1.6 CONJUGATE DIRECTION METHODS

Conjugate direction methods are motivated by a desire to accelerate the convergence rate of steepest descent, while avoiding the overhead associated with Newton's method. They were originally developed for solving the quadratic problem

$$\begin{array}{ll} \text{minimize} & f(x) = \frac{1}{2}x'Qx - b'x \\ \text{subject to} & x \in \Re^n, \end{array} \tag{6.1}$$

where Q is positive definite, or equivalently, for solving the linear system

$$Qx = b. \tag{6.2}$$

They can also be used for solution of the more general system $Ax = b$, where A is invertible but not positive definite, after conversion to the positive definite system $A'Ax = A'b$.

Conjugate direction methods can solve these problems after at most n iterations but they are best viewed as iterative methods, since usually fewer than n iterations are required to attain a sufficiently accurate solution, particularly when n is large. They can also be used to solve nonquadratic optimization problems. For such problems, they do not in general terminate after a finite number of iterations, but still, when properly implemented, they have attractive convergence and rate of convergence properties. We will first develop the methods for quadratic problems and then discuss their application to more general problems.

Given a positive definite $n \times n$ matrix Q, we say that a set of nonzero vectors $d^1, \ldots, d^k$ are *Q-conjugate*, if

$$d^{i\prime}Qd^j = 0, \qquad \text{for all } i \text{ and } j \text{ such that } i \neq j. \tag{6.3}$$

If $d^1, \ldots, d^k$ are Q-conjugate, then they are *linearly independent*, since if one of these vectors, say d^k, were expressed as a linear combination of the others,

$$d^k = \alpha^1 d^1 + \cdots + \alpha^{k-1} d^{k-1},$$

then by multiplication with $d^{k\prime}Q$ we would obtain using the Q-conjugacy of d^k and d^j, $j = 1, \ldots, k-1$,

$$d^{k\prime}Qd^k = \alpha^1 d^{k\prime}Qd^1 + \cdots + \alpha^{k-1} d^{k\prime}Qd^{k-1} = 0,$$

which is impossible since $d^k \neq 0$ and Q is positive definite.

For a given set of n Q-conjugate directions $d^0, \ldots, d^{n-1}$, the corresponding *conjugate direction method* for unconstrained minimization of the quadratic function

$$f(x) = \frac{1}{2}x'Qx - b'x, \tag{6.4}$$

is given by

$$x^{k+1} = x^k + \alpha^k d^k, \qquad k = 0, \ldots, n-1, \tag{6.5}$$

where x^0 is an arbitrary starting vector and α^k is obtained by the line minimization rule

$$f(x^k + \alpha^k d^k) = \min_{\alpha} f(x^k + \alpha d^k). \tag{6.6}$$

The principal result about conjugate direction methods is that *successive iterates minimize f over a progressively expanding linear manifold that eventually includes the global minimum of f*. In particular, for each k, x^{k+1} minimizes f over the linear manifold passing through x^0 and spanned by the conjugate directions $d^0, \ldots, d^k$, that is,

$$x^{k+1} = \arg\min_{x \in M^k} f(x), \tag{6.7}$$

where

$$\begin{aligned} M^k &= \{x \mid x = x^0 + v,\ v \in (\text{subspace spanned by } d^0, \ldots, d^k)\} \\ &= x^0 + (\text{subspace spanned by } d^0, \ldots, d^k). \end{aligned} \tag{6.8}$$

In particular, x^n minimizes f over $\Re^n$.

To show this, note that by Eq. (6.6), we have for all i

$$\left.\frac{\partial f(x^i + \alpha d^i)}{\partial \alpha}\right|_{\alpha = \alpha^i} = \nabla f(x^{i+1})' d^i = 0$$

and, for $i = 0, \ldots, k-1$,

$$\begin{aligned} \nabla f(x^{k+1})' d^i &= (Qx^{k+1} - b)' d^i \\ &= \left(x^{i+1} + \sum_{j=i+1}^{k} \alpha^j d^j \right)' Q d^i - b' d^i \\ &= x^{i+1\prime} Q d^i - b' d^i \\ &= \nabla f(x^{i+1})' d^i, \end{aligned}$$

where we have used the conjugacy of d^i and d^j, $j = i+1, \ldots, k$. Combining the last two equations we obtain

$$\nabla f(x^{k+1})' d^i = 0, \qquad i = 0, \ldots, k, \tag{6.9}$$

so that

$$\left.\frac{\partial f(x^0 + \gamma^0 d^0 + \cdots + \gamma^k d^k)}{\partial \gamma^i}\right|_{\substack{\gamma^j = \alpha^j \\ j=1,\ldots,k}} = 0, \qquad i = 0, \ldots, k,$$

which verifies Eq. (6.7).

