ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

GENERALIZATIONS OF THE TRUST REGION PROBLEM

Jorge J. Moré

Mathematics and Computer Science Division

Preprint MCS-P349-0193

January 1993

This work was supported by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

ABSTRACT

The trust region problem requires the global minimum of a general quadratic function subject to an ellipsoidal constraint. The development of algorithms for the solution of this problem has found applications in nonlinear and combinatorial optimization. In this paper we generalize the trust region problem by allowing a general quadratic constraint. The main results are a characterization of the global minimizer of the generalized trust region problem, and the development of an algorithm that finds an approximate global minimizer in a finite number of iterations.

GENERALIZATIONS OF THE TRUST REGION PROBLEM

Jorge J. Moré

1 Introduction

Trust region methods for the minimization of a function $f : \mathbb{R}^n \to \mathbb{R}$ require an approximate minimizer of a problem of the form

$$\min\{q(x) : \|Dx\|_2 \leq \Delta\}, \tag{1.1}$$

where $q : \mathbb{R}^n \to \mathbb{R}$ is a quadratic model of the function in a neighborhood of the current iterate, D is a nonsingular matrix that specifies the scaling of the variables, and Δ is determined by the trust region method. An interesting and surprising aspect of this problem is that it is possible to obtain a characterization of the global minimizer. This result, due to Gay [5] and Sorensen [17], is that x^* is a global minimizer if and only if $\|Dx^*\| \leq \Delta$, the Kuhn-Tucker condition

$$\nabla q(x^*) + \lambda^* D^T D x^* = 0$$

is satisfied, and the Hessian of the Lagrangian

$$\nabla^2 q(x^*) + \lambda^* D^T D \tag{1.2}$$

is positive semidefinite for some $\lambda^* \geq 0$ such that $\lambda^* = 0$ if $\|Dx^*\| < \Delta$. The surprising aspect of this result is that it holds for any quadratic q; there is no assumption of convexity.

Gay [5] and Sorensen [17] used this characterization result to develop an algorithm for the solution of (1.1) that paid special attention to the case where the Hessian of the Lagrangian (1.2) is nearly singular. Moré and Sorensen [15] developed an improved algorithm that produces approximate global minimizers in a finite number of steps.

The initial application of algorithms for the solution of (1.1) was to trust region methods for unconstrained minimization. See, for example, Gay [6] and Moré and Sorensen [15]. Recent application to other problems includes that by Coleman and Hempel [2] on constrained optimization, Ye [20] on quadratic programming, Pardalos, Ye, and Han [16] on quadratic knapsack problems, and Karmarkar, Resende, and Ramakrishnan [12] on covering problems.

Several extensions of the basic problem (1.1) are of interest. The simplest extension is to consider the problem

$$\min\{q(x) : c(x) \leq 0\}, \tag{1.3}$$

Dedicated to Charles G. Broyden on the occasion of his 60th birthday. Work supported by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

where $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ are quadratic functions defined on $\mathbb{R}^n$. Several special cases of this problem have been considered in the literature. For example, the problem where

$$\min \left\{ \|Ax - b\|_2 : \|Cx - d\|_2 \le \alpha \right\}, \tag{1.4}$$

arises in the regularization of ill-posed problems and in the smoothing of noisy data. Note that this problem is a special case of (1.3) with both q and c convex. Gander [4] considered this problem as a special case of the problem with the equality constraint $c(x) = 0$, and characterized the solution in terms of the Lagrange multiplier. For recent work on this problem, see the work of Golub and von Matt [7].

Critchley [3] considered (1.3) with the equality constraint $c(x) = 0$, with q a strictly convex quadratic, but c a general quadratic. His interest in problem (1.3) derived from applications in the statistical literature. Our interest in (1.3) is motivated by a study of the general problem

$$\min\{q(x) : c_i(x) = 0,\ 1 \le i \le m\},$$

where $c_i : \mathbb{R}^n \to \mathbb{R}$ are quadratic functions. This generalization is important because it includes, as a special case, integer programming problems. Indeed, if $c_i(x) = x_i(1-x_i)$, then $c_i(x) = 0$ if and only if $x_i \in \{0, 1\}$. The case where $m = 2$ and the equality constraints are replaced by inequalities has attracted considerable interest in connection with trust region methods. For recent work on this problem, see Yuan [21, 22], Martínez and Santos [14], Zhang [23], and Heinkenschloß [9].

Our aim in this paper is to characterize the global minimizer of problem (1.3) and to develop an algorithm that determines approximate global minimizers in a finite number of steps. In our terminology, an approximate global minimizer is a vector $x \in \mathbb{R}^n$ such that

$$q(x) \le q^* + \epsilon_q, \qquad |c(x)| \le \epsilon_c,$$

where $\epsilon_q > 0$ and $\epsilon_c > 0$ are tolerances, and $q^* \equiv \min\{q(x) : c(x) \le 0\}$.

We approach the inequality constrained problem (1.3) by first studying the equality constrained problem

$$\min\{q(x) : c(x) = 0\}, \tag{1.5}$$

and then showing that (1.3) can be treated as a special case. We begin with a brief look at a condition that guarantees the existence of a global minimizer for (1.5) and the connection of this condition with matrix pencils. The characterization result for a global minimizer of problems (1.3) and (1.5) appears in Section 3, while a uniqueness result for the minimizer x^* appears in Section 4.

In Section 3 we also show that if the inequality constrained problem (1.3) is well-posed, then there is a global minimizer x^* such that $c(x^*) = 0$, or there is a unique minimizer x^* with $\nabla^2 q$ positive definite, $\nabla q(x^*) = 0$, and $c(x^*) < 0$. This result shows that if we develop an algorithm for the equality constrained problem (1.5), then we can easily extend the algorithm to the inequality constrained problem (1.3).

The essential features of an algorithm for the solution of both (1.3) and (1.5) are developed in Section 3, while Section 5 contains a discussion of the termination criteria that guarantee an approximate global minimizer in a finite number of steps. These results extend and unify the theory associated with problems (1.1) and (1.4) to the general problem (1.3) and (1.5).

2 Preliminaries

We can guarantee that the minimization problem

$$\min\{q(x) : c(x) = 0\}, \tag{2.1}$$

where $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ are quadratic functions defined on $\mathbb{R}^n$, has a global minimizer by assuming that the Hessian of q is positive definite for any direction of zero curvature of c. In this section we explore the connection of this condition with work on matrix pencils.

The definition of a quadratic guarantees that the Hessian matrix is constant. We use the notation

$$\nabla^2 q = A, \qquad \nabla^2 c = C,$$

for the Hessian matrices of q and c. The following result shows that the relationship between A and C is crucial to the existence of a global minimizer for problem (2.1).

Theorem 2.1 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be quadratics, and assume that*

$$\{x \in \mathbb{R}^n : c(x) = 0\}$$

is not empty. If

$$w \neq 0, \quad w^T C w = 0 \quad \Longrightarrow \quad w^T A w > 0, \tag{2.2}$$

then the optimization problem (2.1) *has a global minimizer.*

Proof. A standard compactness argument shows that (2.1) has a global minimizer if we prove that

$$\lim_{k \to \infty} q(x_k) = +\infty$$

for any unbounded $\{x_k\}$ with $c(x_k) = 0$. Assume, on the contrary, that $q(x_k) \leq \beta$. There is no loss of generality in assuming that

$$\lim_{k \to \infty} \frac{x_k}{\|x_k\|} = w.$$

Clearly $w \neq 0$. Since $\{x_k\}$ is an unbounded sequence with $c(x_k) = 0$,

$$\lim_{k \to \infty} \frac{c(x_k)}{\|x_k\|^2} = w^T C w = 0.$$

Similarly, since $\{x_k\}$ is an unbounded sequence with $q(x_k) \leq \beta$,

$$\lim_{k \to \infty} \frac{q(x_k)}{\|x_k\|^2} = w^T A w \leq 0.$$

Thus, w contradicts assumption (2.2) on A and C. ■

The converse of Theorem 2.1 does not hold in general; however, in Section 3 we show that if the optimization problem (2.1) has a global minimizer, then $w^T C w = 0$ implies that $w^T A w \geq 0$. In the remainder of this section we consider the connection of condition (2.2) with the study of matrix pencils, that is, one-parameter families of matrices of the form $A + \lambda C$.

