

Computing a Celis-Dennis-Tapia
Trust Region Step for
Equality Constrained Optimization

Yin Zhang

December, 1988
(revised June, 1989)

TR88-16

Computing a Celis-Dennis-Tapia Trust-region Step for Equality Constrained Optimization *

Yin Zhang[†]

December, 1988

Revised June, 1989

Abstract

We study an approach for minimizing a convex quadratic function subject to two quadratic constraints. This problem stems from computing a trust-region step for an SQP algorithm proposed by Celis, Dennis and Tapia (1984) for equality constrained optimization. Our approach is to reformulate the problem into a univariate nonlinear equation $\phi(\mu) = 0$ where the function $\phi(\mu)$ is continuous, at least piecewise differentiable and monotone. Well-established methods then can be readily applied. We also consider an extension of our approach to a class of non-convex quadratic functions and show that our approach is applicable to reduced Hessian SQP algorithms. Numerical results are presented indicating that our algorithm is reliable, robust and has the potential to be used as a building block to construct trust-region algorithms for small-sized problems in constrained optimization.

Keywords: Constrained optimization, trust-region algorithms, CDT problem, univariate nonlinear equation, Newton's method.

Abbreviated Title: Computing a CDT trust-region step.

*Research supported in part by AFOSR 85-0243, and DOE DEFG05-86ER 25017.

[†]Department of Mathematical Sciences, Rice University, Houston, Texas 77251-1892

1 Introduction

In this paper, we consider solving the following minimization problem

$$\text{minimize } q(d) \equiv g^T d + \frac{1}{2} d^T B d, \quad d \in \mathbf{R}^n \quad (1.1)$$

$$\text{subject to } \|d\| \leq \delta \quad (1.2)$$

$$\text{and } \|A^T d + h\| \leq \theta, \quad (1.3)$$

where $B \in \mathbf{R}^{n \times n}$ is symmetric, $A \in \mathbf{R}^{n \times m}$ ($m \leq n$), $g \in \mathbf{R}^n$, $h \in \mathbf{R}^m$, $\delta > 0$, $\theta > 0$, and throughout this paper the norm $\|\cdot\|$ denotes the ℓ_2 norm. In order to have a meaningful feasible set, we assume that

$$\theta > \theta_{\min} \equiv \min\{\|A^T d + h\| : \|d\| \leq \delta\}. \quad (1.4)$$

Since only a global solution is of interest to us, the term “solution” always implies a global solution.

The above problem comes from applying the successive quadratic programming (SQP) method and a trust-region technique to minimize a function $f(x)$ subject to the equality constraints $h(x) = 0$. At the k -th iteration, we want to obtain the correction step d_k to the current iterate x_k by minimizing a quadratic model $q(d) = g^T d + d^T B d/2$ subject to the linearized constraints $A^T d + h = 0$, where $g = \nabla f(x_k)$, B is the Hessian or an approximate Hessian of the Lagrangian function with respect to x , $A = \nabla h(x_k)$ and $h = h(x_k)$. Meanwhile, we also want to impose the trust-region restriction $\|d\| \leq \delta$. The linearized constraints and the trust-region restriction are not necessarily compatible when $h \neq 0$ (we assume $h \neq 0$ in this paper), so in order to guarantee a non-empty feasible set, we replace the requirement that the linearized constraints be zero by a condition that the norm of the linearized constraints be within a given tolerance level. Eventually, we end up with Problem (1.1)-(1.3). This approach was first proposed by Celis, Dennis and Tapia [2] and later it was also used by Powell and Yuan [11] in their algorithm. For brevity, we shall call Problem (1.1)-(1.3) the CDT problem in this paper.

From the first-order necessary conditions for the CDT problem (the constraint qualification is satisfied as is pointed out by Yuan [14] because the feasible set is convex with

nonempty interior), one can easily deduce that a solution d^* to the CDT problem satisfies

$$(B + \mu^* AA^T + \lambda^* I)d^* = -(g + \mu^* Ah), \quad (1.5)$$

where $\mu^* \geq 0$ and $\lambda^* \geq 0$ are the multipliers of Problem (1.1)-(1.3). The vector d^* is of course feasible. In addition, one has the two complementarity conditions

$$\lambda^* = 0 \text{ and } \|d^*\| < \delta \quad \text{or} \quad \lambda^* \geq 0 \text{ and } \|d^*\| = \delta, \quad (1.6)$$

$$\mu^* = 0 \text{ and } \|A^T d^* + h\| < \theta \quad \text{or} \quad \mu^* \geq 0 \text{ and } \|A^T d^* + h\| = \theta. \quad (1.7)$$

Without the constraint (1.3), Problem (1.1)-(1.2) is the trust-region subproblem for unconstrained optimization and its solution is well understood. The following theorem characterizes global solutions of Problem (1.1)-(1.2). It was proved independently by Gay [4] and Sorensen [13].

