumes that, at the beginning of the interior-point iteration,  $\delta_{old}$  has been initialized to zero.

**Algorithm B.1** (Inertia Correction and Regularization).

Given the current barrier parameter  $\mu$ , constants  $\eta > 0$  and  $\beta < 1$ , and the perturbation  $\delta_{old}$  used in the previous interior-point iteration.

```

Factor (19.25) with  $\delta = \gamma = 0$ .
if (19.25) is nonsingular and its inertia is  $(n + m, l + m, 0)$ 
    compute the primal-dual step; stop;
if (19.25) has zero eigenvalues
    set  $\gamma \leftarrow 10^{-8}\eta\mu^\beta$ ;
if  $\delta_{old} = 0$ 
    set  $\delta \leftarrow 10^{-4}$ ;

else
    set  $\delta \leftarrow \delta_{old}/2$ ;
repeat
    Factor the modified matrix (19.25);
    if the inertia is  $(n + m, l + m, 0)$ 
        Set  $\delta_{old} \leftarrow \delta$ ;
        Compute the primal-dual step (19.12) using the coefficient
            matrix (19.25); stop;
    else
        Set  $\delta \leftarrow 10\delta$ ;
end (repeat)

```

This algorithm has been adapted from a more elaborate procedure described by Wächter and Biegler [301]. All constants used in the algorithm are arbitrary; we have provided typical choices. The algorithm aims to avoid unnecessarily large modifications  $\delta I$  of  $\nabla_{xx}^2 \mathcal{L}$  while trying to minimize the number of matrix factorizations. Excessive modifications degrade the performance of the algorithm because they erase the second derivative information contained in  $\nabla_{xx}^2 \mathcal{L}$ , and cause the step to take on steepest-descent like characteristics. The first trial value ( $\delta = \delta_{old}/2$ ) is based on the previous modification  $\delta_{old}$  because the minimum perturbation  $\delta$  required to achieve the desired inertia will often not vary much from one interior-point iteration to the next.

The heuristics implemented in Algorithm B.1 provide an alternative to those employed in Algorithm 7.3, which were presented in the context of unconstrained optimization. We emphasize, however, that all of these are indeed heuristics and may not always provide adequate safeguards.

# References

- [1] R. K. AHUJA, T. L. MAGNANTI, AND J. B. ORLIN, *Network Flows: Theory, Algorithms, and Applications*, Prentice-Hall, Englewood Cliffs, N.J., 1993.
- [2] H. AKAIKE, *On a successive transformation of probability distribution and its application to the analysis of the optimum gradient method*, *Annals of the Institute of Statistical Mathematics*, 11 (1959), pp. 1–17.
- [3] M. AL-BAALI, *Descent property and global convergence of the Fletcher-Reeves method with inexact line search*, *I.M.A. Journal on Numerical Analysis*, 5 (1985), pp. 121–124.
- [4] E. D. ANDERSEN AND K. D. ANDERSEN, *Presolving in linear programming*, *Mathematical Programming*, 71 (1995), pp. 221–245.
- [5] ———, *The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm*, in *High Performance Optimization*, T. T. H. Frenk, K. Roos and S. Zhang, eds., Kluwer Academic Publishers, 2000, pp. 197–232.
- [6] E. D. ANDERSEN, J. GONDZIO, C. MÉSZÁROS, AND X. XU, *Implementation of interior-point methods for large scale linear programming*, in *Interior Point Methods in Mathematical Programming*, T. Terlaky, ed., Kluwer, 1996, ch. 6, pp. 189–252.
- [7] E. ANDERSON, Z. BAI, C. BISCHOF, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, S. OSTROUCHOV, AND D. SORENSEN, *LAPACK User's Guide*, SIAM, Philadelphia, 1992.

- [8] M. ANITESCU, *On solving mathematical programs with complementarity constraints as nonlinear programs*, SIAM Journal on Optimization, 15 (2005), pp. 1203–1236.
- [9] ARKI CONSULTING AND DEVELOPMENT A/S, *CONOPT version 3*, 2004.
- [10] B. M. AVERICK, R. G. CARTER, J. J. MORÉ, AND G. XUE, *The MINPACK-2 test problem collection*, Preprint MCS-P153-0692, Argonne National Laboratory, 1992.
- [11] P. BAPTIST AND J. STOER, *On the relation between quadratic termination and convergence properties of minimization algorithms, Part II: Applications*, Numerische Mathematik, 28 (1977), pp. 367–392.
- [12] R. H. BARTELS AND G. H. GOLUB, *The simplex method of linear programming using LU decomposition*, Communications of the ACM, 12 (1969), pp. 266–268.
- [13] R. BARTLETT AND L. BIEGLER, *rSQP++: An object-oriented framework for successive quadratic programming*, in Large-Scale PDE-Constrained Optimization, L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, eds., vol. 30, New York, 2003, Springer-Verlag, pp. 316–330. Lecture Notes in Computational Science and Engineering.
- [14] M. BAZARAA, H. SHERALI, AND C. SHETTY, *Nonlinear Programming, Theory and Applications*, John Wiley & Sons, New York, second ed., 1993.
- [15] A. BEN-TAL AND A. NEMIROVSKI, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, MPS-SIAM Series on Optimization, SIAM, 2001.
- [16] H. Y. BENSON, A. SEN, D. F. SHANNO, AND R. J. VANDERBEL, *Interior-point algorithms, penalty methods and equilibrium problems*, Technical Report ORFE-03-02, Operations Research and Financial Engineering, Princeton University, 2003.
- [17] S. BENSON AND J. MORÉ, *A limited-memory variable-metric algorithm for bound constrained minimization*, Numerical Analysis Report P909-0901, ANL, Argonne, IL, USA, 2001.
- [18] D. P. BERTSEKAS, *Constrained Optimization and Lagrange Multiplier Methods*, Academic Press, New York, 1982.
- [19] ———, *Nonlinear Programming*, Athena Scientific, Belmont, MA, second ed., 1999.
- [20] M. BERZ, C. BISCHOF, C. F. CORLISS, AND A. GRIEWANK, eds., *Computational Differentiation: Techniques, Applications, and Tools*, SIAM Publications, Philadelphia, PA, 1996.
- [21] J. BETTS, S. K. ELDERSVELD, P. D. FRANK, AND J. G. LEWIS, *An interior-point nonlinear programming algorithm for large scale optimization*, Technical report MCT TECH-003, Mathematics and Computing Technology, The Boeing Company, P.O. Box 3707, Seattle, WA 98124-2207, 2000.
- [22] J. R. BIRGE AND F. LOUVEAUX, *Introduction to Stochastic Programming*, Springer-Verlag, New York, 1997.
- [23] E. G. BIRGIN, J. M. MARTINEZ, AND M. RAYDAN, *Algorithm 813: SPG software for convex-constrained optimization*, ACM Transactions on Mathematical Software, 27 (2001), pp. 340–349.
- [24] C. BISCHOF, A. BOUARICHA, P. KHADEMI, AND J. J. MORÉ, *Computing gradients in large-scale optimization using automatic differentiation*, INFORMS Journal on Computing, 9 (1997), pp. 185–194.
- [25] C. BISCHOF, A. CARLE, P. KHADEMI, AND A. MAUER, *ADIFOR 2.0: Automatic differentiation of FORTRAN 77 programs*, IEEE Computational Science & Engineering, 3 (1996), pp. 18–32.
- [26] C. BISCHOF, G. CORLISS, AND A. GRIEWANK, *Structured second- and higher-order derivatives through univariate Taylor series*, Optimization Methods and Software, 2 (1993), pp. 211–232.
- [27] C. BISCHOF, P. KHADEMI, A. BOUARICHA, AND A. CARLE, *Efficient computation of gradients and Jacobians by transparent exploitation of sparsity in automatic differentiation*, Optimization Methods and Software, 7 (1996), pp. 1–39.

- [28] C. BISCHOF, L. ROH, AND A. MAUER, *ADIC: An extensible automatic differentiation tool for ANSI-C*, Software—Practice and Experience, 27 (1997), pp. 1427–1456.
- [29] Å. BJÖRCK, *Numerical Methods for Least Squares Problems*, SIAM Publications, Philadelphia, PA, 1996.
- [30] P. T. BOGGS, R. H. BYRD, AND R. B. SCHNABEL, *A stable and efficient algorithm for nonlinear orthogonal distance regression*, SIAM Journal on Scientific and Statistical Computing, 8 (1987), pp. 1052–1078.
- [31] P. T. BOGGS, J. R. DONALDSON, R. H. BYRD, AND R. B. SCHNABEL, *ODRPACK—Software for weighted orthogonal distance regression*, ACM Transactions on Mathematical Software, 15 (1981), pp. 348–364.
- [32] P. T. BOGGS AND J. W. TOLLE, *Convergence properties of a class of rank-two updates*, SIAM Journal on Optimization, 4 (1994), pp. 262–287.
- [33] ———, *Sequential quadratic programming*, Acta Numerica, 4 (1996), pp. 1–51.
- [34] P. T. BOGGS, J. W. TOLLE, AND P. WANG, *On the local convergence of quasi-Newton methods for constrained optimization*, SIAM Journal on Control and Optimization, 20 (1982), pp. 161–171.
- [35] I. BONGARTZ, A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *CUTE: Constrained and unconstrained testing environment*, Research Report, IBM T.J. Watson Research Center, Yorktown Heights, NY, 1993.
- [36] J. F. BONNANS, E. R. PANIER, A. L. TITS, AND J. L. ZHOU, *Avoiding the Maratos effect by means of a nonmonotone line search. II. Inequality constrained problems — feasible iterates*, SIAM Journal on Numerical Analysis, 29 (1992), pp. 1187–1202.
- [37] S. BOYD, L. EL GHAOU, E. FERON, AND V. BALAKRISHNAN, *Linear Matrix Inequalities in Systems and Control Theory*, SIAM Publications, Philadelphia, 1994.
- [38] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, Cambridge, 2003.
- [39] R. P. BRENT, *Algorithms for minimization without derivatives*, Prentice Hall, Englewood Cliffs, NJ, 1973.
- [40] H. M. BÜCKER, G. F. CORLISS, P. D. HOVLAND, U. NAUMANN, AND B. NORRIS, eds., *Automatic Differentiation: Applications, Theory, and Implementations*, vol. 50 of Lecture Notes in Computational Science and Engineering, Springer, New York, 2005.
- [41] R. BULIRSCH AND J. STOER, *Introduction to Numerical Analysis*, Springer-Verlag, New York, 1980.
- [42] J. R. BUNCH AND L. KAUFMAN, *Some stable methods for calculating inertia and solving symmetric linear systems*, Mathematics of Computation, 31 (1977), pp. 163–179.
- [43] J. R. BUNCH AND B. N. PARLETT, *Direct methods for solving symmetric indefinite systems of linear equations*, SIAM Journal on Numerical Analysis, 8 (1971), pp. 639–655.
- [44] J. V. BURKE AND J. J. MORÉ, *Exposing constraints*, SIAM Journal on Optimization, 4 (1994), pp. 573–595.
- [45] W. BURMEISTER, *Die konvergenzordnung des Fletcher-Powell algorithmus*, Zeitschrift für Angewandte Mathematik und Mechanik, 53 (1973), pp. 693–699.
- [46] R. BYRD, J. NOCEDAL, AND R. WALTZ, *Knitro: An integrated package for nonlinear optimization*, Technical Report 18, Optimization Technology Center, Evanston, IL, June 2005.
- [47] R. BYRD, J. NOCEDAL, AND R. A. WALTZ, *Steering exact penalty methods*, Technical Report OTC 2004/07, Optimization Technology Center, Northwestern University, Evanston, IL, USA, April 2004.
- [48] R. H. BYRD, J.-C. GILBERT, AND J. NOCEDAL, *A trust region method based on interior point techniques for nonlinear programming*, Mathematical Programming, 89 (2000), pp. 149–185.