It is easy to visualize the expanding manifold minimization property of Eq. (6.7) when $b = 0$ and $Q = I$ (the identity matrix). In this case, the equal cost surfaces of f are concentric spheres, and the notion of Q-conjugacy reduces to usual orthogonality. By a simple algebraic argument, we see that minimization along n orthogonal directions yields the global minimum of f, that is, the center of the spheres. (This becomes evident once we rotate the coordinate system so that the given n orthogonal directions coincide with the coordinate directions.) The case of a general positive definite Q can be reduced to the case where $Q = I$ by means of a scaling transformation. By setting $y = Q^{1/2}x$, minimizing $\frac{1}{2}x'Qx$ is equivalent to minimizing $\frac{1}{2}\|y\|^2$. If $w^0, \ldots, w^{n-1}$ are any set of orthogonal nonzero vectors in $\Re^n$, the algorithm

$$y^{k+1} = y^k + \alpha^k w^k, \qquad k = 0, \ldots, n-1, \tag{6.10}$$

where

$$\alpha^k = \arg\min_{\alpha} \tfrac{1}{2}\|y^k + \alpha w^k\|^2,$$

terminates in at most n steps with $y^n = 0$. To pass back to the x-coordinate system, we multiply Eq. (6.10) by $Q^{-1/2}$ and obtain

$$x^{k+1} = x^k + \alpha^k d^k, \qquad k = 0, \ldots, n-1,$$

where $d^k = Q^{-1/2}w^k$. The orthogonality property of $w^0, \ldots, w^{n-1}$, that is, $w^{i\prime}w^j = 0$ for $i \neq j$, is equivalent to the requirement that the directions $d^0, \ldots, d^{n-1}$ be Q-conjugate, that is, $d^{i\prime}Qd^j = 0$ for $i \neq j$.

Thus, using the transformation $y = Q^{1/2}x$, we can think of any conjugate direction method for minimizing $\frac{1}{2}x'Qx$ as a method that minimizes $\frac{1}{2}\|y\|^2$ by successive minimization along n orthogonal directions (see Fig. 1.6.1).

Generating Q-Conjugate Directions

Given any set of linearly independent vectors $\xi^0, \ldots, \xi^k$, we can construct a set of mutually Q-conjugate directions $d^0, \ldots, d^k$ such that for all $i = 0, \ldots, k$, we have

$$(\text{subspace spanned by } d^0, \ldots, d^i) = (\text{subspace spanned by } \xi^0, \ldots, \xi^i), \tag{6.11}$$

using the so called *Gram-Schmidt procedure*. Indeed, let us do this recursively, starting with

$$d^0 = \xi^0. \tag{6.12}$$

Suppose that, for some $i < k$, we have selected Q-conjugate $d^0, \ldots, d^i$ so that the above property holds. We then take d^{i+1} to be of the form

$$d^{i+1} = \xi^{i+1} + \sum_{m=0}^{i} c^{(i+1)m} d^m \tag{6.13}$$

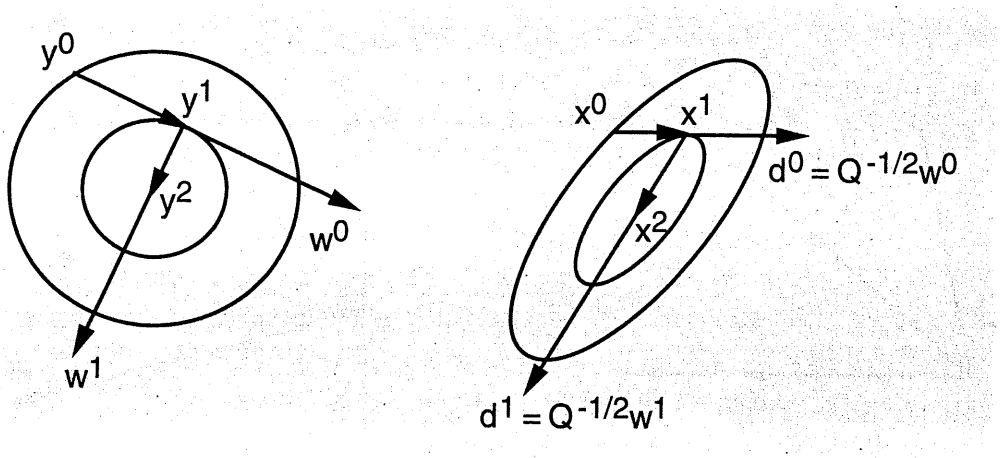

Figure 1.6.1. Geometric interpretation of conjugate direction methods in terms of successive minimization along n orthogonal directions. In (a) the function $\|y\|^2$ is minimized successively along the directions $w^0, \ldots, w^{n-1}$, which are orthogonal in the usual sense ($w^{i\prime}w^j = 0$ for $i \neq j$). When this process is viewed in the coordinate system of variables $x = Q^{-1/2}y$, it yields the conjugate direction method that uses the Q-conjugate directions $d^0, \ldots, d^{n-1}$ with $d^i = Q^{-1/2}w^i$, as shown in (b).

and choose the coefficients $c^{(i+1)m}$ so that d^{i+1} is Q-conjugate to $d^0, \ldots, d^i$. This will be so if for each $j = 0, \ldots, i$,

$$d^{i+1\prime}Qd^j = \xi^{i+1\prime}Qd^j + \left(\sum_{m=0}^{i} c^{(i+1)m}d^m\right)' Qd^j = 0. \tag{6.14}$$

Since $d^0, \ldots, d^i$ are Q-conjugate, we have $d^{m\prime}Qd^j = 0$ if $m \neq j$, and Eq. (6.14) yields

$$c^{(i+1)j} = -\frac{\xi^{i+1\prime}Qd^j}{d^{j\prime}Qd^j}, \qquad j = 0, \ldots, i. \tag{6.15}$$

Note that the denominator $d^{j\prime}Qd^j$ in the above equation is nonzero, since $d^0, \ldots, d^i$ are assumed Q-conjugate and are therefore nonzero. Note also that $d^{i+1} \neq 0$, since otherwise from Eqs. (6.11) and (6.13), ξ^{i+1} would be a linear combination of $\xi^0, \ldots, \xi^i$, contradicting the linear independence of $\xi^0, \ldots, \xi^k$. Finally, note from Eq. (6.13) that ξ^{i+1} lies in the subspace spanned by $d^0, \ldots, d^{i+1}$, while d^{i+1} lies in the subspace spanned by $\xi^0, \ldots, \xi^{i+1}$, since $d^0, \ldots, d^i$ and $\xi^0, \ldots, \xi^i$ span the same space [cf. Eq. (6.11)]. Thus, Eq. (6.11) is satisfied when i is increased to $i+1$ and the Gram-Schmidt procedure defined by Eqs. (6.12), (6.13), and (6.15), has the property claimed. Figure 1.6.2 illustrates the procedure.