Theorem 1.1 (Gay, Sorensen) *A vector d^* is a global solution to Problem (1.1)-(1.2) if and only if $\|d^*\| \leq \delta$ and for some $\lambda^* \geq 0$,*

$$(B + \lambda^* I)d^* = -g, \quad \lambda^*(\|d^*\| - \delta) = 0, \quad (1.8)$$

with $B + \lambda^ I$ positive semi-definite. This λ^* is unique. Moreover, if $B + \lambda^* I$ is positive definite, then d^* is the unique global solution.*

Based on this strong necessary and sufficient optimality condition, very effective Newton type algorithms have been constructed for obtaining a solution or an approximate solution of Problem (1.1)-(1.2) (see [6] for further references).

Unfortunately, with the presence of constraint (1.3) a similar necessary and sufficient optimality condition no longer exists as was recently shown by Yuan (1987) [14]. The positive semi-definiteness of the matrix $(B + \mu^* AA^T + \lambda^* I)$ cannot be guaranteed for general symmetric matrices B . Yuan gave examples demonstrating that this matrix can have one negative eigenvalue when μ^* and λ^* are unique and can even have two negative eigenvalues in an unfavorable situation. The lack of a necessary and sufficient optimality condition and a positive semi-definite Hessian of the Lagrangian (with respect to d) makes it much more

difficult to construct effective algorithms for solving this problem. Nevertheless, Yuan gives sufficient conditions for a solution under the assumption that the Hessian of the Lagrangian with respect to d is positive semi-definite at d^* . The following is Yuan's Theorem 2.5 in [14], with a uniqueness result added by us.

Theorem 1.2 *The conditions (1.5), (1.6), (1.7) and the positive semi-definiteness of the matrix $B + \mu^*AA^T + \lambda^*I$, where μ^* and λ^* are non-negative, are sufficient for a feasible d^* to be a solution of the CDT problem. Moreover, if the matrix $B + \mu^*AA^T + \lambda^*I$ is positive definite, then the solution d^* is unique.*

Proof: We only prove the last statement which was not a part of Yuan's original theorem. Suppose that $B + \mu^*AA^T + \lambda^*I$ is positive definite but d^* is not unique. Then there exists another global solution \hat{d} such that $q(\hat{d}) = q(d^*)$. Since $B + \mu^*AA^T + \lambda^*I$ is positive definite, d^* is the unique minimizer of the quadratic

$$q(d) + \frac{\mu^*}{2}\|A^T d + h\|^2 + \frac{\lambda^*}{2}\|d\|^2.$$

Therefore,

$$\mu^*\|A^T d^* + h\|^2 + \lambda^*\|d^*\|^2 < \mu^*\|A^T \hat{d} + h\|^2 + \lambda^*\|\hat{d}\|^2. \quad (1.9)$$

Since both d^* and \hat{d} are global solutions, from the complementarity conditions (1.6) and (1.7), we have

$$\mu^*\|A^T d^* + h\|^2 = \mu^*\|A^T \hat{d} + h\|^2 = \mu^*\theta^2 \quad \text{and} \quad \lambda^*\|d^*\|^2 = \lambda^*\|\hat{d}\|^2 = \lambda^*\delta^2,$$

which contradicts (1.9). Hence d^* must be unique. \square

While the problem of effectively solving the ℓ_2 -norm CDT problem for a general symmetric matrix B is still open, some algorithms have been recently proposed for other norms or for the ℓ_2 -norm but in less general cases. For example, Celis *et al* [1] suggest solving a modified CDT problem by restricting d to a two-dimensional subspace. Yuan [15] proposes a dual algorithm for solving the CDT problem for a positive definite matrix B . His approach is to solve the dual problem: maximize the Lagrangian function of the CDT problem subject to the constraints that (i) its gradient with respect to d vanishes and (ii) the multipliers are

non-negative. Eliminating the first constraint by substitution, he reduces the problem to maximizing a quadratic concave function of two variables (multipliers) in the first quadrant. His algorithm basically uses Newton's method with a line search but a projected steepest descent direction is used whenever Newton's direction is not feasible. Moreover, special care has to be taken near the boundary to ensure that the iterates stay in the first quadrant. To evaluate the objective function and its Hessian, the Cholesky factors of the matrix $B + \mu AA^T + \lambda I$ have to be computed and stored.

In this paper, we propose a new approach for solving the CDT problem. Our approach is to reformulate the problem into a univariate nonlinear equation $\phi(\mu) = 0$ where the function $\phi(\mu)$ is continuous, at least piecewise differentiable and monotonically decreasing. Well-established methods for root finding then can be readily applied.