Theorem 2.2 *If $A \in \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{n \times n}$ are symmetric matrices, then $A + \lambda C$ is positive definite for some $\lambda \in \mathbb{R}$ if and only if* (2.2) *holds.*

Theorem 2.2 is a classical result on quadratic forms. Hestenes [10, Theorem 6.1] gives an elementary, but somewhat lengthy proof of Theorem 2.2. We will give another proof of this result in Section 3. Uhlig [19] is an excellent source for additional information on results related to Theorem 2.2.

Theorem 2.2 arises in several areas. For example, in the analysis of augmented Lagrangian algorithms, A is the Hessian of the Lagrangian and C is a positive semidefinite matrix associated with the active constraint normals. The proof of Theorem 2.2 when C is semidefinite is short and elementary; see, for example, Bertsekas [1, Lemma 1.25]. Theorem 2.2 also arises in connection with the generalized eigenvalue problem $Ax = \lambda C x$, because if we can find $\lambda_0 \in \mathbb{R}$ such that $A + \lambda_0 C$ is positive definite, then the eigenvalue problem $Ax = \lambda C x$ is equivalent to the positive definite eigenvalue problem

$$(\lambda_0 A - C)x = \nu (A + \lambda_0 C) x.$$

This implies, in particular, that $Ax = \lambda C x$ has n real eigenvalues; some of the eigenvalues may be at infinity if C is singular. Golub and Van Loan [8, Section 8.7] provide additional information on generalized eigenvalue problems.

Theorem 2.3 *Assume that $A \in \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{n \times n}$ are symmetric matrices and that C is indefinite. Then*

$$w^T C w = 0 \quad \Longrightarrow \quad w^T A w \geq 0$$

if and only if $A + \lambda C$ is positive semidefinite for some $\lambda \in \mathbb{R}$.

Proof. If we define $A_\alpha = A + \alpha I$ for $\alpha \geq 0$, then (2.2) holds for A_α, and thus Theorem 2.2 guarantees that $A_\alpha + \lambda_\alpha C$ is positive definite for some $\lambda_\alpha \in \mathbb{R}$. In particular,

$$\lambda_\alpha w^T C w \geq -w^T (A + \alpha I) w$$

for all $w \in \mathbb{R}^n$. This inequality shows that if $\lambda_\alpha \to +\infty$ as $\alpha \to 0$ then C is positive semidefinite. Similarly, we can rule out that $\lambda_\alpha \to -\infty$ as $\alpha \to 0$. Since $\{\lambda_\alpha\}$ is bounded, a subsequence converges to some $\lambda \in \mathbb{R}$, and since $A_\alpha + \lambda_\alpha C$ is positive definite, $A + \lambda C$ is positive semidefinite. ■

This result was established by Krein and Smuljan [13] for Hilbert spaces over the complex field by an entirely different argument. Note that our proof is valid for real vector spaces.

The assumption that C is indefinite is necessary in Theorem 2.3. For example, let

$$A = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad C = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

We can check that $w^T C w = 0$ implies that $w^T A w = 0$. However, $A + \lambda C$ is indefinite for all $\lambda \in \mathbb{R}$. Also note that Theorem 2.3 does not guarantee that $A + \lambda C$ is positive semidefinite for more than one value of λ. For example, if

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \qquad C = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

then A is positive semidefinite, and thus the assumptions of Theorem 2.3 hold. However, $A + \lambda C$ is positive semidefinite only if $\lambda = 0$.

3 Global Minimizers

The main purpose of this section is to establish a characterization result for the global minimizer of the problem

$$\min\{q(x) : c(x) = 0\}, \tag{3.1}$$

where $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ are quadratic functions defined on $\mathbb{R}^n$. We conclude with the characterization result for the inequality constrained version of problem (3.1).

We follow the notational convention of the preceding section by using

$$\nabla^2 q = A, \qquad \nabla^2 c = C,$$

for the Hessian matrices of q and c. We do not assume that either A or C is semidefinite. Interestingly enough, we need to assume that the constraint quadratic c is not a linear function; that is, we assume that $C \neq 0$.

The analysis requires that we rule out certain special cases by assuming a constraint qualification. For (3.1) we assume that

$$\min\{c(x) : x \in \mathbb{R}^n\} < 0 < \max\{c(x) : x \in \mathbb{R}^n\}. \tag{3.2}$$

These assumptions guarantee, in particular, that the feasible set is not empty.

Lemma 3.1 *Let $c : \mathbb{R}^n \to \mathbb{R}$ be a quadratic function defined on $\mathbb{R}^n$. If the feasible set $\{x \in \mathbb{R}^n : c(x) = 0\}$ is not empty, then assumption (3.2) fails if and only if C is semidefinite and there is an x^* such that $c(x^*) = 0$ and $\nabla c(x^*) = 0$. Moreover, if $\nabla c(x) \neq 0$ for some feasible $x \in \mathbb{R}^n$ then assumption (3.2) holds.*

Proof. Assume that (3.2) fails because $c(x) \geq 0$ for all $x \in \mathbb{R}^n$. Since the feasible set is not empty, we must have

$$\min\{c(x) : x \in \mathbb{R}^n\} = 0.$$

Thus, the quadratic c is bounded below and achieves its minimum at some x^*. Hence, C is positive semidefinite, $c(x^*) = 0$, and $\nabla c(x^*) = 0$. A similar argument establishes the result if (3.2) fails because $c(x) \leq 0$ for all $x \in \mathbb{R}^n$.

If $\nabla c(x) \neq 0$ for some feasible $x \in \mathbb{R}^n$ but (3.2) fails because $c(x) \geq 0$ for all $x \in \mathbb{R}^n$ then the above argument shows that the quadratic c achieves its minimum at an x^* such that $c(x^*) = 0$. Hence, c achieves its minimum at any feasible x, and thus $\nabla c(x) = 0$ for any feasible x. ■

This result implies that if assumption (3.2) fails, then the feasible set $\{x \in \mathbb{R}^n : c(x) = 0\}$ is a subspace of $\mathbb{R}^n$ with dimension less than n; if C is definite then only x^* is feasible.

Theorem 3.2 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be quadratic functions defined on $\mathbb{R}^n$. Assume that (3.2) holds and that $\nabla^2 c \neq 0$. A vector x^* is a global minimizer of problem (3.1) if and only if $c(x^*) = 0$ and there is a multiplier $\lambda^* \in \mathbb{R}$ such that the Kuhn-Tucker condition*

$$\nabla q(x^*) + \lambda^* \nabla c(x^*) = 0 \tag{3.3}$$

is satisfied with

$$\nabla^2 q(x^*) + \lambda^* \nabla^2 c(x^*) \tag{3.4}$$

positive semidefinite.

Proof. We first show that if x^* is a global minimizer, then condition (3.3) holds and (3.4) is positive semidefinite. The proof requires consideration of two cases.