It is also worth noting what will happen if the vectors $\xi^0, \ldots, \xi^i$ are linearly independent, but the next vector ξ^{i+1} is linearly dependent on these vectors. In this case it can be seen (compare also with Fig. 1.6.2) that the equations (6.13) and (6.15) are still valid, but the new vector d^{i+1}

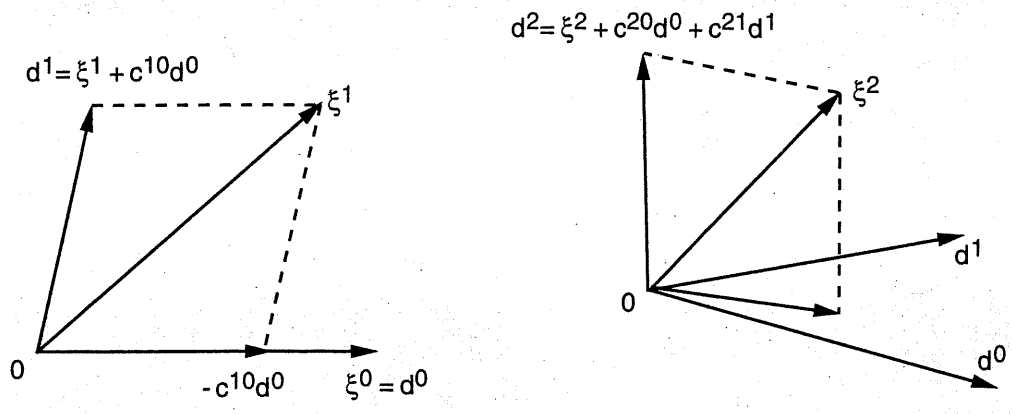

Figure 1.6.2. Illustration of the Gram-Schmidt procedure for generating Q-conjugate directions $d^0, \ldots, d^k$ from a set of linearly independent vectors $\xi^0, \ldots, \xi^k$, so that

$$(\text{subspace spanned by } d^0, \ldots, d^k) = (\text{subspace spanned by } \xi^0, \ldots, \xi^k).$$

Given $d^0, \ldots, d^{i-1}$, the ith direction is obtained as $d^i = \xi^i - \hat{\xi}^i$, where $\hat{\xi}^i$ is a vector on the subspace spanned by $d^0, \ldots, d^{i-1}$ (or $\xi^0, \ldots, \xi^{i-1}$) chosen so that d^i is Q-conjugate to $d^0, \ldots, d^{i-1}$. (It can be shown that the vector $\hat{\xi}^i$ is the projection of ξ^i on this subspace with respect to the norm $\|x\|_Q = \sqrt{x'Qx}$, that is, minimizes $\|\xi^i - x\|_Q$ over all x in this subspace; see Exercise 5.1.)

as given by Eq. (6.13) will be zero.† We can use this property to construct a set of Q-conjugate directions that span the same space as a set of vectors $\xi^0, \ldots, \xi^k$ that are not a priori known to be linearly independent. This construction can be accomplished with an extended version of the Gram-Schmidt procedure that generates directions via Eqs. (6.13) and (6.15), but each time the new direction d^{i+1} as given by Eq. (6.13) turns out to be zero, it is simply discarded rather than added to the set of preceding directions.

The Conjugate Gradient Method

The most important conjugate direction method, the conjugate gradient method, is obtained by applying the Gram-Schmidt procedure to the gradient vectors $\xi^0 = -g^0, \ldots, \xi^{n-1} = -g^{n-1}$, where we use the notation

$$g^k = \nabla f(x^k) = Qx^k - b. \tag{6.16}$$

Thus the conjugate gradient method is defined by

$$x^{k+1} = x^k + \alpha^k d^k, \tag{6.17}$$

† From Eq. (6.13) and the linear independence of $\xi^0, \ldots, \xi^i$, it is seen that d^{i+1} can be uniquely expressed as $d^{i+1} = \sum_{m=0}^{i} \gamma^m d^m$, where γ^m are some scalars. By multiplying this equation with $d^{j\prime}Q$ and by using the Q-conjugacy of the directions $d^0, \ldots, d^i$ and Eq. (6.14), we see that $\gamma^m = 0$ for all $m = 0, \ldots, i$.

where the stepsize α^k is obtained by line minimization, and the direction d^k is obtained by applying the kth step of the Gram-Schmidt procedure to the vector $-g^k$ and the preceding directions $d^0, \ldots, d^{k-1}$. In particular, from the Gram-Schmidt equations (6.13) and (6.15), we have

$$d^k = -g^k + \sum_{j=0}^{k-1} \frac{g^{k\prime} Q d^j}{d^{j\prime} Q d^j} d^j. \tag{6.18}$$

Note here that

$$d^0 = -g^0$$

and that the method terminates with an optimal solution if $g^k = 0$. The method also effectively stops if $d^k = 0$, but we will show that this can only happen if $g^k = 0$.