- [49] R. H. BYRD, N. I. M. GOULD, J. NOCEDAL, AND R. A. WALTZ, *An algorithm for nonlinear optimization using linear programming and equality constrained subproblems*, *Mathematical Programming, Series B*, 100 (2004), pp. 27–48.
- [50] R. H. BYRD, M. E. HRIBAR, AND J. NOCEDAL, *An interior point method for large scale nonlinear programming*, *SIAM Journal on Optimization*, 9 (1999), pp. 877–900.
- [51] R. H. BYRD, H. F. KHALFAN, AND R. B. SCHNABEL, *Analysis of a symmetric rank-one trust region method*, *SIAM Journal on Optimization*, 6 (1996), pp. 1025–1039.
- [52] R. H. BYRD, J. NOCEDAL, AND R. B. SCHNABEL, *Representations of quasi-Newton matrices and their use in limited-memory methods*, *Mathematical Programming, Series A*, 63 (1994), pp. 129–156.
- [53] R. H. BYRD, J. NOCEDAL, AND Y. YUAN, *Global convergence of a class of quasi-Newton methods on convex problems*, *SIAM Journal on Numerical Analysis*, 24 (1987), pp. 1171–1190.
- [54] R. H. BYRD, R. B. SCHNABEL, AND G. A. SCHULTZ, *Approximate solution of the trust regions problem by minimization over two-dimensional subspaces*, *Mathematical Programming*, 40 (1988), pp. 247–263.
- [55] R. H. BYRD, R. B. SCHNABEL, AND G. A. SHULTZ, *A trust region algorithm for nonlinearly constrained optimization*, *SIAM Journal on Numerical Analysis*, 24 (1987), pp. 1152–1170.
- [56] M. R. CELIS, J. E. DENNIS, AND R. A. TAPIA, *A trust region strategy for nonlinear equality constrained optimization*, in *Numerical Optimization*, P. T. Boggs, R. H. Byrd, and R. B. Schnabel, eds., SIAM, 1985, pp. 71–82.
- [57] R. CHAMBERLAIN, C. LEMARECHAL, H. C. PEDERSEN, AND M. J. D. POWELL, *The watchdog technique for forcing convergence in algorithms for constrained optimization*, *Mathematical Programming*, 16 (1982), pp. 1–17.
- [58] S. H. CHENG AND N. J. HIGHAM, *A modified Cholesky algorithm based on a symmetric indefinite factorization*, *SIAM Journal of Matrix Analysis and Applications*, 19 (1998), pp. 1097–1100.
- [59] C. M. CHIN AND R. FLETCHER, *On the global convergence of an SLP-filter algorithm that takes EQP steps*, *Mathematical Programming, Series A*, 96 (2003), pp. 161–177.
- [60] T. D. CHOI AND C. T. KELLEY, *Superlinear convergence and implicit filtering*, *SIAM Journal on Optimization*, 10 (2000), pp. 1149–1162.
- [61] V. CHVÁTAL, *Linear Programming*, W. H. Freeman and Company, New York, 1983.
- [62] F. H. CLARKE, *Optimization and Nonsmooth Analysis*, John Wiley & Sons, New York, 1983 (Reprinted by SIAM Publications, 1990).
- [63] A. COHEN, *Rate of convergence of several conjugate gradient algorithms*, *SIAM Journal on Numerical Analysis*, 9 (1972), pp. 248–259.
- [64] T. F. COLEMAN, *Linearly constrained optimization and projected preconditioned conjugate gradients*, in *Proceedings of the Fifth SIAM Conference on Applied Linear Algebra*, J. Lewis, ed., Philadelphia, USA, 1994, SIAM, pp. 118–122.
- [65] T. F. COLEMAN AND A. R. CONN, *Non-linear programming via an exact penalty-function: Asymptotic analysis*, *Mathematical Programming*, 24 (1982), pp. 123–136.
- [66] T. F. COLEMAN, B. GARBOW, AND J. J. MORÉ, *Software for estimating sparse Jacobian matrices*, *ACM Transactions on Mathematical Software*, 10 (1984), pp. 329–345.
- [67] ———, *Software for estimating sparse Hessian matrices*, *ACM Transactions on Mathematical Software*, 11 (1985), pp. 363–377.
- [68] T. F. COLEMAN AND J. J. MORÉ, *Estimation of sparse Jacobian matrices and graph coloring problems*, *SIAM Journal on Numerical Analysis*, 20 (1983), pp. 187–209.
- [69] ———, *Estimation of sparse Hessian matrices and graph coloring problems*, *Mathematical Programming*, 28 (1984), pp. 243–270.

- [70] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *Testing a class of algorithms for solving minimization problems with simple bounds on the variables*, Mathematics of Computation, 50 (1988), pp. 399–430.
- [71] ———, *Convergence of quasi-Newton matrices generated by the symmetric rank one update*, Mathematical Programming, 50 (1991), pp. 177–195.
- [72] ———, *LANCELOT: a FORTRAN package for large-scale nonlinear optimization (Release A)*, no. 17 in Springer Series in Computational Mathematics, Springer-Verlag, New York, 1992.
- [73] ———, *Numerical experiments with the LANCELOT package (Release A) for large-scale nonlinear optimization*, Report 92/16, Department of Mathematics, University of Namur, Belgium, 1992.
- [74] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *Trust-Region Methods*, MPS-SIAM Series on Optimization, SIAM, 2000.
- [75] A. R. CONN, K. SCHEINBERG, AND P. L. TOINT, *On the convergence of derivative-free methods for unconstrained optimization*, in Approximation Theory and Optimization: Tributes to M. J. D. Powell, A. Iserles and M. Buhmann, eds., Cambridge University Press, Cambridge, UK, 1997, pp. 83–108.
- [76] ———, *Recent progress in unconstrained nonlinear optimization without derivatives*, Mathematical Programming, Series B, 79 (1997), pp. 397–414.
- [77] W. J. COOK, W. H. CUNNINGHAM, W. R. PULLEYBLANK, AND A. SCHRIJVER, *Combinatorial Optimization*, John Wiley & Sons, New York, 1997.
- [78] B. F. CORLISS AND L. B. RALL, *An introduction to automatic differentiation*, in Computational Differentiation: Techniques, Applications, and Tools, M. Berz, C. Bischof, G. F. Corliss, and A. Griewank, eds., SIAM Publications, Philadelphia, PA, 1996, ch. 1.
- [79] T. H. CORMEN, C. E. LEISSERSON, AND R. L. RIVEST, *Introduction to Algorithms*, MIT Press, 1990.
- [80] R. W. COTTLE, J.-S. PANG, AND R. E. STONE, *The Linear Complementarity Problem*, Academic Press, San Diego, 1992.
- [81] R. COURANT, *Variational methods for the solution of problems with equilibrium and vibration*, Bulletin of the American Mathematical Society, 49 (1943), pp. 1–23.
- [82] H. P. CROWDER AND P. WOLFE, *Linear convergence of the conjugate gradient method*, IBM Journal of Research and Development, 16 (1972), pp. 431–433.
- [83] A. CURTIS, M. J. D. POWELL, AND J. REID, *On the estimation of sparse Jacobian matrices*, Journal of the Institute of Mathematics and its Applications, 13 (1974), pp. 117–120.
- [84] J. CZYZYK, S. MEHROTRA, M. WAGNER, AND S. J. WRIGHT, *PCx: An interior-point code for linear programming*, Optimization Methods and Software, 11/12 (1999), pp. 397–430.
- [85] Y. DAI AND Y. YUAN, *A nonlinear conjugate gradient method with a strong global convergence property*, SIAM Journal on Optimization, 10 (1999), pp. 177–182.
- [86] G. B. DANTZIG, *Linear Programming and Extensions*, Princeton University Press, Princeton, NJ, 1963.
- [87] W. C. DAVIDON, *Variable metric method for minimization*, Technical Report ANL-5990 (revised), Argonne National Laboratory, Argonne, IL, 1959.
- [88] ———, *Variable metric method for minimization*, SIAM Journal on Optimization, 1 (1991), pp. 1–17.
- [89] R. S. DEMBO, S. C. EISENSTAT, AND T. STEIHAUG, *Inexact Newton methods*, SIAM Journal on Numerical Analysis, 19 (1982), pp. 400–408.
- [90] J. E. DENNIS, D. M. GAY, AND R. E. WELSCH, *Algorithm 573 — NL2SOL, An adaptive nonlinear least-squares algorithm*, ACM Transactions on Mathematical Software, 7 (1981), pp. 348–368.

- [91] J. E. DENNIS AND J. J. MORÉ, *Quasi-Newton methods, motivation and theory*, SIAM Review, 19 (1977), pp. 46–89.
- [92] J. E. DENNIS AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Non-linear Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1983. Reprinted by SIAM Publications, 1993.
- [93] J. E. DENNIS AND R. B. SCHNABEL, *A view of unconstrained optimization*, in Optimization, vol. 1 of Handbooks in Operations Research and Management, Elsevier Science Publishers, Amsterdam, The Netherlands, 1989, pp. 1–72.
- [94] I. I. DIKIN, *Iterative solution of problems of linear and quadratic programming*, Soviet Mathematics-Doklady, 8 (1967), pp. 674–675.
- [95] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear equations*, ACM Transactions on Mathematical Software, 9 (1983), pp. 302–325.
- [96] I. S. DUFF AND J. K. REID, *The design of MA48: A code for the direct solution of sparse unsymmetric linear systems of equations*, ACM Transactions on Mathematical Software, 22 (1996), pp. 187–226.
- [97] I. S. DUFF, J. K. REID, N. MUNKSGAARD, AND H. B. NEILSEN, *Direct solution of sets of linear equations whose matrix is sparse symmetric and indefinite*, Journal of the Institute of Mathematics and its Applications, 23 (1979), pp. 235–250.
- [98] A. V. FIACCO AND G. P. MCCORMICK, *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*, John Wiley & Sons, New York, N.Y., 1968. Reprinted by SIAM Publications, 1990.
- [99] R. FLETCHER, *A general quadratic programming algorithm*, Journal of the Institute of Mathematics and its Applications, 7 (1971), pp. 76–91.
- [100] ———, *Second order corrections for non-differentiable optimization*, in Numerical Analysis, D. Griffiths, ed., Springer Verlag, 1982, pp. 85–114. Proceedings Dundee 1981.
- [101] ———, *Practical Methods of Optimization*, John Wiley & Sons, New York, second ed., 1987.
- [102] ———, *An optimal positive definite update for sparse Hessian matrices*, SIAM Journal on Optimization, 5 (1995), pp. 192–218.
- [103] ———, *Stable reduced hessian updates for indefinite quadratic programming*, Mathematical Programming, 87 (2000), pp. 251–264.
- [104] R. FLETCHER, A. GROTHEY, AND S. LEYFFER, *Computing sparse Hessian and Jacobian approximations with optimal hereditary properties*, technical report, Department of Mathematics, University of Dundee, 1996.
- [105] R. FLETCHER AND S. LEYFFER, *Nonlinear programming without a penalty function*, Mathematical Programming, Series A, 91 (2002), pp. 239–269.
- [106] R. FLETCHER, S. LEYFFER, AND P. L. TOINT, *On the global convergence of an SLP-filter algorithm*, Numerical Analysis Report NA/183, Dundee University, Dundee, Scotland, UK, 1999.
- [107] R. FLETCHER AND C. M. REEVES, *Function minimization by conjugate gradients*, Computer Journal, 7 (1964), pp. 149–154.
- [108] R. FLETCHER AND E. SAINZ DE LA MAZA, *Nonlinear programming and nonsmooth optimization by successive linear programming*, Mathematical Programming, 43 (1989), pp. 235–256.
- [109] C. FLOUDAS AND P. PARDALOS, eds., *Recent Advances in Global Optimization*, Princeton University Press, Princeton, NJ, 1992.
- [110] J. J. H. FORREST AND J. A. TOMLIN, *Updated triangular factors of the basis to maintain sparsity in the product form simplex method*, Mathematical Programming, 2 (1972), pp. 263–278.