If $\nabla c(x^*) \neq 0$, then the Kuhn-Tucker condition (3.3) holds, so we need to establish that (3.4) is positive semidefinite. In terms of the Lagrangian function

$$\mathcal{L}(x, \lambda) = q(x) + \lambda c(x),$$

the Kuhn-Tucker condition guarantees that $\nabla_x \mathcal{L}(x^*, \lambda^*) = 0$, and thus

$$\mathcal{L}(x, \lambda^*) = \mathcal{L}(x^*, \lambda^*) + \tfrac{1}{2}(x - x^*)^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*)(x - x^*).$$

Since $\mathcal{L}(x, \lambda^*) = q(x)$ if $c(x) = 0$, we obtain that

$$c(x^* + w) = 0 \quad \Longrightarrow \quad w^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*) w \geq 0.$$

We use this implication to show that $\nabla_x^2 \mathcal{L}(x^*, \lambda^*)$ is positive semidefinite. The proof requires that we partition $\mathbb{R}^n$ into four different sets and show that

$$w^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*) w \geq 0 \tag{3.5}$$

holds for vectors in each of these sets. Consider the set

$$S_1 = \left\{ w \in \mathbb{R}^n : \nabla c(x^*)^T w \neq 0, \ w^T C w \neq 0 \right\}.$$

If $w \in S_1$, then $c(x^* + \alpha w) = 0$ for some $\alpha \neq 0$, and thus (3.5) holds. Similarly, if

$$S_2 = \left\{ w \in \mathbb{R}^n : \nabla c(x^*)^T w = 0, \ w^T C w = 0 \right\},$$

then $c(x^* + w) = 0$ for $w \in S_2$, and thus (3.5) holds. Now consider the set

$$S_3 = \left\{ w \in \mathbb{R}^n : \nabla c(x^*)^T w = 0, \ w^T C w \neq 0 \right\}.$$

If $w \in S_3$, define $w_\alpha = w + \alpha v$, where $\nabla c(x^*)^T v \neq 0$. The existence of the vector v is guaranteed by the assumption that $\nabla c(x^*) \neq 0$. A short computation shows that $\nabla c(x^*)^T w_\alpha \neq 0$ and that $w_\alpha^T C w_\alpha \neq 0$ for all $\alpha \neq 0$ sufficiently small. Hence $w_\alpha \in S_1$, and thus

$$w_\alpha^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*) w_\alpha \geq 0. \tag{3.6}$$

Since $\{w_\alpha\}$ converges to w as α converges to zero, (3.5) holds for $w \in S_3$. A similar argument show that (3.5) holds if w belongs to the set

$$S_4 = \left\{ w \in \mathbb{R}^n : \nabla c(x^*)^T w \neq 0, \ w^T C w = 0 \right\}.$$

If $w \in S_4$, we now define $w_\alpha = w + \alpha v$, where $v^T C v \neq 0$. In this case the existence of the vector v is guaranteed by the assumption that $\nabla^2 c \neq 0$, since we can choose v to be an eigenvector of C corresponding to a nonzero eigenvalue. A computation shows that $w_\alpha \in S_1$ for all $\alpha \neq 0$ sufficiently small, and thus (3.6) holds. Since $\{w_\alpha\}$ converges to w as α converges to zero, (3.5) holds for $w \in S_4$.

We have shown that (3.5) holds for all $w \in \mathbb{R}^n$, and thus (3.3) holds if $\nabla c(x^*) \neq 0$. Note that assumptions (3.2) have not been needed so far in the proof.

The proof for the case where $\nabla c(x^*) = 0$ requires that we first show that $\nabla q(x^*) = 0$. We first claim that $q(x^* + \alpha w) \geq q(x^*)$ for all α and $w \in \mathbb{R}^n$ such that $w^T C w = 0$. This claim is easy to establish because $c(x^*) = 0$ and $\nabla c(x^*) = 0$ implies that

$$c(x^* + \alpha w) = \tfrac{1}{2} \alpha^2 w^T C w = 0.$$

Hence, $x^* + \alpha w$ is feasible, and thus $q(x^* + \alpha w) \geq q(x^*)$ because x^* is a global minimizer.

Since $q(x^* + \alpha w) \geq q(x^*)$ for all α and $w \in \mathbb{R}^n$, we must have $\nabla q(x^*)^T w = 0$ whenever $w^T C w = 0$. We now show that $\nabla q(x^*)^T v = 0$ for all eigenvectors v of C. Note that since x^* satisfies $c(x^*) = 0$ and $\nabla c(x^*) = 0$, assumptions (3.2) imply that C is indefinite.

If v is an eigenvector of C corresponding to a zero eigenvalue, then $v^T C v = 0$, and thus $\nabla q(x^*)^T v = 0$. Let v_1 and v_2 be eigenvectors of C corresponding to positive and negative eigenvalues, respectively. We choose the directions of v_1 and v_2 so that $\nabla q(x^*)^T v_i \geq 0$ for $i = 1, 2$, and we choose an $\alpha < 0$ so that

$$(v_1 - \alpha v_2)^T C (v_1 - \alpha v_2) = 0.$$

Hence, $\nabla q(x^*)^T (v_1 - \alpha v_2) = 0$, and since $\nabla q(x^*)^T v_i \geq 0$ and $\alpha < 0$, we must have $\nabla q(x^*)^T v_i = 0$ for $i = 1, 2$. We have shown that $\nabla q(x^*)^T v = 0$ for all eigenvectors v of C. Since C has a complete set of eigenvectors, $\nabla q(x^*) = 0$.

We now show that $w^T C w = 0$ implies that $w^T A w \geq 0$. If $w^T C w = 0$, then

$$c(x^* + w) = c(x^*) + \nabla c(x^*)^T w + \tfrac{1}{2} w^T C w = 0,$$

since $c(x^*) = 0$ and $\nabla c(x^*) = 0$. Hence, $x^* + w$ is feasible, and since $\nabla q(x^*) = 0$ and x^* is the global minimizer,

$$q(x^* + w) - q(x^*) = \tfrac{1}{2} w^T A w \geq 0.$$

Thus, we have shown that $w^T C w = 0$ implies that $w^T A w \geq 0$. Theorem 2.3 now shows that $A + \lambda^* C$ is positive semidefinite for some $\lambda^* \in \mathbb{R}$, that is, $\nabla_x^2 \mathcal{L}(x^*, \lambda^*)$ is positive semidefinite. Moreover, the pair (x^*, λ^*) trivially satisfies the Kuhn-Tucker condition because $\nabla q(x^*) = 0$ and $\nabla c(x^*) = 0$.

The converse is easy to prove. If $c(x^*) = 0$ and there is a multiplier $\lambda^* \in \mathbb{R}$ such that (3.3) holds and (3.4) is positive semidefinite, then the Kuhn-Tucker condition guarantees that $\nabla_x \mathcal{L}(x^*, \lambda^*) = 0$, and thus

$$\mathcal{L}(x, \lambda^*) = \mathcal{L}(x^*, \lambda^*) + \tfrac{1}{2}(x - x^*)^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*)(x - x^*) \geq \mathcal{L}(x^*, \lambda^*).$$

Since $\mathcal{L}(x, \lambda^*) = q(x)$ if $c(x) = 0$, we obtain that $q(x^*) \leq q(x)$ if $c(x) = 0$. Thus, x^* is a global minimizer. ■

The assumption that $\nabla^2 c \neq 0$ in Theorem 3.2 is necessary. For example, the problem

$$\min\{\xi_1^2 - \xi_2^2 : \xi_2 = 0\}$$

has $x^* = (0, 0)$ for a global minimizer and $\lambda^* = 0$ as the unique multiplier. However,

$$\nabla_x^2 \mathcal{L}(x^*, \lambda^*) = \begin{pmatrix} 2 & 0 \\ 0 & -2 \end{pmatrix}$$

is not positive semidefinite. Also note that assumption (3.2) is necessary. For example, the problem

$$\min\{\xi_1^2 + \xi_2 : \xi_2^2 = 0\}$$

has $x^* = (0,0)$ for a global minimizer, but the Kuhn-Tucker condition (3.3) does not hold for any $\lambda^* \in \mathbb{R}$.