This paper is organized as follows. In Section 2, we show how the CDT problem with a convex $q(d)$ can be reduced to a well-behaved univariate nonlinear equation $\phi(\mu) = 0$. In Section 3, we show that our approach can be extended to the CDT problem with a class of non-convex $q(d)$ and give an important application. Numerical results are presented in Section 4. We give some concluding remarks in the last section.

From now on, we shall assume that B is positive definite unless otherwise specified. We have to point out that this is an unsatisfactory assumption since one of the advantages of trust-region strategy is presumably its ability to handle indefinite Hessian (or approximate Hessian) matrices. Nevertheless, when a BFGS type quasi-Newton update is used along with measures to ensure positive definiteness, this assumption on B is still reasonable.

2 Problem reformulation

We first define two functions of the variables $\mu \geq 0$ and $\lambda \geq 0$, namely,

$$H(\mu, \lambda) = B + \mu AA^T + \lambda I \in \mathbf{R}^{n \times n}; \quad (2.1)$$

$$d(\mu, \lambda) = -H(\mu, \lambda)^{-1}(g + \mu Ah) \in \mathbf{R}^n. \quad (2.2)$$

Then, we define three functions of the variable $d \in \mathbf{R}^n$, namely,

$$y(d) = A(A^T d + h) \in \mathbf{R}^n; \quad (2.3)$$

$$\psi(d) = \frac{1}{2}(\|d\|^2 - \delta^2) \in \mathbf{R}; \quad (2.4)$$

$$\phi(d) = \frac{1}{2}(\|A^T d + h\|^2 - \theta^2) \in \mathbf{R}. \quad (2.5)$$

It is evident that the constraints (1.2) and (1.3) are equivalent to $\psi(d) \leq 0$ and $\phi(d) \leq 0$, respectively.

For simplicity, we shall adopt the following notational conventions. Whenever the vector $d(\mu, \lambda)$ defined by (2.2) is substituted into the above three functions of d , we shall write them as functions of (μ, λ) without changing their names. For example,

$$y(\mu, \lambda) \equiv y(d(\mu, \lambda)) = A(A^T d(\mu, \lambda) + h).$$

Furthermore, if λ is given as a function of μ , we shall write all the above five functions as functions of a single variable μ without changing their names. For instance,

$$\phi(\mu) \equiv \phi(d(\mu, \lambda(\mu))). \quad (2.6)$$

It is worth noting that $d = \nabla_d \psi(d)$ and $y(d) = \nabla_d \phi(d)$.

Our approach for solving the CDT problem is motivated by the following observation. If the constraint (1.3) of the CDT problem is active (we will show that this is the case of interest) and if we can define λ as a function of μ such that $\lambda(\mu)$ always satisfies the feasibility condition

$$\|d(\mu, \lambda(\mu))\| \leq \delta \quad (2.7)$$

and the first complementarity condition

$$\lambda(\mu) = 0 \quad \text{or} \quad \psi(\mu, \lambda(\mu)) = 0, \quad (2.8)$$

(we will see that there is a natural way to define such a $\lambda(\mu)$), then the solution of the CDT problem reduces to solving the second complementarity condition (1.7), or equivalently solving the univariate nonlinear equation

$$\phi(\mu) = 0. \quad (2.9)$$

To construct the desired function $\lambda(\mu)$, we notice that by Theorem 1.1, the conditions (2.2), (2.7) and (2.8) are the necessary and sufficient conditions for $d(\mu, \lambda)$ and λ to be the unique solution and the unique Lagrange multiplier, respectively, of the following problem:

$$\begin{aligned} & \text{minimize} && q(d) + \mu\phi(d), \quad d \in \mathbf{R}^n \\ & \text{subject to} && \|d\| \leq \delta. \end{aligned} \tag{2.10}$$

This problem is in the standard form of a trust-region subproblem for unconstrained optimization with its solution and multiplier depending on the parameter μ . We will refer to Problem (2.10) as $\mathcal{P}(\mu)$ to emphasize its dependence on the parameter μ . Obviously, $\mathcal{P}(0)$ is Problem (1.1)-(1.2) and it is natural to use $\mathcal{P}(\infty)$ to denote the problem:

$$\begin{aligned} & \text{minimize} && \phi(d), \quad d \in \mathbf{R}^n \\ & \text{subject to} && \|d\| \leq \delta. \end{aligned} \tag{2.11}$$

The following lemma is a direct consequence of Theorem 1.1 and the positive definiteness assumption on B .