- [111] A. FORSGREN, P. E. GILL, AND M. H. WRIGHT, *Interior methods for nonlinear optimization*, SIAM Review, 44 (2003), pp. 525–597.
- [112] R. FOURER, D. M. GAY, AND B. W. KERNIGHAN, *AMPL: A Modeling Language for Mathematical Programming*, The Scientific Press, South San Francisco, CA, 1993.
- [113] R. FOURER AND S. MEHROTRA, *Solving symmetric indefinite systems in an interior-point method for linear programming*, Mathematical Programming, 62 (1993), pp. 15–39.
- [114] M. P. FRIEDLANDER AND M. A. SAUNDERS, *A globally convergent linearly constrained Lagrangian method for nonlinear optimization*, SIAM Journal on Optimization, 15 (2005), pp. 863–897.
- [115] K. R. FRISCH, *The logarithmic potential method of convex programming*, Technical Report, University Institute of Economics, Oslo, Norway, 1955.
- [116] D. GABAY, *Reduced quasi-Newton methods with feasibility improvement for nonlinearly constrained optimization*, Mathematical Programming Studies, 16 (1982), pp. 18–44.
- [117] U. M. GARCIA-PALOMARES AND O. L. MANGASARIAN, *Superlinearly convergent quasi-Newton methods for nonlinearly constrained optimization problems*, Mathematical Programming, 11 (1976), pp. 1–13.
- [118] D. M. GAY, *More AD of nonlinear AMPL models: computing Hessian information and exploiting partial separability*, in Computational Differentiation: Techniques, Applications, and Tools, M. Berz, C. Bischof, G. F. Corliss, and A. Griewank, eds., SIAM Publications, Philadelphia, PA, 1996, pp. 173–184.
- [119] R.-P. GE AND M. J. D. POWELL, *The convergence of variable metric matrices in unconstrained optimization*, Mathematical Programming, 27 (1983), pp. 123–143.
- [120] A. H. GEBREMEDHIN, F. MANNE, AND A. POTHEN, *What color is your Jacobian? Graph coloring for computing derivatives*, SIAM Review, 47 (2005), pp. 629–705.
- [121] E. M. GERTZ AND S. J. WRIGHT, *Object-oriented software for quadratic programming*, ACM Transactions on Mathematical Software, 29 (2003), pp. 58–81.
- [122] J. GILBERT AND C. LEMARÉCHAL, *Some numerical experiments with variable-storage quasi-Newton algorithms*, Mathematical Programming, Series B, 45 (1989), pp. 407–435.
- [123] J. GILBERT AND J. NOCEDAL, *Global convergence properties of conjugate gradient methods for optimization*, SIAM Journal on Optimization, 2 (1992), pp. 21–42.
- [124] P. E. GILL, G. H. GOLUB, W. MURRAY, AND M. A. SAUNDERS, *Methods for modifying matrix factorizations*, Mathematics of Computation, 28 (1974), pp. 505–535.
- [125] P. E. GILL AND M. W. LEONARD, *Limited-memory reduced-Hessian methods for unconstrained optimization*, SIAM Journal on Optimization, 14 (2003), pp. 380–401.
- [126] P. E. GILL AND W. MURRAY, *Numerically stable methods for quadratic programming*, Mathematical Programming, 14 (1978), pp. 349–372.
- [127] P. E. GILL, W. MURRAY, AND M. A. SAUNDERS, *User’s guide for SNOPT (Version 5.3): A FORTRAN package for large-scale nonlinear programming*, Technical Report NA 97-4, Department of Mathematics, University of California, San Diego, 1997.
- [128] ———, *SNOPT: An SQP algorithm for large-scale constrained optimization*, SIAM Journal on Optimization, 12 (2002), pp. 979–1006.
- [129] P. E. GILL, W. MURRAY, M. A. SAUNDERS, AND M. H. WRIGHT, *User’s guide for SOL/QPSOL*, Technical Report SOL84–6, Department of Operations Research, Stanford University, Stanford, California, 1984.
- [130] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press, 1981.
- [131] ———, *Numerical Linear Algebra and Optimization, Vol. 1*, Addison Wesley, Redwood City, California, 1991.

- [132] D. GOLDFARB, *Curvilinear path steplength algorithms for minimization which use directions of negative curvature*, *Mathematical Programming*, 18 (1980), pp. 31–40.
- [133] D. GOLDFARB AND J. FORREST, *Steepest edge simplex algorithms for linear programming*, *Mathematical Programming*, 57 (1992), pp. 341–374.
- [134] D. GOLDFARB AND J. K. REID, *A practicable steepest-edge simplex algorithm*, *Mathematical Programming*, 12 (1977), pp. 361–373.
- [135] G. GOLUB AND D. O’LEARY, *Some history of the conjugate gradient methods and the Lanczos algorithms: 1948–1976*, *SIAM Review*, 31 (1989), pp. 50–100.
- [136] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, third ed., 1996.
- [137] J. GONDZIO, *HOPDM (version 2.12): A fast LP solver based on a primal–dual interior point method*, *European Journal of Operations Research*, 85 (1995), pp. 221–225.
- [138] ———, *Multiple centrality corrections in a primal–dual method for linear programming*, *Computational Optimization and Applications*, 6 (1996), pp. 137–156.
- [139] J. GONDZIO AND A. GROTHEY, *Parallel interior point solver for structured quadratic programs: Application to financial planning problems*, Technical Report MS-03-001, School of Mathematics, University of Edinburgh, Scotland, 2003.
- [140] N. I. M. GOULD, *On the accurate determination of search directions for simple differentiable penalty functions*, *I.M.A. Journal on Numerical Analysis*, 6 (1986), pp. 357–372.
- [141] ———, *On the convergence of a sequential penalty function method for constrained minimization*, *SIAM Journal on Numerical Analysis*, 26 (1989), pp. 107–128.
- [142] ———, *An algorithm for large scale quadratic programming*, *I.M.A. Journal on Numerical Analysis*, 11 (1991), pp. 299–324.
- [143] N. I. M. GOULD, M. E. HRIBAR, AND J. NOCEDAL, *On the solution of equality constrained quadratic problems arising in optimization*, *SIAM Journal on Scientific Computing*, 23 (2001), pp. 1375–1394.
- [144] N.I.M. GOULD, S. LEYFFER, AND P. L. TOINT, *A multidimensional filter algorithm for nonlinear equations and nonlinear least squares*, *SIAM Journal on Optimization*, 15 (2004), pp. 17–38.
- [145] N. I. M. GOULD, S. LUCIDI, M. ROMA, AND P. L. TOINT, *Solving the trust-region subproblem using the Lanczos method*, *SIAM Journal on Optimization*, 9 (1999), pp. 504–525.
- [146] N. I. M. GOULD, D. ORBAN, AND P. L. TOINT, **GALAHAD**—*a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization*, *ACM Transactions on Mathematical Software*, 29 (2003), pp. 353–372.
- [147] N. I. M. GOULD, D. ORBAN, AND P. L. TOINT, *Numerical methods for large-scale nonlinear optimization*, *Acta Numerica*, 14 (2005), pp. 299–361.
- [148] N. I. M. GOULD AND P. L. TOINT, *An iterative working-set method for large-scale non-convex quadratic programming*, *Applied Numerical Mathematics*, 43 (2002), pp. 109–128.
- [149] ———, *Numerical methods for large-scale non-convex quadratic programming*, in *Trends in Industrial and Applied Mathematics*, A. H. Siddiqi and M. Kočvara, eds., Dordrecht, The Netherlands, 2002, Kluwer Academic Publishers, pp. 149–179.
- [150] A. GRIEWANK, *Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation*, *Optimization Methods and Software*, 1 (1992), pp. 35–54.
- [151] ———, *Automatic directional differentiation of nonsmooth composite functions*, in *Seventh French-German Conference on Optimization*, 1994.
- [152] A. GRIEWANK, *Evaluating Derivatives: Principles and Techniques of Automatic Differentiation*, vol. 19 of *Frontiers in Applied Mathematics*, SIAM, 2000.

- [153] A. GRIEWANK AND G. F. CORLISS, eds., *Automatic Differentiation of Algorithms*, SIAM Publications, Philadelphia, Penn., 1991.
- [154] A. GRIEWANK, D. JUEDES, AND J. UTKE, *ADOL-C, A package for the automatic differentiation of algorithms written in C/C++*, ACM Transactions on Mathematical Software, 22 (1996), pp. 131–167.
- [155] A. GRIEWANK AND P. L. TOINT, *Local convergence analysis of partitioned quasi-Newton updates*, Numerische Mathematik, 39 (1982), pp. 429–448.
- [156] ———, *Partitioned variable metric updates for large structured optimization problems*, Numerische Mathematik, 39 (1982), pp. 119–137.
- [157] J. GRIMM, L. POTTIER, AND N. ROSTAING-SCHMIDT, *Optimal time and minimum space time product for reversing a certain class of programs*, in Computational Differentiation, Techniques, Applications, and Tools, M. Berz, C. Bischof, G. Corliss, and A. Griewank, eds., SIAM, Philadelphia, 1996, pp. 95–106.
- [158] L. GRIPPO, F. LAMPARIELLO, AND S. LUCIDI, *A nonmonotone line search technique for Newton's method*, SIAM Journal on Numerical Analysis, 23 (1986), pp. 707–716.
- [159] C. GUÉRET, C. PRINS, AND M. SEVAUX, *Applications of optimization with Xpress-MP*, Dash Optimization, 2002.
- [160] W. W. HAGER, *Minimizing a quadratic over a sphere*, SIAM Journal on Optimization, 12 (2001), pp. 188–208.
- [161] W. W. HAGER AND H. ZHANG, *A new conjugate gradient method with guaranteed descent and an efficient line search*, SIAM Journal on Optimization, 16 (2005), pp. 170–192.
- [162] ———, *A survey of nonlinear conjugate gradient methods*. To appear in the Pacific Journal of Optimization, 2005.
- [163] S. P. HAN, *Superlinearly convergent variable metric algorithms for general nonlinear programming problems*, Mathematical Programming, 11 (1976), pp. 263–282.
- [164] ———, *A globally convergent method for nonlinear programming*, Journal of Optimization Theory and Applications, 22 (1977), pp. 297–309.
- [165] S. P. HAN AND O. L. MANGASARIAN, *Exact penalty functions in nonlinear programming*, Mathematical Programming, 17 (1979), pp. 251–269.
- [166] HARWELL SUBROUTINE LIBRARY, *A catalogue of subroutines (release 13)*, AERE Harwell Laboratory, Harwell, Oxfordshire, England, 1998.
- [167] M. R. HESTENES, *Multiplier and gradient methods*, Journal of Optimization Theory and Applications, 4 (1969), pp. 303–320.
- [168] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, Journal of Research of the National Bureau of Standards, 49 (1952), pp. 409–436.
- [169] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, SIAM Publications, Philadelphia, 1996.
- [170] J.-B. HIRIART-URRUTY AND C. LEMARECHAL, *Convex Analysis and Minimization Algorithms*, Springer-Verlag, Berlin, New York, 1993.
- [171] P. HOUGH, T. KOLDA, AND V. TORCZON, *Asynchronous parallel pattern search for nonlinear optimization*, SIAM Journal on Optimization, 23 (2001), pp. 134–156.
- [172] ILOG CPLEX 8.0, *User's Manual*, ILOG SA, Gentilly, France, 2002.
- [173] D. JONES, C. PERTUNEN, AND B. STUCKMAN, *Lipschitzian optimization without the Lipschitz constant*, Journal of Optimization Theory and Applications, 79 (1993), pp. 157–181.
- [174] P. KALL AND S. W. WALLACE, *Stochastic Programming*, John Wiley & Sons, New York, 1994.