An interesting aspect of the proof of Theorem 3.2 is that it can be used to give a short proof of Theorem 2.2. We need the part of the proof where it is assumed that $\nabla c(x^*) \neq 0$. For this proof, first note that if (2.2) holds, then there is an $\epsilon > 0$ such that

$$w \neq 0, \quad w^T C w = 0 \quad \Longrightarrow \quad w^T A w \geq \epsilon \|w\|^2.$$

If we consider the optimization problem

$$\min\{x^T(A - \tfrac{1}{2}\epsilon I)x : x^T C x = x_0^T C x_0\},$$

for some $x_0 \in \mathbb{R}^n$ such that $x_0^T C x_0 \neq 0$, then Theorem 2.1 shows that this problem has a global minimizer x^*. Moreover, if $c(x) = x^T C x - x_0^T C x_0$, it is clear that $\nabla c(x^*) = 2Cx^* \neq 0$. Hence, the proof of Theorem 3.2 shows that there is a multiplier λ^* such that

$$A - \tfrac{1}{2}\epsilon I + \lambda^* C$$

is positive semidefinite. Hence, $A + \lambda^* C$ is positive definite as desired.

We have obtained Theorem 3.2 under minimal assumptions. Algorithmic developments, however, require the stability assumption that the set I_{PD} defined by

$$I_{\mathrm{PD}} = \{\lambda \in \mathbb{R} : A + \lambda C \text{ positive definite }\} \tag{3.7}$$

is not empty. We justify this assumption by proving that if I_{PD} is empty, then a small perturbation of problem (3.1) leads to a similar problem without a global minimizer. Theorem 2.2 shows that if I_{PD} is empty, then there is a $w \in \mathbb{R}^n$ such that

$$w \neq 0, \quad w^T C w = 0, \quad w^T A w \leq 0.$$

Now consider the minimization problem

$$\min\{q_\epsilon(x) : c(x) = 0\},$$

where

$$q_\epsilon(x) = q(x) - \tfrac{1}{2}\epsilon(v^T x)^2$$

for any vector $v \in \mathbb{R}^n$ such that $v^T w \neq 0$, and any $\epsilon > 0$. The Hessian of the Lagrangian of this problem is not positive semidefinite, and thus Theorem 3.2 implies that this problem does not have a global minimizer. Since $\nabla^2 q_\epsilon$ is arbitrarily close to $\nabla^2 q$, this proves our claim that if I_{PD} is empty, then a small perturbation of problem (3.1) leads to a similar problem without a global minimizer.

The analogous stability assumption for the inequality constrained version of problem (3.1) is that

$$w \neq 0, \quad w^T C w \leq 0 \quad \Longrightarrow \quad w^T A w > 0. \tag{3.8}$$

We do not elaborate on this assumption; we need (3.8) only to show that the inequality constrained version is essentially a special case of (3.1).

Theorem 3.3 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be quadratic functions defined on $\mathbb{R}^n$, and assume that* (3.8) *holds. If the problem*

$$\min\{q(x) : c(x) \leq 0\} \tag{3.9}$$

is feasible, then there is a global minimizer x^* such that $c(x^*) = 0$, or there is a unique minimizer x^* with A positive definite, $\nabla q(x^*) = 0$, and $c(x^*) < 0$.

Proof. A minor modification of the proof of Theorem 2.1 shows that if (3.8) holds, then the optimization problem (3.9) has a global minimizer. If no global minimizer x^* satisfies $c(x^*) = 0$, then there must be a global minimizer x^* with $c(x^*) < 0$. Hence, $\nabla q(x^*) = 0$ and A must be positive semidefinite. If, on the contrary, A is not positive definite, then there is a $v \neq 0$ with $Av = 0$. Thus, $q(x^* + \alpha v) = q(x^*)$ for all $\alpha \in \mathbb{R}$. We complete the proof by showing that $c(x^* + \alpha v) = 0$ for some $\alpha \in \mathbb{R}$. Since $v^T A v = 0$ and (3.8) holds, we must have $v^T C v > 0$. Thus, $c(x^* + \alpha v) \to +\infty$ as $\alpha \to \infty$. The continuity of c and the assumption that $c(x^*) < 0$ now show that there is some $\alpha \in \mathbb{R}$ with $c(x^* + \alpha v) = 0$, as desired. ■

Theorem 3.3 shows that if we develop an algorithm for the equality constrained problem (3.1), then we can easily extend the algorithm to the inequality constrained problem (3.9). We conclude this section with the characterization result for (3.9).

Theorem 3.4 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be quadratic functions defined on $\mathbb{R}^n$, and assume that*

$$\min\{c(x) : x \in \mathbb{R}^n\} < 0 \tag{3.10}$$

and that $\nabla^2 c \neq 0$. A vector x^* is a global minimizer of (3.9) *if and only if $c(x^*) \leq 0$, the Kuhn-Tucker condition* (3.3) *is satisfied, and the Hessian of the Lagrangian* (3.4) *is positive semidefinite for some $\lambda^* \geq 0$ with $\lambda^* = 0$ if $c(x^*) < 0$.*

Proof. First of all, note that (3.2) can fail only if $c(x) \leq 0$ for all $x \in \mathbb{R}^n$. However, in this case the optimization problem is unconstrained, and the result holds.

If x^* is a global minimizer and $c(x^*) < 0$, then x^* is an unconstrained minimizer, and thus $\nabla^2 q(x^*)$ is positive semidefinite as desired. If $c(x^*) = 0$, the result follows from Theorem 3.2.

The converse follows as in Theorem 3.2. If (3.3) holds and (3.4) is positive semidefinite for some $\lambda^* \geq 0$, then $\nabla_x \mathcal{L}(x^*, \lambda^*) = 0$ and

$$\mathcal{L}(x, \lambda^*) = \mathcal{L}(x^*, \lambda^*) + \tfrac{1}{2}(x - x^*)^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*)(x - x^*) \geq \mathcal{L}(x^*, \lambda^*).$$

Since $\lambda^* \geq 0$, this implies that

$$q(x) \geq q(x) + \lambda^* c(x) = \mathcal{L}(x, \lambda^*) \geq \mathcal{L}(x^*, \lambda^*) = q(x^*)$$

for all x such that $c(x) \leq 0$. Thus, x^* is a global minimizer of problem (3.9). ∎

4 Uniqueness

We only consider conditions that guarantee the uniqueness of x^* since the uniqueness of λ^* is not needed for the developments that follow.

Theorem 4.1 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be quadratic functions defined on $\mathbb{R}^n$, and assume that the Kuhn-Tucker condition* (3.3) *is satisfied and that the Hessian of the Lagrangian* (3.4) *is positive definite. If $c(x^*) = 0$, then x^* is the unique global minimizer of problem* (3.1). *If $\lambda^* \geq 0$ and $c(x^*) \leq 0$, then x^* is the unique global minimizer of problem* (3.9).