The key property of the conjugate gradient method is that the direction formula (6.18) can be greatly simplified. In particular, all but one of the coefficients in the sum of Eq. (6.18) turn out to be zero because, in view of the expanding manifold minimization property, the gradient g^k is orthogonal to the subspace spanned by $d^0, \ldots, d^{k-1}$ [cf. Eq. (6.9)]. We have the following proposition.

Proposition 1.6.1: The directions of the conjugate gradient method are generated by

$$d^0 = -g^0,$$

$$d^k = -g^k + \beta^k d^{k-1}, \qquad k = 1, \ldots, n-1,$$

where β^k is given by

$$\beta^k = \frac{g^{k\prime} g^k}{g^{k-1\prime} g^{k-1}}.$$

Furthermore, the method terminates with an optimal solution after at most n steps.

Proof: We first use induction to show that all the gradients g^k generated up to termination are linearly independent. The result is clearly true for $k = 1$. Suppose that the method has not terminated after k steps, and that $g^0, \ldots, g^{k-1}$ are linearly independent. Then, since the method is by definition a conjugate direction method, we have

$$(\text{subspace spanned by } d^0, \ldots, d^{k-1}) = (\text{subspace spanned by } g^0, \ldots, g^{k-1}) \tag{6.19}$$

[cf. Eq. (6.11)]. There are two possibilities:

(a) $g^k = 0$, in which case the method terminates.

(b) $g^k \neq 0$, in which case the expanding manifold minimization property of the conjugate direction method [cf. Eq. (6.9)] implies that

$$g^k \text{ is orthogonal to } d^0, \ldots, d^{k-1}. \tag{6.20}$$

Since the subspaces spanned by $(d^0, \ldots, d^{k-1})$ and by $(g^0, \ldots, g^{k-1})$ are the same [cf. Eq. (6.19)], we see that

$$g^k \text{ is orthogonal to } g^0, \ldots, g^{k-1}. \tag{6.21}$$

Therefore, g^k is linearly independent of $g^0, \ldots, g^{k-1}$, thus completing the induction.

Since at most n linearly independent gradients can be generated, it follows that the gradient will be zero after at most n iterations and the method will terminate with the minimum of f.

To conclude the proof, we use the orthogonality properties (6.20) and (6.21) to verify that the calculation of the coefficients multiplying d^j in the Gram-Schmidt formula (6.18) can be simplified as stated in the proposition. We have for all j such that $g^j \neq 0$,

$$g^{j+1} - g^j = Q(x^{j+1} - x^j) = \alpha^j Q d^j, \tag{6.22}$$

[cf. Eqs. (6.16) and (6.17)]. We note that $\alpha^j \neq 0$, since if $\alpha^j = 0$ we would have $g^{j+1} = g^j$ implying, in view of Eq. (6.21), that $g^j = 0$. Therefore, we have using Eqs. (6.21) and (6.22)

$$g^{i\prime} Q d^j = \frac{1}{\alpha^j} g^{i\prime}(g^{j+1} - g^j) = \begin{cases} 0 & \text{if } j = 0, \ldots, i-2, \\ \frac{1}{\alpha^j} g^{i\prime} g^i & \text{if } j = i-1, \end{cases}$$

and also that

$$d^{j\prime} Q d^j = \frac{1}{\alpha^j} d^{j\prime}(g^{j+1} - g^j).$$

Substituting the last two relations in the Gram-Schmidt formula (6.18) we obtain

$$d^k = -g^k + \beta^k d^{k-1}, \tag{6.23}$$

where

$$\beta^k = \frac{g^{k\prime} g^k}{d^{k-1\prime}(g^k - g^{k-1})}. \tag{6.24}$$

From Eq. (6.23) we have $d^{k-1} = -g^{k-1} + \beta^{k-1} d^{k-2}$. Using this equation, and the orthogonality of g^k and g^{k-1}, and of d^{k-2} and $g^k - g^{k-1}$ [cf. Eqs. (6.20) and (6.21)], the denominator in Eq. (6.24) is written as $g^{k-1\prime} g^{k-1}$, and the desired formula for β^k follows. **Q.E.D.**

Note that by using the orthogonality of g^k and g^{k-1} the formula

$$\beta^k = \frac{g^{k\prime} g^k}{g^{k-1\prime} g^{k-1}} \tag{6.25}$$

of Prop. 1.6.1 can also be written as

$$\beta^k = \frac{g^{k\prime}(g^k - g^{k-1})}{g^{k-1\prime} g^{k-1}}. \tag{6.26}$$

While the alternative formulas (6.25), and (6.26) produce the same results for quadratic problems, their differences become significant when the conjugate gradient method is extended to nonquadratic problems, as we will discuss shortly.