Lemma 2.1 *Let $\lambda(\mu)$ be the multiplier corresponding to the constraint in $\mathcal{P}(\mu)$, then $\lambda(\mu)$ is a well-defined non-negative function of μ for $\mu \geq 0$. Specifically*

$$\lambda(\mu) = \begin{cases} 0, & \psi(\mu, 0) \leq 0, \\ \Lambda(\mu), & \text{otherwise,} \end{cases} \tag{2.12}$$

where $\Lambda(\mu) \geq 0$ is the implicit function defined by the equation $\psi(\mu, \Lambda) = 0$. Conditions (2.7) and (2.8) are satisfied by $\lambda(\mu)$. Moreover, the solution of $\mathcal{P}(\mu)$ is

$$d(\mu) = d(\mu, \lambda(\mu))$$

which is also a well-defined function of $\mu \geq 0$.

Once the function $\lambda(\mu)$ is defined as in Lemma 2.1, we have the following sufficient optimality condition.

Lemma 2.2 *Let $\lambda(\mu)$ and $d(\mu)$ be as in Lemma 2.1, and $\phi(\mu)$ be defined by (2.6). If $\mu^* \geq 0$ is such that $\phi(\mu^*) = 0$, then $d(\mu^*)$ is the unique solution of the CDT problem.*

Proof: This is straightforward from Theorem 1.2 since $d(\mu^*)$ is obviously feasible, conditions (1.5), (1.6) and (1.7) are all satisfied and $B + \mu^*AA^T + \lambda(\mu^*)I$ is positive definite. \square

We next show that the functions $\lambda(\mu)$, $d(\mu)$ and $\phi(\mu)$ are well-behaved.

Theorem 2.1 *Let $\lambda(\mu)$ and $d(\mu)$ be as in Lemma 2.1, and $\phi(\mu)$ be defined by (2.6).*

1. *The functions $\lambda(\mu)$, $d(\mu)$ and $\phi(\mu)$ are all continuous in $[0, +\infty)$.*
2. *The functions $\lambda(\mu)$, $d(\mu)$ and $\phi(\mu)$ are all differentiable in $[0, +\infty)$ except possibly at zeros of $\psi(\mu, 0)$ which are all isolated.*
3. *The derivatives of $\lambda(\mu)$, $d(\mu)$ and $\phi(\mu)$ are:*

$$\lambda'(\mu) = \begin{cases} 0, & \psi(\mu, 0) < 0, \\ -d(\mu)^T H(\mu)^{-1} y(\mu) / d(\mu)^T H(\mu)^{-1} d(\mu), & \psi(\mu, 0) > 0, \end{cases} \quad (2.13)$$

$$d'(\mu) = -H(\mu)^{-1} y(\mu) - \lambda'(\mu) H(\mu)^{-1} d(\mu), \quad \psi(\mu, 0) \neq 0, \quad (2.14)$$

$$\phi'(\mu) = -y(\mu)^T H(\mu)^{-1} y(\mu) - \lambda'(\mu) y(\mu)^T H(\mu)^{-1} d(\mu), \quad \psi(\mu, 0) \neq 0. \quad (2.15)$$

4. *The function $\phi(\mu)$ is monotonically decreasing in $[0, +\infty)$.*

Proof: From the definition of $d(\lambda)$ and $\phi(\mu)$, we see that they are continuous and differentiable if $\lambda(\mu)$ is continuous and differentiable. Hence it suffices to prove the continuity and differentiability of $\lambda(\mu)$. Now let us assume $\mu_0 \in [0, \infty)$.

First, we note that the function $\psi(\mu, 0)$ is continuous. If $\psi(\mu_0, 0) < 0$, then there exists a neighborhood of μ_0 in $[0, \infty)$ such that $\psi(\mu, 0) < 0$ for μ in that neighborhood. By definition (2.12), $\lambda(\mu) = 0$ for μ in that neighborhood. Thus, $\lambda(\mu)$ is continuous and differentiable at μ_0 , and $\lambda'(\mu_0) = 0$ which gives the expression for $\lambda'(\mu)$ in (2.13) for the case $\psi(\mu, 0) < 0$.