Proof. We only prove this result for problem (3.9) since the proof for problem (3.1) is similar. If (3.3) holds, then $\nabla_x \mathcal{L}(x^*, \lambda^*) = 0$. Thus, since (3.4) is positive definite,

$$\mathcal{L}(x, \lambda^*) = \mathcal{L}(x^*, \lambda^*) + \tfrac{1}{2}(x - x^*)^T \nabla_x^2 \mathcal{L}(x^*, \lambda^*)(x - x^*) > \mathcal{L}(x^*, \lambda^*)$$

for all $x \neq x^*$. Now use the assumptions that $\lambda^* \geq 0$ and that $c(x^*) \leq 0$ to obtain that

$$q(x) \geq q(x) + \lambda^* c(x) = \mathcal{L}(x, \lambda^*) > \mathcal{L}(x^*, \lambda^*) = q(x^*)$$

for all $x \neq x^*$ such that $c(x) \leq 0$. Thus, x^* is the unique global minimizer of problem (3.9). ∎

5 Algorithms

We propose an algorithm for the solution of the optimization problem (3.1) that requires the solution of a sequence of positive definite systems of linear equations. The algorithm is based on a search for the Lagrange multiplier λ^* guaranteed by Theorem 3.2.

Consider the optimization problem (3.1) where $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ are quadratic functions defined on $\mathbb{R}^n$. For algorithmic purposes, we assume that the quadratics q and c are defined by

$$q(x) = \tfrac{1}{2} x^T A x + b^T x, \qquad c(x) = \tfrac{1}{2} x^T C x + d^T x - \delta. \tag{5.1}$$

We also assume that the set I_{PD} defined by (3.7) is not empty. This is a reasonable assumption because the argument in Section 3 showed that if I_{PD} was empty, then the optimization problem (3.1) was not well-posed.

Given $\lambda \in I_{\mathrm{PD}}$, we define $x(\lambda) \in \mathbb{R}^n$ as the solution to the system of linear equations

$$\nabla q[x(\lambda)] + \lambda \nabla c[x(\lambda)] = 0. \tag{5.2}$$

In terms of the Lagrangian function, $x(\lambda)$ is the unique solution of $\nabla_x \mathcal{L}[x(\lambda), \lambda] = 0$; in view of (5.1) we can also think of $x(\lambda)$ as the solution of the system of linear equations

$$(A + \lambda C) x(\lambda) = -(b + \lambda d).$$

The algorithm that we propose in this section is based on finding $\lambda \in I_{\mathrm{PD}}$ such that

$$c[x(\lambda)] = 0.$$

Theorem 3.2 shows that if this is possible, then $x(\lambda)$ is the global minimizer of the optimization problem (3.1).

Theorem 5.1 *If $A \in \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{n \times n}$ are symmetric matrices, then I_{PD} is an interval.*

Proof. We show that I_{PD} is convex and thus an interval. Assume that λ_1 and λ_2 belong to I_{PD} and that $\lambda_1 < \lambda_2$. Since $w^T(A + \theta C)w$ is linear in θ for any $w \in \mathbb{R}^n$,

$$w^T (A + \theta C) w \geq \min\{ w^T (A + \lambda_i) w : i = 1, 2\}, \quad \theta \in (\lambda_1, \lambda_2),$$

and thus $A + \theta C$ is positive definite. Hence, $\theta \in I_{\mathrm{PD}}$ as desired. ∎

Theorem 2.2 gives necessary and sufficient conditions for I_{PD} to be nonempty. Under these conditions Theorem 5.1 shows that I_{PD} is a nonempty open interval. I_{PD} is a finite interval (λ_l, λ_u) if C is indefinite. We can see this by noting that

$$\lambda w^T C w \geq -w^T A w, \quad w \in \mathbb{R}^n$$

for $\lambda \in I_{\mathrm{PD}}$. In a similar manner we can prove that I_{PD} is of the form $(\lambda_l, +\infty)$ if C is positive definite, while $I_{\mathrm{PD}} = (-\infty, \lambda_u)$ if C is negative definite.

For several of our results it will be important to note that each component of $x(\cdot)$ is a rational function for $\lambda \in I_{\mathrm{PD}}$. This result can be established by first recalling the classical result that since $I_{\rm PD}$ is not empty, there is a nonsingular similarity transformation that diagonalizes A and C. For a proof of this result, see Horn and Johnson [11, Theorem 7.6.1]. Thus, there is a nonsingular $P \in \mathbb{R}^{n\times n}$ such that

$$P^T A P = D_a, \qquad P^T C P = D_c,$$

where both D_a and D_c are diagonal matrices. If we use P to change variables, then

$$x(\lambda) = P^{-T} \left(D_a + \lambda D_c \right)^{-1} P^T (b + \lambda d).$$

It is now clear that each component of $x(\cdot)$ is a rational function for $\lambda \in I_{\rm PD}$.

Theorem 5.2 *Let $q : \mathbb{R}^n \to \mathbb{R}$ and $c : \mathbb{R}^n \to \mathbb{R}$ be the quadratic functions defined on $\mathbb{R}^n$ by* (5.1)*, and assume that $I_{\rm PD}$ is not empty. If $x(\lambda) \in \mathbb{R}^n$ is the solution of* (5.2)*, and if the function $\phi : \mathbb{R} \to \mathbb{R}$ is defined on $I_{\rm PD}$ by*

$$\phi(\lambda) = c[x(\lambda)],$$

then ϕ is strictly decreasing on $I_{\rm PD}$ unless $x(\cdot)$ is constant on $I_{\rm PD}$ with

$$\nabla q[x(\lambda)] = 0, \qquad \nabla c[x(\lambda)] = 0,$$

for all λ in $I_{\rm PD}$.

Proof. We first need expressions for $\phi'(\lambda)$ and $x'(\lambda)$. The expression

$$\phi'(\lambda) = \nabla c[x(\lambda)]^T x'(\lambda)$$

for $\phi'(\lambda)$ follows from the definition of ϕ, while an expression for $x'(\lambda)$ is obtained by noting that since $x(\lambda)$ satisfies (5.2),

$$(A + \lambda C) x'(\lambda) = -\nabla c[x(\lambda)]. \tag{5.3}$$

Thus, using the expression for $\phi'(\lambda)$, we obtain that

$$\phi'(\lambda) = -x'(\lambda)^T (A + \lambda C) x'(\lambda).$$

Hence, $\phi'(\lambda) \le 0$ for $\lambda \in I_{\rm PD}$. If ϕ is not strictly decreasing on $I_{\rm PD}$, then $\phi'(\lambda) = 0$ for λ in some subinterval I if $I_{\rm PD}$. The above expression for $\phi'(\lambda)$ then shows that $x'(\lambda) = 0$ on I.

We now wish to conclude that $x'(\lambda) \equiv 0$ on $I_{\rm PD}$. The easiest way to prove this is to note that since each component of $x(\cdot)$ is a rational function for $\lambda \in I_{\rm PD}$, we can have $x'(\lambda) = 0$ on a subinterval I of $I_{\rm PD}$ only if $x'(\lambda) \equiv 0$ on $I_{\rm PD}$.