Preconditioned Conjugate Gradient Method

This method is really the conjugate gradient method implemented in a new coordinate system. Suppose we make a change of variables, $x = Sy$, where S is an invertible symmetric $n \times n$ matrix, and we apply the conjugate gradient method to the equivalent problem

$$\begin{aligned} &\text{minimize} \quad h(y) = f(Sy) = \tfrac{1}{2} y' SQSy - b'Sy \\ &\text{subject to} \quad y \in \Re^n. \end{aligned}$$

The method is described by

$$y^{k+1} = y^k + \alpha^k \tilde{d}^k, \tag{6.27}$$

where α^k is obtained by line minimization and $\tilde{d}^k$ is generated by [cf. Eqs. (6.23) and (6.25)]

$$\tilde{d}^0 = -\nabla h(y^0), \qquad \tilde{d}^k = -\nabla h(y^k) + \beta^k \tilde{d}^{k-1}, \qquad k = 1, \ldots, n-1, \tag{6.28}$$

where

$$\beta^k = \frac{\nabla h(y^k)' \nabla h(y^k)}{\nabla h(y^{k-1})' \nabla h(y^{k-1})}. \tag{6.29}$$

Setting $x^k = Sy^k$, $\nabla h(y^k) = Sg^k$, $d^k = S\tilde{d}^k$, and $H = S^2$, we obtain from Eqs. (6.27)-(6.29) the equivalent method

$$x^{k+1} = x^k + \alpha^k d^k, \tag{6.30}$$

$$d^0 = -Hg^0, \qquad d^k = -Hg^k + \beta^k d^{k-1}, \qquad k = 1, \ldots, n-1, \tag{6.31}$$

where

$$\beta^k = \frac{g^{k\prime} H g^k}{g^{k-1\prime} H g^{k-1}}, \tag{6.32}$$

and α^k is obtained by line minimization.

The method described by the above equations is called the *preconditioned conjugate gradient method with scaling matrix H*. To see that this

method is a conjugate direction method, note that since $\nabla^2 h(y) = SQS$, the vectors $\tilde{d}^0, \ldots, \tilde{d}^{n-1}$ are (SQS)-conjugate. Since $d^k = S\tilde{d}^k$, we obtain that $d^0, \ldots, d^{n-1}$ are Q-conjugate. Therefore, the scaled method terminates with the minimum of f after at most n iterations, just as the ordinary conjugate gradient method. The motivation for scaling is to improve the rate of convergence within an n-iteration cycle (see the following analysis). This is important for a nonquadratic problem, but it may be important even for a quadratic problem if n is large and we want to obtain an approximate solution without waiting for the method to terminate.

Application to Nonquadratic Problems

The conjugate gradient method can be applied to the nonquadratic problem

$$\begin{aligned} &\text{minimize} \quad f(x) \\ &\text{subject to} \quad x \in \Re^n. \end{aligned}$$

It takes the form

$$x^{k+1} = x^k + \alpha^k d^k, \tag{6.33}$$

where α^k is obtained by line minimization

$$f(x^k + \alpha^k d^k) = \min_{\alpha} f(x^k + \alpha d^k), \tag{6.34}$$

and d^k is generated by

$$d^k = -\nabla f(x^k) + \beta^k d^{k-1}. \tag{6.35}$$

The most common way to compute β^k is [cf. Eq. (6.26)]

$$\beta^k = \frac{\nabla f(x^k)'\big(\nabla f(x^k) - \nabla f(x^{k-1})\big)}{\nabla f(x^{k-1})'\nabla f(x^{k-1})}. \tag{6.36}$$

The direction d^k generated by the formula $d^k = -\nabla f(x^k) + \beta^k d^{k-1}$ is a direction of descent, since from Eq. (6.34) we obtain $\nabla f(x^k)' d^{k-1} = 0$, so that

$$\nabla f(x^k)' d^k = -\big\|\nabla f(x^k)\big\|^2 + \beta^k \nabla f(x^k)' d^{k-1} = -\big\|\nabla f(x^k)\big\|^2.$$

For nonquadratic problems, the formula (6.36) is typically superior to alternative formulas such as

$$\beta^k = \frac{\nabla f(x^k)'\nabla f(x^k)}{\nabla f(x^{k-1})'\nabla f(x^{k-1})}, \tag{6.37}$$

[cf. Eq. 6.25)]. A heuristic explanation is that due to nonquadratic terms in the objective function and possibly inaccurate line searches, conjugacy

of the generated directions is progressively lost and a situation may arise where the method "jams" in the sense that the generated direction d^k is nearly orthogonal to the gradient $\nabla f(x^k)$. When this occurs, we have $\nabla f(x^{k+1}) \simeq \nabla f(x^k)$. In that case, the scalar β^{k+1}, generated by

$$\beta^{k+1} = \frac{\nabla f(x^{k+1})'\big(\nabla f(x^{k+1}) - \nabla f(x^k)\big)}{\nabla f(x^k)'\nabla f(x^k)},$$

will be nearly zero and the next direction $d^{k+1} = -\nabla f(x^{k+1}) + \beta^{k+1}d^k$ will be close to $-\nabla f(x^{k+1})$ thereby breaking the jam. By contrast, when Eq. (6.37) is used, under the same circumstances the method typically continues to jam.

Regardless of the direction update formula used, one must deal with the loss of conjugacy that results from nonquadratic terms in the cost function. The conjugate gradient method is often employed in problems where the number of variables n is large, and it is not unusual for the method to start generating nonsensical and inefficient directions of search after a few iterations. For this reason it is important to operate the method in cycles of conjugate direction steps, with the first step in the cycle being a steepest descent step. Some possible restarting policies are:

(a) Restart with a steepest descent step n iterations after the preceding restart.

(b) Restart with a steepest descent step k iterations after the preceding restart with $k < n$. This is recommended when the problem has special structure so that the resulting method has good convergence rate (see the following Prop. 1.6.2).

(c) Restart with a steepest descent step if either n iterations have taken place since the preceding restart or if

$$\big|\nabla f(x^k)'\nabla f(x^{k-1})\big| > \gamma\big\|\nabla f(x^{k-1})\big\|^2, \tag{6.38}$$

where γ is a fixed scalar with $0 < \gamma < 1$. The above relation is a test on loss of conjugacy, for if the generated directions were conjugate then we would have $\nabla f(x^k)'\nabla f(x^{k-1}) = 0$.