Similarly, if $\psi(\mu_0) > 0$, then there exists a neighborhood of μ_0 in $[0, \infty)$ such that $\psi(\mu, 0) > 0$ for μ in that neighborhood. By (2.12), $\lambda(\mu) = \Lambda(\mu)$ for μ in that neighborhood. In view of Theorem 1.1, there exists a unique Λ_0 such that $\psi(\mu_0, \Lambda_0) = 0$. Through direct calculation, we have

$$\frac{\partial \psi(\mu, \lambda)}{\partial \mu} = -d(\mu, \lambda)^T H(\mu, \lambda)^{-1} y(\mu, \lambda),$$

and

$$\frac{\partial \psi(\mu, \lambda)}{\partial \lambda} = -d(\mu, \lambda)^T H(\mu, \lambda)^{-1} d(\mu, \lambda).$$

The latter is negative at (μ_0, Λ_0) since $H(\mu_0, \Lambda_0)$ is positive definite and

$$\|d(\mu_0, \Lambda_0)\| = \delta > 0.$$

Therefore $\Lambda(\mu)$ is well-defined and differentiable in a neighborhood of μ_0 by the well-known implicit function theorem (see [9], for example). In addition,

$$\Lambda'(\mu) = - \left[\frac{\partial \psi(\mu, \lambda)}{\partial \lambda} \right]^{-1} \frac{\partial \psi(\mu, \lambda)}{\partial \mu}$$

which gives the expression for $\lambda'(\mu)$ in (2.13) in the case $\psi(\mu, 0) > 0$.

Let μ_0 be such that $\psi(\mu_0, 0) = 0$ and $\{\mu_j\}_{j=1}^{\infty}$ be any sequence that converges to μ_0 . To prove the continuity of $\lambda(\mu)$ at μ_0 , we need to show that

$$\lim_{j \rightarrow \infty} \lambda(\mu_j) = \lambda(\mu_0) = 0.$$

Without loss of generality, we assume that $\psi(\mu_j, 0) > 0$ for all $j \geq 1$ (otherwise $\lambda(\mu_j) = 0$) which implies

$$\|d(\mu_j, \lambda(\mu_j))\| = \|d(\mu_0, 0)\| = \delta$$

for all j . Let j go to infinity and let $\lambda_{\max} = \limsup_{j \rightarrow \infty} \lambda(\mu_j)$. We have

$$\|d(\mu_0, \lambda_{\max})\| = \|d(\mu_0, 0)\|,$$

or equivalently,

$$(g + \mu_0 Ah)^T [H(\mu_0, 0)^{-2} - H(\mu_0, \lambda_{\max})^{-2}] (g + \mu_0 Ah) = 0. \quad (2.16)$$

There exists an orthogonal matrix Q such that

$$H(\mu_0, 0) = Q^T [\text{diag}(\sigma_i)_{i=1}^n] Q$$

where $\sigma_i > 0$, $i = 1, 2, \dots, n$. Consequently

$$H(\mu_0, \lambda_{\max}) = H(\mu_0, 0) + \lambda_{\max} I = Q^T [\text{diag}(\sigma_i)_{i=1}^n + \lambda_{\max} I] Q.$$

It follows from (2.16) that

$$\sum_{i=1}^n \left(\frac{1}{\sigma_i^2} - \frac{1}{(\sigma_i + \lambda_{\max})^2} \right) p_i^2 = 0$$

where p_i is the i -th element of $Q(g + \mu_0 Ah)$. Since the p_i 's cannot be all zero otherwise $d(\mu_0, 0) = 0$ which contradicts $\psi(\mu_0, 0) = 0$, so we must have at least one index i ($1 \leq i \leq n$) such that

$$\frac{1}{\sigma_i} = \frac{1}{\sigma_i + \lambda_{\max}}.$$

This implies $\lambda_{\max} = 0$ and leads to

$$\lim_{j \rightarrow \infty} \lambda(\mu_j) = \lambda(\mu_0) = 0,$$

recalling that $\lambda_{\max} = \limsup_{j \rightarrow \infty} \lambda(\mu_j)$ and $\lambda(\mu) \geq 0$. Therefore, $\lambda(\mu)$ is continuous at μ_0 .

So far we have proved that the function $\lambda(\mu)$ is continuous in $[0, \infty)$ and also differentiable except possibly at points where $\psi(\mu, 0) = 0$. Since $\psi(\mu, 0)$ is a non-constant rational function which is real analytic in $[0, \infty)$, all its zeros have to be isolated by the well-known theory of analytic functions. Hence $\lambda(\mu)$, and in turn $d(\mu) = d(\mu, \lambda(\mu))$ and $\phi(\mu) = \phi(\mu, \lambda(\mu))$ which are differentiable with respect to both $\mu \geq 0$ and $\lambda \geq 0$, are continuous and at least piecewise differentiable.

Now we calculate the derivatives for $d(\mu)$ and $\phi(\mu)$ when $\psi(\mu, 0) \neq 0$. Differentiating both sides of the equation

$$(B + \mu AA^T + \lambda(\mu)I)d(\mu) = -(g + \mu Ah)$$

and rearranging the terms, we have

$$H(\mu)d'(\mu) = -y(\mu) - \lambda'(\mu)d(\mu)$$

which leads to the expression (2.14) for $d'(\mu)$. For $\phi(\mu)$, we have

$$\phi'(\mu) = \nabla_d \phi(d(\mu))^T d'(\mu) = y(\mu)^T d'(\mu)$$

which gives (2.15).