- [175] N. KARMAKAR, *A new polynomial-time algorithm for linear programming*, *Combinatorics*, 4 (1984), pp. 373–395.
- [176] C. KELLER, N. I. M. GOULD, AND A. J. WATHEN, *Constraint preconditioning for indefinite linear systems*, *SIAM Journal on Matrix Analysis and Applications*, 21 (2000), pp. 1300–1317.
- [177] C. T. KELLEY, *Iterative Methods for Linear and Nonlinear Equations*, SIAM Publications, Philadelphia, PA, 1995.
- [178] C. T. KELLEY, *Detection and remediation of stagnation in the Nelder-Mead algorithm using a sufficient decrease condition*, *SIAM Journal on Optimization*, 10 (1999), pp. 43–55.
- [179] ———, *Iterative Methods for Optimization*, no. 18 in *Frontiers in Applied Mathematics*, SIAM Publications, Philadelphia, PA, 1999.
- [180] L. G. KHACHYAN, *A polynomial algorithm in linear programming*, *Soviet Mathematics Doklady*, 20 (1979), pp. 191–194.
- [181] H. F. KHALFAN, R. H. BYRD, AND R. B. SCHNABEL, *A theoretical and experimental study of the symmetric rank one update*, *SIAM Journal on Optimization*, 3 (1993), pp. 1–24.
- [182] V. KLEE AND G. J. MINTY, *How good is the simplex algorithm?* in *Inequalities*, O. Shisha, ed., Academic Press, New York, 1972, pp. 159–175.
- [183] T. G. KOLDA, R. M. LEWIS, AND V. TORCZON, *Optimization by direct search: New perspectives on some classical and modern methods*, *SIAM Review*, 45 (2003), pp. 385–482.
- [184] M. KOČVARA AND M. STINGL, *PENNON, a code for nonconvex nonlinear and semidefinite programming*, *Optimization Methods and Software*, 18 (2003), pp. 317–333.
- [185] H. W. KUHN AND A. W. TUCKER, *Nonlinear programming*, in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, J. Neyman, ed., Berkeley, CA, 1951, University of California Press, pp. 481–492.
- [186] J. W. LAGARIAS, J. A. REEDS, M. H. WRIGHT, AND P. E. WRIGHT, *Convergence properties of the Nelder-Mead simplex algorithm in low dimensions*, *SIAM Journal on Optimization*, 9 (1998), pp. 112–147.
- [187] S. LANG, *Real Analysis*, Addison-Wesley, Reading, MA, second ed., 1983.
- [188] C. L. LAWSON AND R. J. HANSON, *Solving Least Squares Problems*, Prentice-Hall, Englewood Cliffs, NJ, 1974.
- [189] C. LEMARÉCHAL, *A view of line searches*, in *Optimization and Optimal Control*, W. Oettli and J. Stoer, eds., no. 30 in *Lecture Notes in Control and Information Science*, Springer-Verlag, 1981, pp. 59–78.
- [190] K. LEVENBERG, *A method for the solution of certain non-linear problems in least squares*, *Quarterly of Applied Mathematics*, 2 (1944), pp. 164–168.
- [191] S. LEYFFER, G. LOPEZ-CALVA, AND J. NOCEDAL, *Interior methods for mathematical programs with complementarity constraints*, technical report 8, Optimization Technology Center, Northwestern University, Evanston, IL, 2004.
- [192] C. LIN AND J. MORÉ, *Newton’s method for large bound-constrained optimization problems*, *SIAM Journal on Optimization*, 9 (1999), pp. 1100–1127.
- [193] C. LIN AND J. J. MORÉ, *Incomplete Cholesky factorizations with limited memory*, *SIAM Journal on Scientific Computing*, 21 (1999), pp. 24–45.
- [194] D. C. LIU AND J. NOCEDAL, *On the limited-memory BFGS method for large scale optimization*, *Mathematical Programming*, 45 (1989), pp. 503–528.
- [195] D. LUENBERGER, *Introduction to Linear and Nonlinear Programming*, Addison Wesley, second ed., 1984.

- [196] L. LUKŠAN AND J. VLČEK, *Indefinitely preconditioned inexact Newton method for large sparse equality constrained nonlinear programming problems*, Numerical Linear Algebra with Applications, 5 (1998), pp. 219–247.
- [197] *Macsyma User's Guide*, second ed., 1996.
- [198] O. L. MANGASARIAN, *Nonlinear Programming*, McGraw-Hill, New York, 1969. Reprinted by SIAM Publications, 1995.
- [199] N. MARATOS, *Exact penalty function algorithms for finite dimensional and control optimization problems*, PhD thesis, University of London, 1978.
- [200] M. MARAZZI AND J. NOCEDAL, *Wedge trust region methods for derivative free optimization*, Mathematical Programming, Series A, 91 (2002), pp. 289–305.
- [201] H. M. MARKOWITZ, *Portfolio selection*, Journal of Finance, 8 (1952), pp. 77–91.
- [202] ———, *The elimination form of the inverse and its application to linear programming*, Management Science, 3 (1957), pp. 255–269.
- [203] D. W. MARQUARDT, *An algorithm for least squares estimation of non-linear parameters*, SIAM Journal, 11 (1963), pp. 431–441.
- [204] D. Q. MAYNE AND E. POLAK, *A superlinearly convergent algorithm for constrained optimization problems*, Mathematical Programming Studies, 16 (1982), pp. 45–61.
- [205] L. MCLINDEN, *An analogue of Moreau's proximation theorem, with applications to the nonlinear complementarity problem*, Pacific Journal of Mathematics, 88 (1980), pp. 101–161.
- [206] N. MEGIDDO, *Pathways to the optimal set in linear programming*, in Progress in Mathematical Programming: Interior-Point and Related Methods, N. Megiddo, ed., Springer-Verlag, New York, NY, 1989, ch. 8, pp. 131–158.
- [207] S. MEHROTRA, *On the implementation of a primal-dual interior point method*, SIAM Journal on Optimization, 2 (1992), pp. 575–601.
- [208] S. MIZUNO, M. TODD, AND Y. YE, *On adaptive step primal-dual interior-point algorithms for linear programming*, Mathematics of Operations Research, 18 (1993), pp. 964–981.
- [209] J. L. MORALES AND J. NOCEDAL, *Automatic preconditioning by limited memory quasi-newton updating*, SIAM Journal on Optimization, 10 (2000), pp. 1079–1096.
- [210] J. J. MORÉ, *The Levenberg-Marquardt algorithm: Implementation and theory*, in Lecture Notes in Mathematics, No. 630—Numerical Analysis, G. Watson, ed., Springer-Verlag, 1978, pp. 105–116.
- [211] ———, *Recent developments in algorithms and software for trust region methods*, in Mathematical Programming: The State of the Art, Springer-Verlag, Berlin, 1983, pp. 258–287.
- [212] ———, *A collection of nonlinear model problems*, in Computational Solution of Nonlinear Systems of Equations, vol. 26 of Lectures in Applied Mathematics, American Mathematical Society, Providence, RI, 1990, pp. 723–762.
- [213] J. J. MORÉ AND D. C. SORENSEN, *On the use of directions of negative curvature in a modified Newton method*, Mathematical Programming, 16 (1979), pp. 1–20.
- [214] ———, *Computing a trust region step*, SIAM Journal on Scientific and Statistical Computing, 4 (1983), pp. 553–572.
- [215] ———, *Newton's method*, in Studies in Numerical Analysis, vol. 24 of MAA Studies in Mathematics, The Mathematical Association of America, 1984, pp. 29–82.
- [216] J. J. MORÉ AND D. J. THUENTE, *Line search algorithms with guaranteed sufficient decrease*, ACM Transactions on Mathematical Software, 20 (1994), pp. 286–307.
- [217] J. J. MORÉ AND S. J. WRIGHT, *Optimization Software Guide*, SIAM Publications, Philadelphia, 1993.

- [218] B. A. MURTAGH AND M. A. SAUNDERS, *MINOS 5.1 User's guide*, Technical Report SOL-83-20R, Stanford University, 1987.
- [219] K. G. MURTY AND S. N. KABADI, *Some NP-complete problems in quadratic and nonlinear programming*, *Mathematical Programming*, 19 (1987), pp. 200–212.
- [220] S. G. NASH, *Newton-type minimization via the Lanczos method*, *SIAM Journal on Numerical Analysis*, 21 (1984), pp. 553–572.
- [221] ———, *SUMT (Revisited)*, *Operations Research*, 46 (1998), pp. 763–775.
- [222] U. NAUMANN, *Optimal accumulation of Jacobian matrices by elimination methods on the dual computational graph*, *Mathematical Programming*, 99 (2004), pp. 399–421.
- [223] J. A. NELDER AND R. MEAD, *A simplex method for function minimization*, *The Computer Journal*, 8 (1965), pp. 308–313.
- [224] G. L. NEMHAUSER AND L. A. WOLSEY, *Integer and Combinatorial Optimization*, John Wiley & Sons, New York, 1988.
- [225] A. S. NEMIROVSKII AND D. B. YUDIN, *Problem complexity and method efficiency*, John Wiley & Sons, New York, 1983.
- [226] Y. E. NESTEROV AND A. S. NEMIROVSKII, *Interior Point Polynomial Methods in Convex Programming*, SIAM Publications, Philadelphia, 1994.
- [227] G. N. NEWSAM AND J. D. RAMSDELL, *Estimation of sparse Jacobian matrices*, *SIAM Journal on Algebraic and Discrete Methods*, 4 (1983), pp. 404–418.
- [228] J. NOCEDAL, *Updating quasi-Newton matrices with limited storage*, *Mathematics of Computation*, 35 (1980), pp. 773–782.
- [229] ———, *Theory of algorithms for unconstrained optimization*, *Acta Numerica*, 1 (1992), pp. 199–242.
- [230] J. M. ORTEGA AND W. C. RHEINOLDT, *Iterative solution of nonlinear equations in several variables*, Academic Press, New York and London, 1970.
- [231] M. R. OSBORNE, *Nonlinear least squares—the Levenberg algorithm revisited*, *Journal of the Australian Mathematical Society, Series B*, 19 (1976), pp. 343–357.
- [232] ———, *Finite Algorithms in Optimization and Data Analysis*, John Wiley & Sons, New York, 1985.
- [233] M. L. OVERTON, *Numerical Computing with IEEE Floating Point Arithmetic*, SIAM, Philadelphia, PA, 2001.
- [234] C. C. PAIGE AND M. A. SAUNDERS, *LSQR: An algorithm for sparse linear equations and sparse least squares*, *ACM Transactions on Mathematical Software*, 8 (1982), pp. 43–71.
- [235] C. H. PAPADIMITRIOU AND K. STEIGLITZ, *Combinatorial Optimization: Algorithms and Complexity*, Prentice Hall, Englewood Cliffs, NJ, 1982.
- [236] E. POLAK, *Optimization: Algorithms and Consistent Approximations*, no. 124 in *Applied Mathematical Sciences*, Springer, 1997.
- [237] E. POLAK AND G. RIBIÈRE, *Note sur la convergence de méthodes de directions conjuguées*, *Revue Française d'Informatique et de Recherche Opérationnelle*, 16 (1969), pp. 35–43.
- [238] B. T. POLYAK, *The conjugate gradient method in extremal problems*, *U.S.S.R. Computational Mathematics and Mathematical Physics*, 9 (1969), pp. 94–112.
- [239] M. J. D. POWELL, *An efficient method for finding the minimum of a function of several variables without calculating derivatives*, *Computer Journal*, 91 (1964), pp. 155–162.
- [240] ———, *A method for nonlinear constraints in minimization problems*, in *Optimization*, R. Fletcher, ed., Academic Press, New York, NY, 1969, pp. 283–298.