Note that in all these restart procedures the steepest descent iteration serves as a spacer step and guarantees global convergence (Prop. 1.2.6 in Section 1.2). If the scaled version of the conjugate gradient method is used, then a scaled steepest descent iteration is used to restart a cycle. The scaling matrix may change at the beginning of a cycle but should remain unchanged during the cycle.

An important practical issue relates to the line search accuracy that is necessary for efficient computation. On one hand, an accurate line search is needed to limit the loss of direction conjugacy and the attendant deterioration of convergence rate. On the other hand, insisting on a very accurate

line search can be computationally expensive. Some trial and error may therefore be required in practice. For a discussion of implementations that are tolerant of line search inaccuracies see [Per78] and [Sha78]. For a computational study comparing different implementations, see [PaG86].

Conjugate Gradient-Like Methods for Linear Systems*

The conjugate gradient method can be used to solve the linear system of equations

$$Ax = b,$$

where A is an invertible $n \times n$ matrix and b is a given vector in $\Re^n$. One way to do this is to apply the conjugate gradient method to the positive definite quadratic optimization problem

$$\begin{aligned} &\text{minimize} \quad \tfrac{1}{2}x'A'Ax - b'Ax \\ &\text{subject to} \quad x \in \Re^n, \end{aligned}$$

which corresponds to the equivalent linear system $A'Ax = A'b$. This, however, has several disadvantages, including the need to form the matrix $A'A$, which may have a much less favorable sparsity structure than A.

An alternative possibility is to introduce the vector z defined by

$$x = A'z$$

and to solve the system $AA'z = b$ or equivalently the positive definite quadratic problem

$$\begin{aligned} &\text{minimize} \quad \tfrac{1}{2}z'AA'z - b'z \\ &\text{subject to} \quad z \in \Re^n, \end{aligned}$$

whose cost function gradient is zero at z if and only if $AA'z = b$. By streamlining the computations, it is possible to write the conjugate gradient method for the preceding problem directly in terms of the vector x, and without explicitly forming the product AA'. The resulting method is known as *Craig's method*; it is given by the following iteration where H is a positive definite symmetric preconditioning matrix:

$$x^{k+1} = x^k + \alpha^k d^k, \qquad \alpha^k = \frac{r^{k\prime} r^k}{d^{k\prime} d^k},$$

where the vectors r^k and d^k are generated by the recursions

$$r^{k+1} = r^k + \alpha^k HAd^k, \qquad d^{k+1} = -A'Hr^{k+1} + \frac{r^{k+1\prime} r^{k+1}}{r^{k\prime} r^k} d^k$$

with the initial conditions

$$r^0 = H(Ax^0 - b), \qquad d^0 = -A'Hr^0.$$

The verification of these equations is left for the reader.

There are other conjugate gradient-like methods for the system $Ax = b$, which are not really equivalent to the conjugate gradient method for any quadratic optimization problem. One possibility, due to [SaS86], known as the *Generalized Minimum Residual* method (GMRES), is to start with a vector x^0 and obtain x^k as the vector that minimizes $\|Ax - b\|^2$ over the linear manifold $x^0 + S^k$, where

$$S^k = (\text{subspace spanned by the vectors } r, Ar, A^2r, \ldots, A^{k-1}r),$$

and r is the initial residual

$$r = Ax^0 - b.$$

This successive subspace minimization process can be efficiently implemented, but we will not get into the details further (see [SaS86]). It can be shown that x^k is a solution of the system $Ax = b$ if and only if $A^k r$ belongs to the subspace S^k (write the minimization of $\|Ax - b\|^2$ over $x^0 + S^k$ as the equivalent minimization of $\|\xi - r\|^2$ over all ξ in the subspace AS^k). Thus if none of the vectors $x^0, \ldots, x^{n-2}$ is a solution, the subspace S^{n-1} is equal to $\Re^n$, implying that x^{n-1} is an unconstrained minimum of $\|Ax - b\|^2$, and therefore solves the system $Ax = b$. It follows that the method will terminate after at most n iterations.

GMRES can be viewed as a conjugate gradient method only in the special case where A *is positive definite and symmetric*. In that case it can be shown that the method is equivalent to a preconditioned conjugate gradient method applied to the quadratic cost $\|Ax - b\|^2$. This is based on the expanding subspace minimization property of the conjugate gradient method (see also Exercise 6.4). Note, however, that GMRES can be used for any matrix A that is invertible.

Rate of Convergence of the Conjugate Gradient Method*

There are a number of convergence rate results for the conjugate gradient method. Since the method terminates in at most n steps for a quadratic cost, one would expect that when viewed in cycles of n steps, its rate of convergence for a nonquadratic cost would be comparable to the rate of Newton's method. Indeed there are results which roughly state that if the method is restarted every n iterations and $\{x^k\}$ converges to a nonsingular local minimum x^*, then the error $e^k = \|x^{nk} - x^*\|$ converges superlinearly. (Note that here the error is considered at the end of cycles of n iterations rather than at the end of each iteration.) Such results are reassuring but not terribly interesting because the conjugate gradient method is most useful in problems where n is large (see the discussion at the end of Section 1.7), and for such problems, one hopes that practical

convergence will occur after fewer than n iterations. Therefore, the single-step rate of convergence of the method is more interesting than its rate of convergence in terms of n-step cycles. The following analysis gives a result of this type, based on an interpretation of the conjugate gradient method as an optimal process.