Finally, we need to show that $\phi(\mu)$ is monotonically decreasing in $[0, \infty)$. It suffices to show that $\phi'(\mu) \leq 0$ whenever the derivative exists because $\phi'(\mu)$ is at least piecewise differentiable. Substituting $\lambda'(\mu)$ into (2.15), we have

$$\phi'(\mu) = -y(\mu)^T H(\mu)^{-1} y(\mu) \leq 0$$

when $\psi(\mu, 0) < 0$. For $\psi(\mu, 0) > 0$,

$$\phi'(\mu) = -y(\mu)^T H(\mu)^{-1} y(\mu) \left(1 - \frac{(d(\mu)^T H(\mu)^{-1} y(\mu))^2}{d(\mu)^T H(\mu)^{-1} d(\mu) y(\mu)^T H(\mu)^{-1} y(\mu)} \right) \leq 0 \quad (2.17)$$

by the Cauchy-Schwarz inequality. This completes the proof. \square

Now we are ready to reformulate the CDT problem into a well-behaved univariate non-linear equation.

Theorem 2.2 *Let the function $\phi(\mu)$ be given as in Theorem 2.1. If $\phi(0) \leq 0$, then $d(0)$ is the solution of the CDT problem; otherwise, there exists a $\mu^* \geq 0$ such that $\phi(\mu^*) = 0$ (which implies that the constraint (1.3) is active). The vector $d(\mu^*)$ is the unique solution of the CDT problem.*

Proof: Since $d(0)$ is already the solution of Problem (1.1)-(1.2), if $\phi(0) \leq 0$, then the constraint (1.3) is also satisfied which implies that $d(0)$ is the solution of the CDT problem.

Now let us assume $\phi(0) > 0$ and let $d(\infty)$ denote a solution of $\mathcal{P}(\infty)$ (i.e., Problem 2.11). In view of the definitions of θ_{\min} and $\phi(\mu)$ (see (1.4) and (2.6)), we have

$$\phi(d(\infty)) = (\theta_{\min}^2 - \theta^2)/2 < 0.$$

We first show $\phi(\mu) < 0$ for μ large enough. Suppose that this not true, i.e., $\phi(\mu) \geq 0$ for all μ . Since $d(\mu)$ is the solution of $\mathcal{P}(\mu)$, we have

$$q(d(\mu)) + \mu\phi(\mu) \leq q(d(\infty)) + \mu\phi(d(\infty)).$$

It then follows that

$$q(d(\mu)) - q(d(\infty)) \leq \mu\phi(d(\infty)) = \mu(\theta_{\min}^2 - \theta^2)/2.$$

The right-hand side tends to $-\infty$ as μ goes to $+\infty$. The left-hand side, however, is bounded below by $q(d(0)) - q(d(\infty))$. This is a contradiction. Therefore, $\phi(\mu) < 0$ for μ large enough. In fact, $\phi(\mu) < 0$ for $\mu \geq \bar{\mu}$ where

$$\bar{\mu} = \frac{2[q(d(\infty)) - q(d(0))]}{\theta^2 - \theta_{\min}^2}. \quad (2.18)$$

Since $\phi(0) > 0$, $\phi(\mu) < 0$ for some $\mu > 0$, and $\phi(\mu)$ is continuous and monotonically decreasing by Theorem 2.1, there must exist a $\mu^* > 0$ such that $\phi(\mu^*) = 0$. Finally, $d(\mu^*)$ is the solution of the CDT problem by Lemma 2.2. \square

We note that because of its monotonically decreasing property, $\phi(\mu)$ cannot have more than one isolated zero. But it could have more than one zero. In that situation, there must exist a unique interval in which $\phi(\mu)$ is identically zero. This happens only if both constraints are active and d^* and $y(d^*)$ are linearly dependent as shown by Yuan [14]. From (2.17), we see that $\phi'(\mu) = 0$ whenever $d(\mu)$ and $y(\mu)$ are linearly dependent.

It is of interest to interpret our formulation geometrically. Clearly, $d(\mu)$ is a continuous curve in \mathbf{R}^n . The solution $d(\mu^*)$ is the point on the curve where the curve intersects the surface of $\phi(d) = 0$ which is in general an elliptical cylinder because the rank of A is generally less than n for equality constrained optimization problems.