We have shown that $x(\cdot)$ is constant on $I_{\rm PD}$. Hence, (5.3) shows that $\nabla c[x(\lambda)] = 0$ on $I_{\rm PD}$, and thus (5.2) yields that $\nabla q[x(\lambda)] = 0$ on $I_{\rm PD}$. ■

The possibility in Theorem 5.2 that $x(\cdot)$ is constant on I_{PD} cannot be ruled out. For example, if A and C are two matrices such that I_{PD} is not empty, and we define b and d by

$$Ax_0 = -b, \qquad Cx_0 = -d$$

for some $x_0 \in \mathbb{R}^n$, then $x(\lambda) \equiv x_0$ for $\lambda \in I_{\mathrm{PD}}$. We can reverse this construction because if $x(\lambda) \equiv x_0$ for $\lambda \in I_{\mathrm{PD}}$, then x_0 satisfies $Ax_0 = -b$ and $Cx_0 = -d$. This argument shows that small perturbations on the data that defines q and c lead to a case where $x(\cdot)$ is not constant.

Theorem 5.2 can be used to find a global minimizer x^* of (3.1) in the case where $A + \lambda^* C$ is positive definite because in this case $\lambda^* \in I_{\mathrm{PD}}$ is a solution of $\phi(\lambda) = 0$. We now outline an algorithm for finding λ^* in this case.

In the search for λ^* we first need to decide whether I_{PD} is empty or determine $\lambda_0 \in I_{\mathrm{PD}}$. We start with bounds λ_l and λ_u such that $I_{\mathrm{PD}} \subset (\lambda_l, \lambda_u)$. For example, the bounds

$$\lambda_l = \max\left\{ -\frac{a_{i,i}}{c_{i,i}} : c_{i,i} > 0 \right\}, \qquad \lambda_u = \min\left\{ -\frac{a_{i,i}}{c_{i,i}} : c_{i,i} < 0 \right\},$$

could be used, where $\lambda_l = -\infty$ if $c_{i,i} \le 0$ for all i, and $\lambda_u = \infty$ if $c_{i,i} \ge 0$ for all i. We now show how to update (λ_l, λ_u) given $\lambda_0 \in (\lambda_l, \lambda_u)$.

If $\lambda_0 \notin I_{\mathrm{PD}}$, then $A + \lambda_0 C$ is not positive definite, and thus we can compute an $x \in \mathbb{R}^n$ such that

$$x^T(A + \lambda_0 C)x \le 0, \qquad x \ne 0. \tag{5.4}$$

A vector x that satisfies (5.4) can be computed, as in the algorithms of Gay [5] and Moré and Sorensen [15], by noting that during the Cholesky decomposition of $A + \lambda_0 C$ it is possible to find $\delta \ge 0$ such that the leading submatrix of order $l \le n$ of

$$A + \lambda_0 C + \delta e_l e_l^T$$

is singular, and $x \in \mathbb{R}^n$ such that $(A + \lambda_0 C + \delta e_l e_l^T)x = 0$ with $x_l = 1$ and $x_i = 0$ for $i > l$. Clearly, x satisfies (5.4).

Given a vector $x \in \mathbb{R}^n$ that satisfies (5.4), note that if $x^T C x < 0$, then

$$x^T(A + \lambda C)x \le x^T(A + \lambda_0 C)x \le 0$$

for $\lambda > \lambda_0$. Thus we can set $\lambda_u = \lambda_0$. Similarly, if $x^T C x > 0$, then $\lambda_l = \lambda_0$. Given a tolerance $\epsilon > 0$, we can use this updating procedure to reduce the length of (λ_l, λ_u) until we determine some $\lambda_0 \in I_{\mathrm{PD}}$, or we determine that the length of I_{PD} is less than ϵ.

Once we determine $\lambda_0 \in I_{\mathrm{PD}}$, we can continue to isolate λ^* by updating λ_l and λ_u such that $\lambda^* \in (\lambda_l, \lambda_u)$. This is easily done because if $\lambda_0 \in I_{\mathrm{PD}}$ and $\phi(\lambda_0) < 0$, then we set $\lambda_u = \lambda_0$, but if $\phi(\lambda_0) > 0$, then $\lambda_l = \lambda_0$. This procedure is valid even if $x(\cdot)$ is constant for $\lambda \in I_{\mathrm{PD}}$.

The algorithm that we have outlined can be used to generate a sequence $\{\lambda_k\}$ that converges to λ^*. Moreover, if $A + \lambda^* C$ is positive definite, then $\{x(\lambda_k)\}$ converges to $x(\lambda^*)$ with $\phi(\lambda^*) = 0$. Hence, $x(\lambda^*)$ is the global solution of (3.1).

We have not specified any particular algorithm to generate the sequence $\{\lambda_k\}$. If C is positive (negative) semidefinite, it is not difficult to show that ϕ is convex (concave) on I_{PD}, and thus a safeguarded version of Newton's method is a reasonable choice. This is the approach used by Gay [5] and Moré and Sorensen [15]. If C is indefinite, then ϕ may not be convex or concave, so extra care is needed. For example, if

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad C = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \qquad b = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad d = 0,$$

then $I_{\mathrm{PD}} = (-1, 1)$, but ϕ'' changes sign in $(-1, 1)$.

The situation where $A + \lambda^* C$ is singular is more delicate. In this case $\phi(\lambda) = 0$ may not have a solution in I_{PD}, and $\{x(\lambda_k)\}$ may not converge to the solution of (3.1); this is the reason why Moré and Sorensen [15], in their study of the problem with $C = I$, called this case the *hard case.*

Theorem 5.3 *Let A and C be symmetric matrices, and assume that I_{PD} is not empty. If an endpoint λ^* of I_{PD} is finite, then $A + \lambda^* C$ is positive semidefinite and singular. Moreover, if $I_{\mathrm{PD}} = (\lambda_l^*, \lambda_u^*)$, then*

$$[\lambda_l^*, \lambda_u^*] = I_{\mathrm{PSD}} \equiv \{\lambda \in \mathbb{R} : A + \lambda C \textit{ positive semidefinite } \}.$$

Proof. If λ^* is an endpoint of I, then it is clear that $A + \lambda^* C$ is positive semidefinite. If $A + \lambda^* C$ is not singular, then it must be positive definite. However, then $A + \lambda C$ is also positive definite in a neighborhood of λ^*. This conclusion is not possible because it contradicts the definition of I_{PD}.

We now show that if $I_{\mathrm{PD}} = (\lambda_l^*, \lambda_u^*)$, then $I_{\mathrm{PSD}} = [\lambda_l^*, \lambda_u^*]$. Clearly, $I_{\mathrm{PD}} \subset I_{\mathrm{PSD}}$. Moreover, just as in Theorem 5.1, we can prove that I_{PSD} is an interval. Thus, if $I_{\mathrm{PSD}} \neq [\lambda_l^*, \lambda_u^*]$, then $A + \lambda C$ must be positive semidefinite and singular on a nontrivial interval. This conclusion implies that the polynomial

$$p(\lambda) = \det(A + \lambda C)$$

vanishes in a nontrivial interval. Hence, $p(\lambda) \equiv 0$, and thus $A + \lambda C$ is singular for all λ, contradicting the assumption that I_{PD} is not empty. ■

Theorem 5.3 shows that if $A + \lambda^* C$ is singular, then λ^* must be an endpoint of I_{PD}. We now prove that λ^* is an endpoint of I_{PD} such that the limit

$$\lim_{\lambda \to \lambda^*} \phi(\lambda)$$

exists. We are certainly guaranteed the existence of this limit for one of the endpoints of $I_{\rm PD}$. In fact, since we are assuming that $\phi(\lambda) = 0$ does not have a solution in $I_{\rm PD}$, the function ϕ does not change sign in $I_{\rm PD}$. Hence, if $\phi(\lambda) > 0$ on $I_{\rm PD}$, then

$$\lim_{\lambda \to \lambda_u^*} \phi(\lambda) = \inf\{\phi(\lambda) : \lambda_l^* < \lambda < \lambda_u^*\},$$

because ϕ is decreasing in $I_{\rm PD}$. Similarly, if $\phi(\lambda) < 0$ on $I_{\rm PD}$, then

$$\lim_{\lambda \to \lambda_l^*} \phi(\lambda) = \sup\{\phi(\lambda) : \lambda_l^* < \lambda < \lambda_u^*\}.$$

These calculations suggest that $\lambda^* = \lambda_u^*$ if $\phi(\lambda) > 0$ on $I_{\rm PD}$, and that $\lambda^* = \lambda_l^*$ if $\phi(\lambda) < 0$ on $I_{\rm PD}$. We need the following result to establish this claim.