Assume that the cost is positive definite quadratic of the form

$$f(x) = \tfrac{1}{2}x'Qx.$$

(To simplify the following exposition, we have assumed that the linear term $b'x$ is zero, but with minor modifications, the following analysis holds also when $b \neq 0$.) Let g^i denote as usual the gradient $\nabla f(x^i)$ and consider an algorithm of the form

$$\begin{aligned} x^1 &= x^0 + \gamma^{00}g^0, \\ x^2 &= x^0 + \gamma^{10}g^0 + \gamma^{11}g^1, \\ &\cdots \qquad \cdots \\ x^{k+1} &= x^0 + \gamma^{k0}g^0 + \cdots + \gamma^{kk}g^k, \end{aligned} \tag{6.39}$$

where γ^{ij} are arbitrary scalars. Since $g^i = Qx^i$, we see that for suitable scalars c^{ki}, the above algorithm can be written for all k as

$$x^{k+1} = x^0 + c^{k0}Qx^0 + c^{k1}Q^2x^0 + \cdots + c^{kk}Q^{k+1}x^0 = \big(I + QP^k(Q)\big)x^0,$$

where P^k is a polynomial of degree k.

Among algorithms of the form (6.39), the conjugate gradient method is optimal in the sense that for every k, it minimizes $f(x^{k+1})$ over all sets of coefficients $\gamma^{k0}, ..., \gamma^{kk}$. It follows from the equation above that in the conjugate gradient method we have, for every k,

$$f(x^{k+1}) = \min_{P^k} \tfrac{1}{2}x^{0\prime}Q\big(I + QP^k(Q)\big)^2x^0. \tag{6.40}$$

Let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of Q, and let $e_1, \ldots, e_n$ be corresponding orthogonal eigenvectors, normalized so that $\|e_i\| = 1$. Since $e_1, \ldots, e_n$ form a basis, any vector $x^0 \in \Re^n$ can be written as

$$x^0 = \sum_{i=1}^{n} \xi_i e_i$$

for some scalars ξ_i. Since

$$Qx^0 = \sum_{i=1}^{n} \xi_i Qe_i = \sum_{i=1}^{n} \xi_i \lambda_i e_i,$$

we have, using the orthogonality of $e_1, \ldots, e_n$ and the fact $\|e_i\| = 1$,

$$f(x^0) = \tfrac{1}{2}x^{0\prime}Qx^0 = \tfrac{1}{2}\left(\sum_{i=1}^{n}\xi_i e_i\right)'\left(\sum_{i=1}^{n}\xi_i\lambda_i e_i\right) = \tfrac{1}{2}\sum_{i=1}^{n}\lambda_i\xi_i^2.$$

Applying the same process to Eq. (6.40), we obtain for *any* polynomial P^k of degree k

$$f(x^{k+1}) \le \tfrac{1}{2}\sum_{i=1}^{n}\big(1+\lambda_i P^k(\lambda_i)\big)^2\lambda_i\xi_i^2,$$

and it follows that

$$f(x^{k+1}) \le \max_i\big(1+\lambda_i P^k(\lambda_i)\big)^2 f(x^0), \qquad \forall\, P^k,\ k. \tag{6.41}$$

One can use this relationship for different choices of polynomials P^k to obtain a number of convergence rate results. We provide one such result, which shows that the first k conjugate gradient iterations in an n-iteration cycle eliminate the effect of the k largest eigenvalues of Q.

Proposition 1.6.2: Assume that Q has $n-k$ eigenvalues in an interval $[a,b]$ with $a>0$, and the remaining k eigenvalues are greater than b. Then for every x^0, the vector x^{k+1} generated after $k+1$ steps of the conjugate gradient method satisfies

$$f(x^{k+1}) \le \left(\frac{b-a}{b+a}\right)^2 f(x^0). \tag{6.42}$$

This relation also holds for the preconditioned conjugate gradient method (6.30)-(6.32) if the eigenvalues of Q are replaced by those of $H^{1/2}QH^{1/2}$.

Proof: Let $\lambda_1, \ldots, \lambda_k$ be the eigenvalues of Q that are greater than b and consider the polynomial P^k defined by

$$1+\lambda P^k(\lambda) = \frac{2}{(a+b)\lambda_1\cdots\lambda_k}\left(\frac{a+b}{2}-\lambda\right)(\lambda_1-\lambda)\cdots(\lambda_k-\lambda). \tag{6.43}$$

Since $1+\lambda_i P^k(\lambda_i) = 0$, we have, using Eqs. (6.41), (6.43), and a simple calculation,

$$f(x^{k+1}) \le \max_{a\le\lambda\le b}\frac{\big(\lambda-\frac{1}{2}(a+b)\big)^2}{\big(\frac{1}{2}(a+b)\big)^2}f(x^0) = \left(\frac{b-a}{b+a}\right)^2 f(x^0).$$

Q.E.D.

One consequence of the above proposition is that if the eigenvalues of Q take only k distinct values then the conjugate gradient method will find the minimum of the quadratic function f in at most k iterations. (Take $a = b$.) Some other possibilities are explored in the problem section.

It is worth mentioning two more rate of convergence results regarding the conjugate gradient method as applied to the positive definite quadratic function

$$f(x) = \tfrac{1}{2}(x - x^*)'Q(x - x^*).$$

Let M and m be the largest and smallest eigenvalues of f, respectively. Then, for any starting point x^0 and any iteration index k, it can be shown (see [Pol87]) that

$$\|x^k - x^*\| \le 2\left(\frac{M}{m}\right)^{1/2}\left(\frac{\sqrt{M}-\sqrt{m}}{\sqrt{M}+\sqrt{m}}\right)^k \|x^0 - x^*\|, \tag{6.44}$$

$$f(x^k) \le \frac{M\|x^0 - x^*\|^2}{2(2k+1)^2}. \tag{6.45}$$

These relations again suggest a more favorable convergence rate than the one of steepest descent; compare with the results of Section 1.3.