- [241] ———, *A hybrid method for nonlinear equations*, in Numerical Methods for Nonlinear Algebraic Equations, P. Rabinowitz, ed., Gordon & Breach, London, 1970, pp. 87–114.
- [242] ———, *Problems related to unconstrained optimization*, in Numerical Methods for Unconstrained Optimization, W. Murray, ed., Academic Press, 1972, pp. 29–55.
- [243] ———, *On search directions for minimization algorithms*, Mathematical Programming, 4 (1973), pp. 193–201.
- [244] ———, *Convergence properties of a class of minimization algorithms*, in Nonlinear Programming 2, O. L. Mangasarian, R. R. Meyer, and S. M. Robinson, eds., Academic Press, New York, 1975, pp. 1–27.
- [245] ———, *Some convergence properties of the conjugate gradient method*, Mathematical Programming, 11 (1976), pp. 42–49.
- [246] ———, *Some global convergence properties of a variable metric algorithm for minimization without exact line searches*, in Nonlinear Programming, SIAM-AMS Proceedings, Vol. IX, R. W. Cottle and C. E. Lemke, eds., SIAM Publications, 1976, pp. 53–72.
- [247] ———, *A fast algorithm for nonlinearly constrained optimization calculations*, in Numerical Analysis Dundee 1977, G. A. Watson, ed., Springer Verlag, Berlin, 1977, pp. 144–157.
- [248] ———, *Restart procedures for the conjugate gradient method*, Mathematical Programming, 12 (1977), pp. 241–254.
- [249] ———, *Algorithms for nonlinear constraints that use Lagrangian functions*, Mathematical Programming, 14 (1978), pp. 224–248.
- [250] ———, *The convergence of variable metric methods for nonlinearly constrained optimization calculations*, in Nonlinear Programming 3, Academic Press, New York and London, 1978, pp. 27–63.
- [251] ———, *On the rate of convergence of variable metric algorithms for unconstrained optimization*, Technical Report DAMTP 1983/NA7, Department of Applied Mathematics and Theoretical Physics, Cambridge University, 1983.
- [252] ———, *Variable metric methods for constrained optimization*, in Mathematical Programming: The State of the Art, Bonn, 1982, Springer-Verlag, Berlin, 1983, pp. 288–311.
- [253] ———, *Nonconvex minimization calculations and the conjugate gradient method*, Lecture Notes in Mathematics, 1066 (1984), pp. 122–141.
- [254] ———, *The performance of two subroutines for constrained optimization on some difficult test problems*, in Numerical Optimization, P. T. Boggs, R. H. Byrd, and R. B. Schnabel, eds., SIAM Publications, Philadelphia, 1984.
- [255] ———, *Convergence properties of algorithms for nonlinear optimization*, SIAM Review, 28 (1986), pp. 487–500.
- [256] ———, *Direct search algorithms for optimization calculations*, Acta Numerica, 7 (1998), pp. 287–336.
- [257] ———, *UOBYQA: unconstrained optimization by quadratic approximation*, Mathematical Programming, Series B, 92 (2002), pp. 555–582.
- [258] ———, *On trust-region methods for unconstrained minimization without derivatives*, Mathematical Programming, 97 (2003), pp. 605–623.
- [259] ———, *Least Frobenius norm updating of quadratic models that satisfy interpolation conditions*, Mathematical Programming, 100 (2004), pp. 183–215.
- [260] ———, *The NEWUOA software for unconstrained optimization without derivatives*, Numerical Analysis Report DAMPT 2004/NA05, University of Cambridge, Cambridge, UK, 2004.

- [261] M. J. D. POWELL AND P. L. TOINT, *On the estimation of sparse Hessian matrices*, SIAM Journal on Numerical Analysis, 16 (1979), pp. 1060–1074.
- [262] R. L. RARDIN, *Optimization in Operations Research*, Prentice Hall, Englewood Cliffs, NJ, 1998.
- [263] F. RENDL AND H. WOLKOWICZ, *A semidefinite framework for trust region subproblems with applications to large scale minimization*, Mathematical Programming, 77 (1997), pp. 273–299.
- [264] J. M. RESTREPO, G. K. LEAF, AND A. GRIEWANK, *Circumventing storage limitations in variational data assimilation studies*, SIAM Journal on Scientific Computing, 19 (1998), pp. 1586–1605.
- [265] K. RITTER, *On the rate of superlinear convergence of a class of variable metric methods*, Numerische Mathematik, 35 (1980), pp. 293–313.
- [266] S. M. ROBINSON, *A quadratically convergent algorithm for general nonlinear programming problems*, Mathematical Programming, 3 (1972), pp. 145–156.
- [267] ———, *Perturbed Kuhn-Tucker points and rates of convergence for a class of nonlinear programming algorithms*, Mathematical Programming, 7 (1974), pp. 1–16.
- [268] ———, *Generalized equations and their solutions. Part II: Applications to nonlinear programming*, Mathematical Programming Study, 19 (1982), pp. 200–221.
- [269] R. T. ROCKAFELLAR, *The multiplier method of Hestenes and Powell applied to convex programming*, Journal of Optimization Theory and Applications, 12 (1973), pp. 555–562.
- [270] ———, *Lagrange multipliers and optimality*, SIAM Review, 35 (1993), pp. 183–238.
- [271] J. B. ROSEN AND J. KREUSER, *A gradient projection algorithm for nonlinear constraints*, in Numerical Methods for Non-Linear Optimization, F. A. Lootsma, ed., Academic Press, London and New York, 1972, pp. 297–300.
- [272] ———, *Iterative Methods for Sparse Linear Systems*, SIAM Publications, Philadelphia, PA, second ed., 2003.
- [273] Y. SAAD AND M. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [274] H. SCHEEL AND S. SCHOLTES, *Mathematical programs with complementarity constraints: Stationarity, optimality and sensitivity*, Mathematics of Operations Research, 25 (2000), pp. 1–22.
- [275] T. SCHLICK, *Modified Cholesky factorizations for sparse preconditioners*, SIAM Journal on Scientific Computing, 14 (1993), pp. 424–445.
- [276] R. B. SCHNABEL AND E. ESKOW, *A new modified Cholesky factorization*, SIAM Journal on Scientific Computing, 11 (1991), pp. 1136–1158.
- [277] R. B. SCHNABEL AND P. D. FRANK, *Tensor methods for nonlinear equations*, SIAM Journal on Numerical Analysis, 21 (1984), pp. 815–843.
- [278] G. SCHULLER, *On the order of convergence of certain quasi-Newton methods*, Numerische Mathematik, 23 (1974), pp. 181–192.
- [279] G. A. SCHULTZ, R. B. SCHNABEL, AND R. H. BYRD, *A family of trust-region-based algorithms for unconstrained minimization with strong global convergence properties*, SIAM Journal on Numerical Analysis, 22 (1985), pp. 47–67.
- [280] G. A. F. SEBER AND C. J. WILD, *Nonlinear Regression*, John Wiley & Sons, New York, 1989.
- [281] T. STEihaug, *The conjugate gradient method and trust regions in large scale optimization*, SIAM Journal on Numerical Analysis, 20 (1983), pp. 626–637.
- [282] J. STOER, *On the relation between quadratic termination and convergence properties of minimization algorithms. Part I: Theory*, Numerische Mathematik, 28 (1977), pp. 343–366.

- [283] K. TANABE, *Centered Newton method for mathematical programming*, in System Modeling and Optimization: Proceedings of the 13th IFIP conference, vol. 113 of Lecture Notes in Control and Information Systems, Berlin, 1988, Springer-Verlag, pp. 197–206.
- [284] M. J. TODD, *Potential reduction methods in mathematical programming*, Mathematical Programming, Series B, 76 (1997), pp. 3–45.
- [285] ———, *Semidefinite optimization*, Acta Numerica, 10 (2001), pp. 515–560.
- [286] ———, *Detecting infeasibility in infeasible-interior-point methods for optimization*, in Foundations of Computational Mathematics, Minneapolis, 2002, F. Cucker, R. DeVore, P. Olver, and E. Suli, eds., Cambridge University Press, Cambridge, 2004, pp. 157–192.
- [287] M. J. TODD AND Y. YE, *A centered projective algorithm for linear programming*, Mathematics of Operations Research, 15 (1990), pp. 508–529.
- [288] P. L. TOINT, *On sparse and symmetric matrix updating subject to a linear equation*, Mathematics of Computation, 31 (1977), pp. 954–961.
- [289] ———, *Towards an efficient sparsity exploiting Newton method for minimization*, in Sparse Matrices and Their Uses, Academic Press, New York, 1981, pp. 57–87.
- [290] L. TREFETHEN AND D. BAU, *Numerical Linear Algebra*, SIAM, Philadelphia, PA, 1997.
- [291] M. ULBRICH, S. ULBRICH, AND L. N. VICENTE, *A globally convergence primal-dual interior-point filter method for nonlinear programming*, Mathematical Programming, Series B, 100 (2004), pp. 379–410.
- [292] L. VANDENBERGHE AND S. BOYD, *Semidefinite programming*, SIAM Review, 38 (1996), pp. 49–95.
- [293] R. J. VANDERBEI, *Linear Programming: Foundations and Extensions*, Springer Verlag, New York, second ed., 2001.
- [294] R. J. VANDERBEI AND D. F. SHANNO, *An interior point algorithm for nonconvex nonlinear programming*, Computational Optimization and Applications, 13 (1999), pp. 231–252.
- [295] A. VARDI, *A trust region algorithm for equality constrained minimization: convergence properties and implementation*, SIAM Journal of Numerical Analysis, 22 (1985), pp. 575–591.
- [296] S. A. VAVASIS, *Quadratic programming is NP*, Information Processing Letters, 36 (1990), pp. 73–77.
- [297] ———, *Nonlinear Optimization*, Oxford University Press, New York and Oxford, 1991.
- [298] A. WÄCHTER, *An interior point algorithm for large-scale nonlinear optimization with applications in process engineering*, PhD thesis, Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA, 2002.
- [299] A. WÄCHTER AND L. T. BIEGLER, *Failure of global convergence for a class of interior point methods for nonlinear programming*, Mathematical Programming, 88 (2000), pp. 565–574.
- [300] ———, *Line search filter methods for nonlinear programming: Motivation and global convergence*, SIAM Journal on Optimization, 16 (2005), pp. 1–31.
- [301] ———, *On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming*, Mathematical Programming, 106 (2006), pp. 25–57.
- [302] H. WALKER, *Implementation of the GMRES method using Householder transformations*, SIAM Journal on Scientific and Statistical Computing, 9 (1989), pp. 815–825.
- [303] R. A. WALTZ, J. L. MORALES, J. NOCEDAL, AND D. ORBAN, *An interior algorithm for nonlinear optimization that combines line search and trust region steps*, Tech. Rep. 2003-6, Optimization Technology Center, Northwestern University, Evanston, IL, USA, June 2003.
- [304] WATERLOO MAPLE SOFTWARE, INC, *Maple V software package*, 1994.