Theorem 5.4 *Let A and C be symmetric matrices, and assume that $I_{\rm PD}$ is not empty. If an endpoint λ^* of $I_{\rm PD}$ is finite and the limit*

$$\lim_{\lambda \to \lambda^*} \phi(\lambda) \tag{5.5}$$

exists, then the limit

$$\lim_{\lambda \to \lambda^*} x(\lambda)$$

also exists.

Proof. We have already noted that the function $x(\cdot)$ is a rational function for $\lambda \in I_{\rm PD}$ so we need to show that $\{x(\lambda_k)\}$ is bounded when $\{\lambda_k\}$ is any sequence in $I_{\rm PD}$ converging to λ^*; that is, we need to show that $x(\cdot)$ does not have a pole at λ^*. Assume, on the contrary, that $\{\|x(\lambda_k)\|\}$ is unbounded, and define

$$z_k = \frac{x(\lambda_k)}{\|x(\lambda_k)\|}.$$

Then $\|z_k\| = 1$, so that we can assume, without loss of generality, that $\{z_k\}$ converges to some vector z with $\|z\| = 1$.

Since $\{\phi(\lambda_k)\}$ is bounded and $\phi(\lambda_k) = c[x(\lambda_k)]$, we obtain that $z^T C z = 0$. Moreover, since

$$(A + \lambda_k C) z_k = -\frac{(b + \lambda_k d)}{\|x(\lambda_k)\|},$$

we also obtain that $(A + \lambda^* C) z = 0$. We have already shown that $z^T C z = 0$, so this yields that $z^T A z = 0$. However, this is not possible if $I_{\rm PD}$ is not empty. ■

We now prove that $\lambda^* = \lambda_u^*$ if $\phi(\lambda) > 0$ on $I_{\rm PD}$, and that $\lambda^* = \lambda_l^*$ if $\phi(\lambda) < 0$ on $I_{\rm PD}$. We have already shown that the limit (5.5) exists for this choice of λ^*, and thus Theorem 5.4 shows that

$$\lim_{\lambda \to \lambda^*} x(\lambda) = x^*$$

for some $x^* \in \mathbb{R}^n$, and hence

$$(A + \lambda^* C)x^* = -(b + \lambda^* d).$$

Although we also have that $A + \lambda^* C$ is positive semidefinite, x^* is not necessarily the solution of (3.1) because we cannot guarantee that $c(x^*) = 0$. However, we claim that if z_* is chosen so that

$$(A + \lambda^* C)z_* = 0, \qquad z_* \neq 0,$$

then $x^* + \alpha^* z_*$ solves (3.1) for some $\alpha^* \in \mathbb{R}$.

First of all, note that we can obtain z_* because $A + \lambda^* C$ is singular. Now assume, to be definite, that $\phi(\lambda) < 0$ on I_{PD}, and thus $\lambda^* = \lambda_l^*$ is the smallest endpoint of I_{PD}. Since

$$c(x^* + \alpha z_*) = c(x^*) + \alpha \nabla c(x^*)^T z_* + \tfrac{1}{2}\alpha^2 z_*^T C z_*,$$

it is clear that there is an $\alpha^* \in \mathbb{R}$ such that $c(x^* + \alpha^* z_*) = 0$ if we can show that $c(x^*) \leq 0$ and that $z_*^T C z_* > 0$. Since $\phi(\lambda) < 0$ on I_{PD}, and $\{x(\lambda)\}$ converges to x^*, it is clear that $c(x^*) \leq 0$. If, on the contrary, $z_*^T C z_* \leq 0$, then

$$z_*^T(A + \lambda C)z_* = z_*^T(A + \lambda^* C)z_* + (\lambda - \lambda^*) z_*^T C z_* \leq 0$$

for $\lambda > \lambda^* = \lambda_l^*$. This is not possible because $A + \lambda C$ is positive definite for $\lambda \in (\lambda_l^*, \lambda_u^*)$. Hence, $z_*^T C z_* > 0$ as desired.

6 Approximate Global Minimizers

The results in Section 5 show how to generate a sequence $\{\lambda_k\}$ in I_{PD} that converges to the Lagrange multiplier λ^* of the optimization problem (3.1). Moreover, if $A + \lambda^* C$ is positive definite then $\{x(\lambda_k)\}$ converges to the solution x^* of (3.1), while if $A + \lambda^* C$ is positive semidefinite and singular, then $\{x(\lambda_k)\}$ also converges to some x^*, but now a solution of (3.1) is obtained by computing $x^* + \alpha^* z_*$, where

$$(A + \lambda^* C)z_* = 0, \qquad z_* \neq 0,$$

and $\alpha^* \in \mathbb{R}$ satisfies $c[x^* + \alpha^* z_*] = 0$. In this section we address the question of how to determine approximate global minimizers $x \in \mathbb{R}^n$ such that, given tolerances $\epsilon_q > 0$ and $\epsilon_c > 0$,

$$q(x) \leq q^* + \epsilon_q, \qquad |c(x)| \leq \epsilon_c, \tag{6.1}$$

where

$$q^* \equiv \min\{q(x) : c(x) = 0\}.$$

We first analyze the case where $A + \lambda^* C$ is positive definite.

Theorem 6.1 *If* $\lambda \in I_{\rm PD}$*, then*

$$q[x(\lambda)] + \lambda c[x(\lambda)] \le q^*.$$

Proof. The proof is simple; just note that (5.2) implies that $x(\lambda)$ is the global minimizer of the quadratic $q + \lambda c$, and thus

$$q[x(\lambda)] + \lambda c[x(\lambda)] = \min\{q(x) + \lambda c(x) : x \in \mathbb{R}^n\} \le \min\{q(x) : c(x) = 0\} = q^*.$$

■

If $A + \lambda^* C$ is positive definite, then we can compute $\lambda \in I_{\rm PD}$ such that $\phi(\lambda) = c[x(\lambda)]$ is small. Theorem 6.1 shows that $x(\lambda)$ satisfies the termination criteria (6.1) if we determine $\lambda \in I_{\rm PD}$ such that

$$-\lambda c[x(\lambda)] \le \epsilon_q, \qquad |c[x(\lambda)]| \le \epsilon_c.$$

These inequalities are satisfied for all $\lambda \in I_{\rm PD}$ sufficiently close to λ^*.