EXERCISES

6.1

Show that the Gram-Schmidt procedure has the projection property stated in Fig. 1.6.2.

6.2 [Ber74]

Let Q have the form

$$Q = M + \sum_{i=1}^{k} v_i v_i',$$

where M is positive definite, and v_i are some vectors in $\Re^n$. Show that the vector x^{k+1} generated after $k+1$ steps of the conjugate gradient method satisfies

$$f(x^{k+1}) \le \left(\frac{b-a}{b+a}\right)^2 f(x^0),$$

where a and b are the smallest and largest eigenvalues of M, respectively. Show also that the vector x^{k+1} generated by the preconditioned conjugate gradient method with $H = M^{-1}$ minimizes f. *Hint*: Use the interlocking eigenvalues lemma [Prop. A.18(d) in Appendix A].

6.3 (Hessian with Clustered Eigenvalues [Ber82a])

Assume that Q has all its eigenvalues concentrated at k intervals of the form

$$[z_i - \delta_i, z_i + \delta_i], \qquad i = 1, \ldots, k,$$

where we assume that $\delta_i \geq 0$, $i = 1, \ldots, k$, $0 < z_1 - \delta_1$, and

$$0 < z_1 < z_2 < \cdots < z_k, \qquad z_i + \delta_i \leq z_{i+1} - \delta_{i+1}, \qquad i = 1, \ldots, k-1.$$

Show that the vector x^{k+1} generated after $k+1$ steps of the conjugate gradient method satisfies

$$f(x^{k+1}) \leq Rf(x^0),$$

where

$$R = \max\left\{\frac{\delta_1^2}{z_1^2}, \frac{\delta_2^2(z_2+\delta_2-z_1)^2}{z_1^2 z_2^2}, \ldots, \frac{\delta_k^2(z_k+\delta_k-z_1)^2(z_k+\delta_k-z_2)^2\cdots(z_k+\delta_k-z_{k-1})^2}{z_1^2 z_2^2 \cdots z_k^2}\right\}.$$

6.4

Consider the conjugate gradient method applied to the minimization of $f(x) = \frac{1}{2}x'Qx - b'x$, where Q is positive definite and symmetric. Show that the iterate x^k minimizes f over the linear manifold

$$x^0 + (\text{subspace spanned by } g^0, Qg^0, \ldots, Q^{k-1}g^0),$$

where $g^0 = \nabla f(x^0)$.

6.5

Let f be positive definite quadratic. Consider the following method: The first iteration is a steepest descent iteration with the stepsize determined by line minimization. For $k = 2, \ldots, n$, the kth iteration finds x^k that minimizes f over the two-dimensional subspace spanned by g^k and $x^k - x^{k-1}$. Show that this method is equivalent to the conjugate gradient method.

6.6

Suppose that $d^0, \ldots, d^k$ are Q-conjugate directions, let $x^1, \ldots, x^{k+1}$ be the vectors generated by the corresponding conjugate direction method, and assume that $x^{i+1} \neq x^i$ for all $i = 0, \ldots, k$. Show that a vector d^{k+1} is Q-conjugate to $d^0, \ldots, d^k$ if and only if $d^{k+1} \neq 0$ and d^{k+1} is orthogonal to the gradient differences $g^{i+1} - g^i$, $i = 0, \ldots, k$.

6.7

Describe the behavior of the conjugate gradient method for a positive semidefinite quadratic function. Consider the case where there is no optimal solution and the case where there are infinitely many optimal solutions.

6.8

Let $f(x) = \frac{1}{2}x'Qx - b'x$, where Q is positive definite and symmetric. Suppose that x_1 and x_2 minimize f over linear manifolds that are parallel to subspaces S_1 and S_2, respectively. Show that if $x_1 \neq x_2$, then $x_1 - x_2$ is Q-conjugate to all vectors in the intersection of S_1 and S_2. Use this property to construct a conjugate direction method that does not evaluate gradients and uses only line minimizations.

1.7 QUASI-NEWTON METHODS

Quasi-Newton methods are gradient methods of the form

$$x^{k+1} = x^k + \alpha^k d^k, \tag{7.1}$$

$$d^k = -D^k \nabla f(x^k), \tag{7.2}$$

where D^k is a positive definite matrix, which may be adjusted from one iteration to the next so that the direction d^k tends to approximate the Newton direction. Some of these methods are quite popular because they typically converge fast, while avoiding the second derivative calculations associated with Newton's method. Their main drawback relative to the conjugate gradient method is that they require storage of the matrix D^k as well as the matrix-vector multiplication overhead associated with the calculation of the direction d^k (see the subsequent discussion).

An important idea for many quasi-Newton methods is that two successive iterates x^k, x^{k+1} together with the corresponding gradients $\nabla f(x^k)$, $\nabla f(x^{k+1})$, yield curvature information by means of the approximate relation

$$q^k \approx \nabla^2 f(x^{k+1}) p^k, \tag{7.3}$$

where

$$p^k = x^{k+1} - x^k, \tag{7.4}$$

$$q^k = \nabla f(x^{k+1}) - \nabla f(x^k). \tag{7.5}$$