- [305] L. T. WATSON, *Numerical linear algebra aspects of globally convergent homotopy methods*, SIAM Review, 28 (1986), pp. 529–545.
- [306] R. B. WILSON, *A simplicial algorithm for concave programming*, PhD thesis, Graduate School of Business Administration, Harvard University, 1963.
- [307] D. WINFIELD, *Function and functional optimization by interpolation in data tables*, PhD thesis, Harvard University, Cambridge, USA, 1969.
- [308] W. L. WINSTON, *Operations Research*, Wadsworth Publishing Co., third ed., 1997.
- [309] P. WOLFE, *A duality theorem for nonlinear programming*, Quarterly of Applied Mathematics, 19 (1961), pp. 239–244.
- [310] ———, *The composite simplex algorithm*, SIAM Review, 7 (1965), pp. 42–54.
- [311] S. WOLFRAM, *The Mathematica Book*, Cambridge University Press and Wolfram Media, Inc., third ed., 1996.
- [312] L. A. WOLSEY, *Integer Programming*, Wiley–Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, New York, NY, 1998.
- [313] M. H. WRIGHT, *Interior methods for constrained optimization*, in Acta Numerica 1992, Cambridge University Press, Cambridge, 1992, pp. 341–407.
- [314] ———, *Direct search methods: Once scorned, now respectable*, in Numerical Analysis 1995 (Proceedings of the 1995 Dundee Biennial Conference in Numerical Analysis), Addison Wesley Longman, 1996, pp. 191–208.
- [315] S. J. WRIGHT, *Applying new optimization algorithms to model predictive control*, in Chemical Process Control-V, J. C. Kantor, ed., CACHE, 1997.
- [316] ———, *Primal-Dual Interior-Point Methods*, SIAM Publications, Philadelphia, PA, 1997.
- [317] ———, *Modified Cholesky factorizations in interior-point algorithms for linear programming*, SIAM Journal on Optimization, 9 (1999), pp. 1159–1191.
- [318] S. J. WRIGHT AND J. N. HOLT, *An inexact Levenberg-Marquardt method for large sparse nonlinear least squares problems*, Journal of the Australian Mathematical Society, Series B, 26 (1985), pp. 387–403.
- [319] E. A. YILDIRIM AND S. J. WRIGHT, *Warm-start strategies in interior-point methods for linear programming*, SIAM Journal on Optimization, 12 (2002), pp. 782–810.
- [320] Y. YUAN, *On the truncated conjugate-gradient method*, Mathematical Programming, Series A, 87 (2000), pp. 561–573.
- [321] Y. ZHANG, *Solving large-scale linear programs with interior-point methods under the Matlab environment*, Optimization Methods and Software, 10 (1998), pp. 1–31.
- [322] C. ZHU, R. H. BYRD, P. LU, AND J. NOCEDAL, *Algorithm 778: L-BFGS-B, FORTRAN subroutines for large scale bound constrained optimization*, ACM Transactions on Mathematical Software, 23 (1997), pp. 550–560.

# Index

- Accumulation point, *see* Limit point
- Active set, 308, 323, 336, 342
- Affine scaling
  - direction, 395, 398, 414
  - method, 417
- Alternating variables method, *see also*
  - Coordinate search method, 104, 230
- Angle test, 41
- Applications
  - design optimization, 1
  - finance, 7
  - portfolio optimization, 1, 449–450, 492
  - transportation, 4
- Armijo line search, *see* Line search, Armijo
- Augmented Lagrangian function, 423
  - as merit function, 436
  - definition, 514
  - exactness of, 517–518
  - example, 516
- Augmented Lagrangian method, 422, 498, 514–526
  - convergence, 518–519
  - framework for, 515
  - implementation, 519–523
  - LANCELOT, 175, 519–522
  - motivation, 514–515
- Automatic differentiation, 170, 194
  - adjoint variables, 208, 209
  - and graph-coloring algorithms, 212, 216–218
  - checkpointing, 210
  - common expressions, 211
  - computational graph, 205–206, 208, 210, 211, 213, 215

- Automatic (*cont.*)
  - computational requirements, 206–207, 210, 214, 216, 219
  - forward mode, 206–207, 278
  - forward sweep, 206, 208, 210, 213–215, 219
  - foundations in elementary arithmetic, 194, 204
  - Hessian calculation
    - forward mode, 213–215
    - interpolation formulae, 214–215
    - reverse mode, 215–216
  - intermediate variables, 205–209, 211, 212, 218
  - Jacobian calculation, 210–213
    - forward mode, 212
    - reverse mode, 212–213
  - limitations of, 216–217
  - reverse mode, 207–210
  - reverse sweep, 208–210, 218
  - seed vectors, 206, 207, 212, 213, 216
  - software, 194, 210, 217
- Backtracking, 37, 240
- Barrier functions, 566, 583
- Barrier method, 563–566
  - primal, 583
- Basic variables, 429
- Basis matrix, 429–431
- BFGS method, 24, 29, 136–143
  - damping, 537
  - implementation, 142–143
  - properties, 141–142, 161
  - self-correction, 142
  - skipping, 143, 537
- Bound-constrained optimization, 97, 485–490
- BQPD, 490
- Broyden class, *see* Quasi-Newton method, Broyden class
- Broyden's method, 273, 274, 284, 285, 302, 634
  - derivation of, 279–281
  - limited-memory variants, 283
  - rate of convergence, 281–283
  - statement of algorithm, 281
- Byrd–Omojokun method, 547, 579
- Calculus of variations, 9
- Cancellation error, *see* Floating-point arithmetic, cancellation
- Cauchy point, 71–73, 76, 77, 93, 100, 170, 172, 262, 486
  - calculation of, 71–72, 96
  - for nonlinear equations, 291–292
  - role in global convergence, 77–79
- Cauchy sequence, 618
- Cauchy–Schwarz inequality, 75, 99, 151, 600
- Central path, 397–399, 417
  - for nonlinear problems, 565, 584, 594
  - neighborhoods of, 399–401, 403, 406, 413
- Chain rule, 29, 194, 204, 206–208, 213, 625, 627, 629
- Cholesky factorization, 87, 141, 143, 161, 251, 259, 289, 292, 454, 599, 608–609, 617
  - incomplete, 174
  - modified, 48, 51–54, 63, 64, 76
    - bounded modified factorization property, 48
    - sparse, 412–413
    - stability of, 53, 617
- Classification of algorithms, 422
- Combinatorial difficulty, 424
- Complementarity condition, 70, 313, 321, 333, 397
  - strict, 321, 337, 342, 533, 565, 591
- Complementarity problems
  - linear (LCP), 415
  - nonlinear (NCP), 417
- Complexity of algorithms, 388–389, 393, 406, 415, 417
- Conditioning, *see also* Matrix, condition number, 426, 430–432, 616–617
  - ill conditioned, 29, 502, 514, 586, 616
  - well conditioned, 616
- Cone, 621
- Cone of feasible directions, *see* Tangent cone
- Conjugacy, 25, 102
- Conjugate direction method, 103

- expanding subspace minimization, 106, 172, 173
- termination of, 103
- Conjugate gradient method, 71, 101–132, 166, 170–173, 253, 278
  - $n$ -step quadratic convergence, 133
  - clustering of eigenvalues, 116
  - effect of condition number, 117
  - expanding subspace minimization, 112
  - Fletcher–Reeves, *see* Fletcher–Reeves method
  - for reduced system, 459–461
  - global convergence, 40
  - Hestenes–Stiefel, 123
  - Krylov subspace, 113
  - modified for indefiniteness, 169–170
  - nonlinear, 25, 121–131
  - numerical performance, 131
  - optimal polynomial, 113
  - optimal process, 112
  - Polak–Ribière, *see* Polak–Ribière method
  - practical version, 111
  - preconditioned, 118–119, 170, 460
  - projected, 461–463, 548, 571, 581, 593
  - rate of convergence, 112
  - relation to limited-memory, 180
  - restarts, 124
  - superlinear convergence, 132
  - superquadratic, 133
  - termination, 115, 124
- Constrained optimization, 6
  - nonlinear, 4, 6, 211, 293, 356, 421, 498, 500
- Constraint qualifications, 315–320, 333, 338–340, 350
  - linear independence (LICQ), 320, 321, 323, 339, 341, 358, 464, 503, 517, 533, 557, 565, 591
  - Mangasarian–Fromovitz (MFCQ), 339–340
- Constraints, 2, 307
  - bounds, 434, 519, 520
  - equality, 305
  - inequality, 305
- Continuation methods for nonlinear equations, 274, 303
  - application to KKT conditions for nonlinear optimization, 565
  - convergence of, 300–301
  - formulation as initial-value ODE, 297–299
  - motivation, 296–297
  - predictor–corrector method, 299–300
  - zero path, 296–301, 303
    - divergence of, 300–301
    - tangent, 297–300
    - turning point, 296, 297, 300
- Convergence, rate of, 619–620
  - $n$ -step quadratic, 133
  - linear, 262, 619, 620
  - quadratic, 23, 29, 49, 168, 257, 619, 620
  - sublinear, 29
  - superlinear, 23, 29, 73, 132, 140, 142, 160, 161, 168, 262–265, 414, 619, 620
  - superquadratic, 133
- Convex combination, 621
- Convex hull, 621
- Convex programming, 7, 8, 335
- Convexity, 7–8
  - of functions, 8, 16–17, 28, 250
  - of sets, 8, 28, 352
  - strict, 8
- Coordinate descent method, *see* Alternating variables method, 233
- Coordinate relaxation step, 431
- Coordinate search method, 135, 230–231
- CPLEX, 490
- Critical cone, 330
  
- Data-fitting problems, 11–12, 248**
- Degeneracy, 465
  - of basis, 366, 369, 372, 382
  - of linear program, 366
- Dennis and Moré characterization, 47
- Descent direction, 21, 29, 30
- DFP method, 139
- Differential equations
  - ordinary, 299
  - partial, 216, 302
- Direct sum, 603
- Directional derivative, 206, 207, 437, 628–629
- Discrete optimization, 5–6

- Dual slack variables, 359
- Dual variables, *see also* Lagrange multipliers, 359
- Duality, 350
  - in linear programming, 359–362
  - in nonlinear programming, 343–349
  - weak, 345, 361
- Eigenvalues, 84, 252, 337, 599, 603, 613
  - negative, 77, 92
  - of symmetric matrix, 604
- Eigenvectors, 84, 252, 603
- Element function, 186
- Elimination of variables, 424
  - linear equality constraints, 428–433
  - nonlinear, 426–428
  - when inequality constraints are present, 434
- Ellipsoid algorithm, 389, 393, 417
- Error
  - absolute, 614
  - relative, 196, 251, 252, 614, 617
  - truncation, 216
- Errors-in-variables models, 265
- Feasibility restoration, 439–440
- Feasible sequences, 316–325, 332–333, 336
  - limiting directions of, 316–325, 329, 333
- Feasible set, 3, 305, 306, 338
  - geometric properties of, 340–341
  - primal, 358
  - primal-dual, 397, 399, 405, 414
- Filter method, 437–440
- Filters, 424, 437–440, 575, 589
  - for interior-point methods, 575
- Finite differencing, 170, 193–204, 216, 268, 278
  - and graph-coloring algorithms, 202–204
  - and noise, 221
  - central-difference formula, 194, 196–197, 202, 217
  - forward-difference formula, 195, 196, 202, 217
  - gradient approximation, 195–197
  - graph-coloring algorithms and, 200–201
  - Hessian approximation, 201–204
  - Jacobian approximation, 197–201, 283
- First-order feasible descent direction, 310–315
- First-order optimality conditions, *see also* Karush–Kuhn–Tucker (KKT) conditions, 90, 275, 307–329, 340, 352
  - derivation of, 315–329
  - examples, 308–315, 317–319, 321–322
  - fundamental principle of, 325–326
  - unconstrained optimization, 14–15, 513
- Fixed-regressor model, 248
- Fletcher–Reeves method, 102, 121–131
  - convergence of, 125
  - numerical performance, 131
- Floating-point arithmetic, 216, 614–615, 617
  - cancellation, 431, 615
  - double-precision, 614
  - roundoff error, 195, 217, 251, 615
  - unit roundoff, 196, 217, 614
- Floating-point numbers, 614
  - exponent, 614
  - fractional part, 614
- Forcing sequence, *see* Newton’s method, inexact, forcing sequence
- Function
  - continuous, 623–624
  - continuously differentiable, 626, 631
  - derivatives of, 625–630
  - differentiable, 626
  - Lipschitz continuous, 624, 630
  - locally Lipschitz continuous, 624
  - one-sided limit, 624
  - univariate, 625
- Functions
  - smooth, 10, 14, 306–307, 330
- Fundamental theorem of algebra, 603
- Gauss–Newton method, 254–258, 263, 266, 275
  - connection to linear least squares, 255
  - line search in, 254
  - performance on large-residual problems, 262
- Gaussian elimination, 51, 430, 455, 609
  - sparse, 430, 433
  - stability of, 617
  - with row partial pivoting, 607, 617