Theorem 6.2 *If* $\lambda \in I_{\rm PD}$ *and*

$$c[x(\lambda) + \alpha_\lambda z_\lambda] = 0,$$

then

$$q^* \le q[x(\lambda) + \alpha_\lambda z_\lambda] \le q^* + \tfrac{1}{2}\alpha_\lambda^2 z_\lambda^T (A + \lambda C) z_\lambda.$$

Proof. The first inequality follows from the assumption that $c[x(\lambda)+\alpha_\lambda z_\lambda] = 0$ and from the definition of q^*. The proof of the second inequality is obtained in terms of the Lagrangian function

$$\mathcal{L}(x, \lambda) = q(x) + \lambda c(x).$$

Since the pair $(x(\lambda), \lambda)$ satisfies (5.2), $\nabla_x \mathcal{L}[x(\lambda), \lambda] = 0$, and thus

$$\mathcal{L}[x(\lambda) + \alpha_\lambda z_\lambda, \lambda] = \mathcal{L}[x(\lambda), \lambda] + \tfrac{1}{2}\alpha_\lambda^2 z_\lambda^T (A + \lambda C) z_\lambda.$$

We now make use of the assumption that $c[x(\lambda) + \alpha_\lambda z_\lambda] = 0$ to obtain

$$q[x(\lambda) + \alpha_\lambda z_\lambda] = q[x(\lambda)] + \lambda c[x(\lambda)] + \tfrac{1}{2}\alpha_\lambda^2 z_\lambda^T (A + \lambda C) z_\lambda.$$

The result now follows from Theorem 6.1. ■

Theorem 6.2 is applicable if λ is sufficiently close to λ^* and $A + \lambda^* C$ is positive semidefinite and singular. However, Theorem 6.2 is also applicable if we find $\lambda \in I_{\rm PD}$ and z_λ such that $\frac{1}{2} z_\lambda^T (A + \lambda C) z_\lambda$ is smaller than ϵ_q. Thus, if there is an $\alpha_\lambda \in \mathbb{R}$ such that $c[x(\lambda) + \alpha_\lambda z_\lambda] = 0$, then the termination criteria (6.1) are satisfied for any $\epsilon_c > 0$.

The theorems of this section extend results of Moré and Sorensen [15] to the case where the quadratic function c is arbitrary. Theorem 6.2 is an extension of Lemma 3.4 in that paper, while Theorem 6.1 is a generalization of Lemma 3.13.

7 Concluding Remarks

During the final stages of the preparation of this manuscript I attended the Panamerican Workshop on Applied and Computational Mathematics (Universidad Simon Bolivar, Caracas, January 10–15, 1993), where Henry Wolkowicz gave me a copy of his recently completed manuscript [18] with Ronald Stern on the problem

$$\min\{q(x) : c_l \le x^T C x \le c_u\}, \tag{7.1}$$

and on the application of this problem to nonsymmetric eigenvalue perturbations. Although the two manuscripts are related, there are several differences in our approach and results.

If we consider the one-sided case of (7.1), then it is clear that (7.1) is a special case of (1.3); however, note that if $\nabla^2 c = C$ is nonsingular, then a change of variables transforms (1.3) into (7.1).

The characterization result in [18] assumes that $\nabla c(x^*) \ne 0$, or that $c_l < 0 < c_u$. As pointed out in Section 3, this assumption is stronger than (3.2) for the equality constrained case where $c_l = c_u$. Note, in particular, that if C is indefinite, then (3.2) is automatically satisfied, but that this is not the case for the assumptions in [18]. Similar remarks apply to the inequality constrained case where $c_l = -\infty$.

The introduction of the two-sided bounds in (7.1) is an interesting variation. We can handle this variation, and the more general problem

$$\min\{q(x) : c_l \le c(x) \le c_u\},$$

with the material in Sections 5 and 6. The main difference is that instead of looking for a solution of $\phi(\lambda) = 0$, we would search for either a solution $\lambda^* \ge 0$ in I_{PD} of $\phi(\lambda^*) = c_u$, or a solution $\lambda^* \le 0$ in I_{PD} of $\phi(\lambda^*) = c_l$. Since ϕ is either constant or strictly decreasing, we can update the interval (λ_l, λ_u) of Section 5 and isolate the appropriate solution.

Our treatments of the algorithm also differ. Stern and Wolkowicz [18] follow the development of Moré and Sorensen [15] and outline an algorithm that exploits the Cholesky factorization of $A + \lambda C$, while our treatment in Sections 5 and 6 is independent of the method used to determine the solution $x(\lambda)$ of the linear system (5.2). We also note that the application of their algorithm to (1.3) requires a change of variables to reduce the problem to the standard form (7.1). In our approach this change of variables is not needed.

Acknowledgments. My enjoyable discussions with Henry Wolkowicz at the Panamerican Workshop on Applied and Computational Mathematics are gratefully acknowledged. Gail Pieper deserves special thanks for her careful reading of the manuscript.

References

[1] D. P. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*, Academic Press, 1982.

[2] T. F. Coleman and C. Hempel, *Computing a trust region step for a penalty function*, SIAM J. Sci. Statist. Comput., 11 (1990), pp. 180–201.

[3] F. Critchley, *On the minimization of a positive definite quadratic form under an arbitrary quadratic constraint*, Technical Report, University of Warwick, Coventry, England, 1990.

[4] W. Gander, *Least squares with a quadratic constraint*, Numer. Math., 36 (1981), pp. 291–307.

[5] D. M. Gay, *Computing optimal locally constrained steps*, SIAM J. Sci. Statist. Comput., 2 (1981), pp. 186–197.

[6] ——, *Subroutines for unconstrained minimization using a model/trust region approach*, ACM Trans. Math. Software, 9 (1983), pp. 503–524.

[7] G. Golub and U. von Matt, *Quadratically constrained least squares and quadratic problems*, Numer. Math., 59 (1991), pp. 561–580.

[8] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, 1989.

[9] M. Heinkenschloss, *On the solution of a two ball trust region subproblem*, Technical Report 92-16, Universität Trier, Trier, Germany, 1992.

[10] M. R. Hestenes, *Optimization Theory*, John Wiley & Sons, 1975.

[11] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 1985.

[12] N. Karmarkar, M. G. C. Resende, and K. G. Ramakrishnan, *An interior point algorithm to solve computationally difficult set covering problems*, Math. Programming, 52 (1991), pp. 597–618.

[13] M. G. Krein and J. L. Smuljan, *Plus-operators in a space with an indefinite metric*, Amer. Math. Soc. Transl., 85 (1969), pp. 93–113.

[14] J. M. Martínez and S. A. Santos, *A trust region method for minimization on arbitrary domains*, Technical Report, State University of Campinas, Campinas, Brazil, 1991.

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

[16] P. M. Pardalos, Y. Ye, and C.-G. Han, *Algorithms for the solution of quadratic knapsack-problems*, Linear Algebra Appl., 152 (1991), pp. 69–91.

[17] D. C. Sorensen, *Newton's method with a model trust region modification*, SIAM J. Numer. Anal., 19 (1982), pp. 409–426.

[18] R. J. Stern and H. Wolkowicz, *Indefinite trust region subproblems and nonsymmetric eigenvalue perturbations*, Technical Report SOR 93-1, Princeton University, Princeton, New Jersey, 1993.

[19] F. Uhlig, *A recurring theorem about pairs of quadratic forms and extensions: A survey*, Linear Algebra Appl., 25 (1979), pp. 219–237.

[20] Y. Ye, *On affine scaling algorithms for nonconvex quadratic programming*, Math. Programming, 56 (1992), pp. 285–300.

[21] Y. Yuan, *On a subproblem of trust region algorithms for constrained optimization*, Math. Programming, 47 (1990), pp. 53–63.

[22] ——, *A dual algorithm for minimizing a quadratic function with two quadratic constraints*, J. Comput. Math., 9 (1991), pp. 348–359.

[23] Y. Zhang, *Computing a Celis-Dennis-Tapia trust-region step for equality constrained optimization*, Math. Programming, 55 (1992), pp. 109–124.