3 Extension to non-convex quadratics

Since the second-order sufficient condition for equality constrained optimization only requires the Hessian of Lagrangian with respect to x to be positive definite on the null space of constraint gradients, the positive definiteness requirement for B is often (and rightly) viewed as too strong. It is therefore desirable to extend our formulation to the CDT problem with non-convex $q(d)$, namely, B may be indefinite. The following theorem provides us such an extension.

Theorem 3.1 *Let the following two conditions hold:*

1. *The matrix $B \in \mathbf{R}^{n \times n}$ in the CDT problem is positive definite on the null space of A^T , i.e., $p^T B p > 0$ for all non-zero $p \in \mathbf{R}^n$ such that $A^T p = 0$.*

2. There exists a $\hat{\mu} \in [0, \mu^*)$, where μ^* is the multiplier μ^* associated with the constraint (1.3), such that $B + \hat{\mu}AA^T$ is positive semi-definite.

Then the functions $\lambda(\mu)$, $d(\mu)$ and $\phi(\mu)$ can be defined for $\mu > \hat{\mu}$ as in Theorem 2.1 and they are all continuous and at least piecewise differentiable in $(\hat{\mu}, \infty)$. Moreover, $\phi(\mu)$ is monotonically decreasing in $(\hat{\mu}, \infty)$. If $\phi(\hat{\mu}) > 0$, then there exist a $\mu^* > \hat{\mu}$ such that $\phi(\mu^*) = 0$ and the solution of the CDT problem is $d(\mu^*)$.

Proof: We omit the proof because it is analogous to the proof of Theorem 2.2 noting that $H(\mu, \lambda)$ is positive definite for all μ in $(\hat{\mu}, \infty)$ and all λ in $[0, \infty)$. \square

Although the second condition in the theorem seems to be difficult to verify, this theorem does have an important application. It is well-known that in SQP type algorithms for constrained optimization, the reduced Hessian on the null space of the active constraint gradients is *the* essential piece of the second-order derivative information for fast local convergence (see [10] for an explanation). It is a very popular approach to just use the reduced Hessian to generate SQP steps (see [3] and [8], to cite a few examples). The advantages are: (i) a smaller matrix is stored and handled and (ii) near a solution the reduced Hessian is usually positive definite. Let us assume now that a reduced Hessian approach is used in the CDT problem where A is of full rank, then we have $B = ZMZ^T$, where $M \in \mathbf{R}^{m \times m}$ is an approximation to the reduced Hessian and is positive definite and $Z \in \mathbf{R}^{n \times (n-m)}$ is a basis for the null space of A^T (assuming that A is of full rank). It is easy to verify that such a matrix B satisfies the two conditions in Theorem 3.1 with $\hat{\mu} = 0$. So if $d(0)$ is not a solution, then the CDT problem can be reduced to the zero finding problem $\phi(\mu) = 0$.

Unfortunately, it is not possible to directly extend our approach to general symmetric matrices B . This is because in our approach, $H(\mu, \lambda)$ is always kept positive definite, if $H(\mu^*, \lambda^*)$ is in fact indefinite, as it may be, and if d^* and $y(d^*)$ are linearly independent which implies that (μ^*, λ^*) is unique, then there is no way that the necessary condition (1.5) can be satisfied by a positive definite matrix $H(\mu, \lambda)$.

4 Numerical Tests

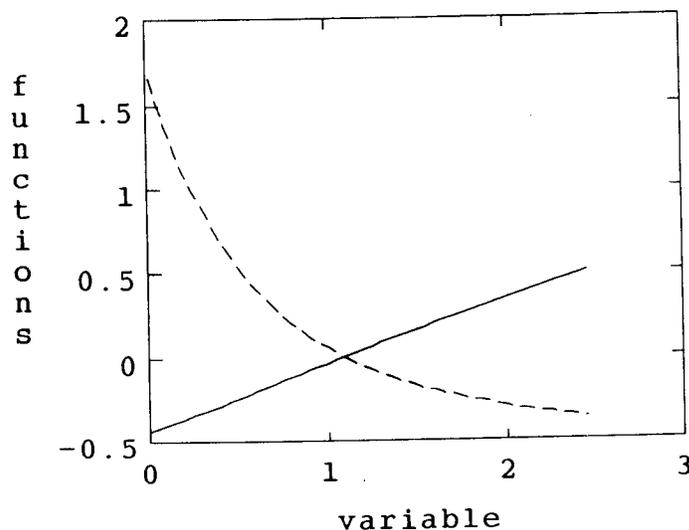
We have reformulated the CDT problem with a convex quadratic objective function $q(d)$ into a problem of zero-finding for an at least piecewise differentiable and monotonically decreasing function $\phi(\mu)$. The zero-finding problem for univariate functions is perhaps among the oldest problems considered in numerical analysis. There are many well-established methods for solving this problem with guaranteed convergence.