- Global convergence, 77–92, 261, 274
- Global minimizer, 12–13, 16, 17, 502, 503
- Global optimization, 6–8, 422
- Global solution, *see also* Global minimizer, 6, 69–70, 89–91, 305, 335, 352
- GMRES, 278, 459, 492, 571
- Goldstein condition, 36, 48
- Gradient, 625
  - generalized, 18
- Gradient projection method, 464, 485–490, 492, 521
- Group partial separability, *see* Partially separable function, group partially separable
- Hessian, 14, 19, 20, 23, 26, 626
  - average, 138, 140
- Homotopy map, 296
- Homotopy methods, *see* Continuation methods for nonlinear equations
- Implicit filtering, 240–242
- Implicit function theorem, 324, 630–631
- Inexact Newton method, *see* Newton’s method, inexact
- Infeasibility measure, 437
- Inner product, 599
- Integer programming, 5, 416
  - branch-and-bound algorithm, 6
- Integral equations, 302
- Interior-point methods, *see* Primal-dual interior-point methods
  - nonlinear, *see* Nonlinear interior-point method
- Interlacing eigenvalue theorem, 613
- Interpolation conditions, 223
- Invariant subspace, *see* Partially separable optimization, invariant subspace
- Iterative refinement, 463
- Jacobian, 246, 254, 256, 269, 274, 324, 395, 504, 627, 630
- Karmarkar’s algorithm, 389, 393, 417
- Karush–Kuhn–Tucker (KKT) conditions, 330, 332, 333, 335–337, 339, 350, 354, 503, 517, 520, 528
  - for general constrained problem, 321
  - for linear programming, 358–360, 367, 368, 394–415
  - for linear programming, 394
- KNITRO, 490, 525, 583, 592
- Krylov subspace, 108
  - method, 459
- L-BFGS algorithm, 177–180, 183
- Lagrange multipliers, 310, 330, 333, 337, 339, 341–343, 353, 358, 360, 419, 422
  - estimates of, 503, 514, 515, 518, 521, 522, 584
- Lagrangian function, 90, 310, 313, 320, 329, 330, 336
  - for linear program, 358, 360
  - Hessian of, 330, 332, 333, 335, 337, 358
- LANCELOT, 520, 525, 592
- Lanczos method, 77, 166, 175–176
- LAPACK, 607
- Least-squares multipliers, 581
- Least-squares problems, linear, 250–254
  - normal equations, 250–251, 255, 259, 412
  - sensitivity of solutions, 252
  - solution via QR factorization, 251–252
  - solution via SVD, 252–253
- Least-squares problems, nonlinear, 12, 210
  - applications of, 246–248
  - Dennis–Gay–Welsch algorithm, 263–265
  - Fletcher–Xu algorithm, 263
  - large-residual problems, 262–265
  - large-scale problems, 257
  - scaling of, 260–261
  - software for, 263, 268
  - statistical justification of, 249–250
  - structure, 247, 254
- Least-squares problems, total, 265
- Level set, 92, 261
- Levenberg–Marquardt method, 258–262, 266, 289
  - as trust-region method, 258–259, 292
  - for nonlinear equations, 292
  - implementation via orthogonal transformations, 259–260
  - inexact, 268

- Levenberg–Marquardt (*cont.*)
  - local convergence of, 262
  - performance on large-residual problems, 262
- lim inf, lim sup, 618–619
- Limit point, 28, 79, 92, 99, 502, 503, 618, 620
- Limited-memory method, 25, 176–185, 190
  - compact representation, 181–184
  - for interior-point method, 575, 597
  - L-BFGS, 176–180, 538
  - memoryless BFGS method, 180
  - performance of, 179
  - relation to CG, 180
  - scaling, 178
  - SRI, 183
  - two-loop recursion, 178
- Line search, *see also* Step length selection
  - Armijo, 33, 48, 240
  - backtracking, 37
  - curvature condition, 33
  - Goldstein, 36
  - inexact, 31
  - Newton’s method with, 22–23
  - quasi-Newton methods with, 23–25
  - search directions, 20–25
  - strong Wolfe conditions, *see* Wolfe conditions, strong
  - sufficient decrease, 33
  - Wolfe conditions, *see* Wolfe conditions
- Line search method, 19–20, 30–48, 66, 67, 71, 230–231, 247
  - for nonlinear equations, 271, 285, 287–290
    - global convergence of, 287–288
    - poor performance of, 288–289
- Linear programming, 4, 6, 7, 9, 293
  - artificial variables, 362, 378–380
  - basic feasible points, 362–366
  - basis  $\mathcal{B}$ , 362–368, 378
  - basis matrix, 363
  - dual problem, 359–362
  - feasible polytope, 356
    - vertices of, 365–366
  - fundamental theorem of, 363–364
  - infeasible, 356, 357
  - nonbasic matrix, 367
  - primal solution set, 356
  - slack and surplus variables, 356, 357, 362, 379, 380
  - splitting variables, 357
  - standard form, 356–357
  - unbounded, 356, 357, 369
  - warm start, 410, 416
- Linearly constrained Lagrangian methods, 522–523, 527
  - MINOS, 523, 527
- Linearly dependent, 337
- Linearly independent, 339, 503, 504, 517, 519, 602
- Lipschitz continuity, *see also* Function, Lipschitz continuous, 80, 93, 256, 257, 261, 269, 276–278, 287, 294
- Local minimizer, 12, 14, 273
  - isolated, 13, 28
  - strict, 13, 14, 16, 28, 517
  - weak, 12
- Local solution, *see also* Local minimizer, 6, 305–306, 316, 325, 329, 332, 340, 342, 352, 513
  - isolated, 306
  - strict, 306, 333, 335, 336
  - strong, 306
- Log-barrier function, 417, 597
  - definition, 583–584
  - difficulty of minimizing, 584–585
  - example, 586
  - ill conditioned Hessian of, 586
- Log-barrier method, 498, 584
- LOQQ, 490, 592
- LSQR method, 254, 268, 459, 492, 571
- LU factorization, 606–608
- Maratos effect, 440–446, 543, 550
  - example of, 440, 543
  - remedies, 442
- Matlab, 416
- Matrix
  - condition number, 251, 601–602, 604, 610, 616
  - determinant, 154, 605–606
  - diagonal, 252, 412, 429, 599
  - full-rank, 298, 300, 504, 609
  - identity, 599

- indefinite, 76
- inertia, 55, 454
- lower triangular, 599, 606, 607
- modification, 574
- nonsingular, 325, 337, 601, 612
- null space, 298, 324, 337, 430, 432, 603, 608, 609
- orthogonal, 251, 252, 337, 432, 599, 604, 609
- permutation, 251, 429, 606
- positive definite, 15, 16, 23, 28, 68, 76, 337, 599, 603, 609
- positive semidefinite, 8, 15, 70, 415, 599
- projection, 462
- range space, 430, 603
- rank-deficient, 253
- rank-one, 24
- rank-two, 24
- singular, 337
- sparse, 411, 413, 607
  - Cholesky factorization, 413
- symmetric, 24, 68, 412, 599, 603
- symmetric indefinite, 413
- symmetric positive definite, 608
- trace, 154, 605
- transpose, 599
- upper triangular, 251, 337, 599, 606, 607, 609
- Maximum likelihood estimate, 249
- Mean value theorem, 629–630
- Merit function, *see also* Penalty function, 435–437, 446
  - $\ell_1$ , 293, 435–436, 513, 540–543, 550
  - choice of parameter, 543
- exact, 435–436
  - definition of, 435
  - nonsmoothness of, 513
- Fletcher’s augmented Lagrangian, 436, 540
- for feasible methods, 435
- for nonlinear equations, 273, 285–287, 289, 290, 293, 296, 301–303, 505
- for SQP, 540–543
- Merit functions, 424, 575
- Method of multipliers, *see* Augmented Lagrangian method
- MINOS, *see also* Linearly constrained Lagrangian methods, 523, 525, 592
- Model-based methods for derivative-free optimization, 223–229
  - minimum Frobenius change, 228
- Modeling, 2, 9, 11, 247–249
- Monomial basis, 227
- MOSEK, 490
- Multiobjective optimization, 437
- Negative curvature direction, 49, 50, 63, 76, 169–172, 175, 489, 491
- Neighborhood, 13, 14, 28, 256, 621
- Network optimization, 358
- Newton’s method, 25, 247, 254, 257, 263
  - for log-barrier function, 585
  - for nonlinear equations, 271, 274–277, 281, 283, 285, 287–290, 294, 296, 299, 302
  - cycling, 285
  - inexact, 277–279, 288
  - for quadratic penalty function, 501, 506
  - global convergence, 40
  - Hessian-free, 165, 170
  - in one variable, 84–87, 91, 633
  - inexact, 165–168, 171, 213
    - forcing sequence, 166–169, 171, 277
  - large scale
    - LANCELOT, 175
    - line search method, 49
    - TRON, 175
  - modified, 48–49
    - adding a multiple of I, 51
    - eigenvalue modification, 49–51
- Newton–CG, 202
  - line search, 168–170
  - preconditioned, 174–175
  - trust-region, 170–175
- Newton–Lanczos, 175–176, 190
- rate of convergence, 44, 76, 92, 166–168, 275–277, 281–282, 620
- scale invariance, 27
- Noise in function evaluation, 221–222
- Nondifferentiable optimization, 511
- Nonlinear equations, 197, 210, 213, 633
  - degenerate solution, 274, 275, 283, 302
  - examples of, 271–272, 288–289, 300–301

- Nonlinear (*cont.*)
  - merit function, *see* Merit function, for nonlinear equations
  - multiple solutions, 273–274
  - nondegenerate solution, 274
  - quasi-Newton methods, *see* Broyden’s method
  - relationship to least squares, 271–272, 275, 292–293, 302
  - relationship to optimization, 271
  - relationship to primal-dual interior-point methods, 395
  - solution, 271
  - statement of problem, 270–271
- Nonlinear interior-point method, 423, 563–593
  - barrier formulation, 565
  - feasible version, 576
  - global convergence, 589
  - homotopy formulation, 565
  - superlinear convergence, 591
  - trust-region approach, 578
- Nonlinear least-squares, *see* Least-squares problems, nonlinear
- Nonlinear programming, *see* Constrained optimization, nonlinear
- Nonmonotone strategy, 18, 444–446
  - relaxed steps, 444
- Nonnegative orthant, 97
- Nonsmooth functions, 6, 17–18, 306, 307, 352
- Nonsmooth penalty function, *see* Penalty function, nonsmooth
- Norm
  - dual, 601
  - Euclidean, 25, 51, 251, 280, 302, 600, 601, 605, 610
  - Frobenius, 50, 138, 140, 601
  - matrix, 601–602
  - vector, 600–601
- Normal cone, 340–341
- Normal distribution, 249
- Normal subproblem, 580
- Null space, *see* Matrix, null space
- Numerical analysis, 355
  
- Objective function, 2, 10, 304
- One-dimensional minimization, 19, 56
  
- OOPS, 490
- OOQP, 490
- Optimality conditions, *see also* First-order optimality conditions, Second-order optimality conditions, 2, 9, 305
  - for unconstrained local minimizer, 14–17
- Order notation, 631–633
- Orthogonal distance regression, 265–267
  - contrast with least squares, 265–266
  - structure, 266–267
- Orthogonal transformations, 251, 259–260
  - Givens, 259, 609
  - Householder, 259, 609
  
- Partially separable function, 25, 186–189, 211
  - automatic detection, 211
  - definition, 211
- Partially separable optimization, 165
  - BFGS, 189
  - compactifying matrix, 188
  - element variables, 187
  - quasi-Newton method, 188
  - SR1, 189
- Penalty function, *see also* Merit function, 498
  - $\ell_1$ , 507–513
  - exact, 422–423, 507–513
  - nonsmooth, 497, 507–513
  - quadratic, *see also* Quadratic penalty method, 422, 498–507, 525–527, 586
    - difficulty of minimizing, 501–502
    - Hessian of, 505–506
    - relationship to augmented Lagrangian, 514
    - unbounded, 500
- Penalty parameter, 435, 436, 498, 500, 501, 507, 514, 521, 525
  - update, 511, 512
- PENNON, 526
- Pivoting, 251, 617
- Polak–Ribière method, 122
  - convergence of, 130
- Polak–Ribière method
  - numerical performance, 131

- Polynomial bases, 226
  - monomials, 227
- Portfolio optimization, *see* Applications, portfolio optimization
- Preconditioners, 118–120
  - banded, 120
  - constraint, 463
  - for constrained problems, 462
  - for primal-dual system, 571
  - for reduced system, 460
  - incomplete Cholesky, 120
  - SSOR, 120
- Preprocessing, *see* Presolving
- Presolving, 385–388
- Primal interior-point method, 570
- Primal-dual interior-point methods, 389, 597
  - centering parameter, 396, 398, 401, 413
  - complexity of, 393, 406, 415
  - contrasts with simplex method, 356, 393
  - convex quadratic programs, 415
  - corrector step, 414
  - duality measure, 395, 398
  - infeasibility detection, 411
  - linear algebra issues, 411–413
  - Mehrotra’s predictor-corrector algorithm, 393, 407–411
  - path-following algorithms, 399–414
    - long-step, 399–406
    - predictor-corrector (Mizuno–Todd–Ye) algorithm, 413
    - short-step, 413
  - potential function, 414
    - Tanabe–Todd–Ye, 414
  - potential-reduction algorithms, 414
  - predictor step, 413
  - quadratic programming, 480–485
  - relationship to Newton’s method, 394, 395
  - starting point, 410–411
- Primal-dual system, 567
- Probability density function, 249
- Projected conjugate gradient method, *see* Conjugate gradient method, projected
- Projected Hessian, 558
  - two-sided, 559
- Proximal point method, 523
- QMR method, 459, 492, 571
- QPA, 490
- QPOPT, 490
- QR factorization, 251, 259, 290, 292, 298, 337, 432, 433, 609–610
  - cost of, 609
  - relationship to Cholesky factorization, 610
- Quadratic penalty method, *see also* Penalty function, quadratic, 497, 501–502, 514
  - convergence of, 502–507
- Quadratic programming, 422, 448–492
  - active-set methods, 467–480
  - big M method, 473
  - blocking constraint, 469
  - convex, 449
  - cycling, 477
  - duality, 349, 490
  - indefinite, 449, 467, 491–492
  - inertia controlling methods, 491, 492
  - initial working set, 476
  - interior-point method, 480–485
  - nonconvex, *see* Quadratic programming, indefinite
  - null-space method, 457–459
  - optimal active set, 467
  - optimality conditions, 464
  - phase I, 473
  - Schur-complement method, 455–456
  - software, 490
  - strictly convex, 349, 449, 472, 477–478
  - termination, 477–478
  - updating factorizations, 478
  - working set, 468–478
- Quasi-Newton approximate Hessian, 23, 24, 73, 242, 634
- Quasi-Newton method, 25, 165, 247, 263, 501, 585
  - BFGS, *see* BFGS method, 263
  - bounded deterioration, 161
  - Broyden class, 149–152
  - curvature condition, 137

- Quasi-Newton (*cont.*)  
 DFP, *see* DFP method, 190, 264  
 for interior-point method, 575  
 for nonlinear equations, *see* Broyden's method  
 for partially separable functions, 25  
 global convergence, 40  
 large-scale, 165–189  
 limited memory, *see* Limited memory method  
 rate of convergence, 46, 620  
 secant equation, 24, 137, 139, 263–264, 280, 634  
 sparse, *see* Sparse quasi-Newton method
- Range space, *see* Matrix, range space
- Regularization, 574
- Residuals, 11, 245, 262–265, 269  
 preconditioned, 462  
 vector of, 18, 197, 246
- Restoration phase, 439
- Robust optimization, 7
- Root, *see* Nonlinear equations, solution
- Rootfinding algorithm, *see also* Newton's method, in one variable, 259, 260, 633  
 for trust-region subproblem, 84–87
- Rosenbrock function  
 extended, 191
- Roundoff error, *see* Floating-point arithmetic, roundoff error
- Row echelon form, 430
- $S\ell_1$ QP method, 293, 549
- Saddle point, 28, 92
- Scale invariance, 27, 138, 141  
 of Newton's method, *see* Newton's method, scale invariance
- Scaling, 26–27, 95–97, 342–343, 585  
 example of poor scaling, 26–27  
 matrix, 96
- Schur complement, 456, 611
- Secant method, *see also* Quasi-Newton method, 280, 633, 634
- Second-order correction, 442–444, 550
- Second-order optimality conditions, 330–337, 342, 602  
 for unconstrained optimization, 15–16  
 necessary, 92, 331  
 sufficient, 333–336, 517, 557
- Semidefinite programming, 415
- Sensitivity, 252, 616
- Sensitivity analysis, 2, 194, 341–343, 350, 361
- Separable function, 186
- Separating hyperplane, 327
- Sequential linear-quadratic programming (SLQP), 293, 423, 534
- Sequential quadratic programming, 423, 512, 523, 529–560  
 Byrd–Omojokun method, 547  
 derivation, 530–533  
 full quasi-Newton Hessian, 536  
 identification of optimal active set, 533  
 IQP vs. EQP, 533  
 KKT system, 275  
 least-squares multipliers, 539  
 line search algorithm, 545  
 local algorithm, 532  
 Newton–KKT system, 531  
 null-space, 538  
 QP multipliers, 538  
 rate of convergence, 557–560  
 reduced-Hessian approximation, 538–540  
 relaxation constraints, 547  
 $S\ell_1$ QP method, *see*  $S\ell_1$ QP method  
 step computation, 545  
 trust-region method, 546–549  
 warm start, 545
- Set  
 affine, 622  
 affine hull of, 622  
 bounded, 620  
 closed, 620  
 closure of, 621  
 compact, 621  
 interior of, 621  
 open, 620  
 relative interior of, 622, 623
- Sherman–Morrison–Woodbury formula, 139, 140, 144, 162, 283, 377, 612–613
- Simplex method  
 as active-set method, 388

- basis  $\mathcal{B}$ , 365
- complexity of, 388–389
- cycling, 381–382
  - lexicographic strategy, 382
  - perturbation strategy, 381–382
- degenerate steps, 372, 381
- description of single iteration, 366–372
- discovery of, 355
- dual simplex, 366, 382–385
- entering index, 368, 370, 372, 375–378
- finite termination of, 368–370
- initialization, 378–380
- leaving index, 368, 370
- linear algebra issues, 372–375
- Phase I/Phase II, 378–380
- pivoting, 368
- pricing, 368, 370, 375–376
  - multiple, 376
  - partial, 376
- reduced costs, 368
- revised, 366
- steepest-edge rule, 376–378
- Simulated annealing, 221
- Singular values, 255, 604
- Singular-value decomposition (SVD), 252, 269, 303, 603–604
- Slack variables, *see also* Linear programming, slack/surplus variables, 424, 519
- SNOPT, 536, 592
- Software
  - BQPD, 490
  - CPLEX, 490
  - for quadratic programming, 490
  - IPOPT, 183, 592
  - KNITRO, 183, 490, 525, 592
  - L-BFGS-B, 183
  - LANCELOT, 520, 525, 592
  - LOQQ, 490, 592
  - MINOS, 523, 525, 592
  - MOSEK, 490
  - OOPS, 490
  - OOQP, 490
  - PENNON, 526
  - QPA, 490
  - QPOPT, 490
  - SNOPT, 592
  - TRON, 175
  - VE09, 490
  - XPRESS-MP, 490
- Sparse quasi-Newton method, 185–186, 190
- SR1 method, 24, 144, 161
  - algorithm, 146
  - for constrained problems, 538, 540
  - limited-memory version, 177, 181, 183
  - properties, 147
  - safeguarding, 145
  - skipping, 145, 160
- Stability, 616–617
- Starting point, 18
- Stationary point, 15, 28, 289, 436, 505
- Steepest descent direction, 20, 21, 71, 74
- Steepest descent method, 21, 25–27, 31, 73, 95, 585
  - rate of convergence, 42, 44, 620
- Step length, 19, 30
  - unit, 23, 29
- Step length selection, *see also* Line search, 56–62
  - bracketing phase, 57
  - cubic interpolation, 59
  - for Wolfe conditions, 60
  - initial step length, 59
  - interpolation in, 57
  - selection phase, 57
- Stochastic optimization, 7
- Stochastic simulation, 221
- Strict complementarity, *see* Complementarity condition, strict
- Subgradient, 17
- Subspace, 602
  - basis, 430, 603
  - orthonormal, 432
  - dimension, 603
  - spanning set, 603
- Sufficient reduction, 71, 73, 79
- Sum of absolute values, 249
- Sum of squares, *see* Least-squares problems, nonlinear
- Symbolic differentiation, 194
- Symmetric indefinite factorization, 455, 570, 610–612
  - Bunch–Kaufman, 612

- Symmetric (*cont.*)
  - Bunch–Parlett, 611
  - modified, 54–56, 63
  - sparse, 612
- Symmetric rank-one update, *see* SR1 method
- Tangent, 315–325
- Tangent cone, 319, 340–341
- Taylor series, 15, 22, 28, 29, 67, 274, 309, 330, 332, 334, 502
- Taylor’s theorem, 15, 21–23, 80, 123, 138, 167, 193–195, 197, 198, 202, 274, 280, 294, 323, 325, 332, 334, 341, 630
  - statement of, 14
- Tensor methods, 274
  - derivation, 283–284
- Termination criterion, 92
- Triangular substitution, 433, 606, 609, 617
- Truncated Newton method, *see* Newton’s method, Newton-CG, line-search
- Trust region
  - boundary, 69, 75, 95, 171–173
  - box-shaped, 19, 293
  - choice of size for, 67, 81
  - elliptical, 19, 67, 95, 96, 100
  - radius, 20, 26, 68, 69, 73, 258, 294
  - spherical, 95, 258
- Trust-region method, 19–20, 69, 77, 79, 80, 82, 87, 91, 247, 258, 633
  - contrast with line search method, 20, 66–67
  - dogleg method, 71, 73–77, 79, 84, 91, 95, 99, 173, 291–293, 548
  - double-dogleg method, 99
  - for derivative-free optimization, 225
  - for nonlinear equations, 271, 273, 285, 290–296
    - global convergence of, 292–293
    - local convergence of, 293–296
  - global convergence, 71, 73, 76–92, 172
  - local convergence, 92–95
  - Newton variant, 26, 68, 92
  - software, 98
  - Steihaug’s approach, 77, 170–173, 489
  - strategy for adjusting radius, 69
  - subproblem, 19, 25–26, 68, 69, 72, 73, 76, 77, 91, 95–97, 258
    - approximate solution of, 68, 71
    - exact solution of, 71, 77, 79, 83–92
    - hard case, 87–88
    - nearly exact solution of, 95, 292–293
  - two-dimensional subspace
    - minimization, 71, 76–77, 79, 84, 95, 98, 100
- Unconstrained optimization, 6, 352, 427, 432, 499, 501
  - of barrier function, 584
- Unit ball, 91
- Unit roundoff, *see* Floating-point arithmetic, unit roundoff
- Variable metric method, *see* Quasi-Newton method
- Variable storage method, *see* Limited memory method
- VE09, 490
- Watchdog technique, 444–446
- Weakly active constraints, 342
- Wolfe conditions, 33–36, 48, 78, 131, 137, 138, 140–143, 146, 160, 179, 255, 287, 290
  - scale invariance of, 36
  - strong, 34, 35, 122, 125, 126, 128, 131, 138, 142, 162, 179
- XPRESS-MP, 490
- Zoutendijk condition, 38–41, 128, 156, 287