We have found that instead of solving $\phi(\mu) = 0$, it is generally easier to solve the equivalent equation $\Phi(\mu) = 0$, where

$$\Phi(\mu) = \frac{1}{\|A^T d(\mu) + h\|} - \frac{1}{\theta}. \quad (4.1)$$

The function $\Phi(\mu)$ is still a monotone (this time increasing) function with the same smoothness properties as $\phi(\mu)$. It is easy to verify that $\Phi'(\mu) = -\phi'(\mu)/\|A^T d(\mu) + h\|^3$. The nice thing about $\Phi(\mu)$ is that it usually exhibits a lower degree of nonlinearity than $\phi(\mu)$ does. Figure 1 shows the behavior of the two functions for a random problem. It has been observed from numerous examples that this kind of phenomenon is not atypical.

Figure 1: Functions Φ (solid line) and ϕ (dashed line) vs. Variable μ



4.1 An Implementation of Newton's method

We have implemented an algorithm for solving $\Phi(\mu) = 0$, which is basically Newton's method with a form of safeguarding, as given below.

Algorithm 1 (Solving $\Phi(\mu) = 0$) *Given tolerance τ_1 , set $\mu = 0$, $\mu_\ell = 0$, $\mu_r = \infty$.*

Step 1 *Solve $\mathcal{P}(\mu)$ to obtain $d(\mu)$ and evaluate $\Phi(\mu)$.*

Step 2 *If $\mu = 0$ and $\Phi(\mu) > 0$, or $\|A^T d(\mu) + h\| - \theta / \theta \leq \tau_1$, let $d^* = d(\mu)$ and exit.*

Step 3 *If $\Phi(\mu) < 0$, then $\mu_\ell = \mu$ and $\Phi_\ell = \Phi(\mu)$; else $\mu_r = \mu$ and $\Phi_r = \Phi(\mu)$.*

Step 4 *Compute the Newton step: $\mu := \mu - \Phi(\mu) / \Phi'(\mu)$.*

Step 5 *If $\mu \notin (\mu_\ell, \mu_r)$, use the Regula Falsi step: $\mu := \mu_\ell - \Phi_\ell(\mu_r - \mu_\ell) / (\Phi_r - \Phi_\ell)$.*

Step 6 *Go to Step 1.*

A cause of concern about the above algorithm is probably the need to solve the standard trust-region subproblem $\mathcal{P}(\mu)$ which may be considered to be expensive. The solution to $\mathcal{P}(\mu)$ is now well-understood and very efficient methods have been developed during the last decade (see [6] and [7], for example), especially for positive definite matrices B . The method perhaps most often used is to solve the equation

$$\Psi(\mu, \lambda) \equiv \frac{1}{\|d(\mu, \lambda)\|} - \frac{1}{\delta} = 0, \quad (4.2)$$

for λ , while μ is fixed, using Newton's method which was proposed by Reinsch [12] and Hebden [5] independently. A basic algorithm for solving the above equation can be found in [7, p.p.47], which involves the Cholesky factorization of $H(\mu, \lambda)$. In our case, since $H(\mu, \lambda)$ is positive definite, the success of this algorithm is guaranteed. In practice, it has been reported [6] that on the average less than two iterations (factorizations) are needed to obtain an approximate solution to $\lambda(\mu)$. Since there is no need for a line search in our algorithm and no need for special measures to maintain feasibility as are required by Yuan's algorithm [15], the need for solving $\mathcal{P}(\mu)$ does not seem to be a formidable overhead.

We have used Newton's method to solve the problem $\mathcal{P}(\mu)$ (i.e., $\Psi(\mu, \lambda) = 0$ for a fixed μ) as given in Moré and Sorensen's paper [7, p.p.47], but we have added the restriction $\mu \geq 0$ and used a practical stopping criterion. The algorithm is as follows.

Algorithm 2 (Solving $\Psi(\mu, \lambda) = 0$ for a fixed μ) Given μ and tolerance τ_2 , set λ .

Step 1 Let $H = B + \mu AA^T + \lambda I$ and solve $Hd = -g$ for d .

Step 2 If $\lambda = 0$ and $\|d\| \leq \delta$, or $|\|d\| - \delta|/\delta \leq \tau_2$, then let $d^* = d$ and exit.

Step 3 Compute the restricted Newton step:

$$\lambda := \max \left(0, \lambda + \frac{d^T d}{d^T H^{-1} d} \frac{\|d\| - \delta}{\delta} \right).$$

Step 4 Go to Step 1.

Interested readers are referred to Moré and Sorensen's paper [7] for more details about the trust-region subproblem for unconstrained optimization.

We note from (2.15) that in order to compute the derivative $\Phi'(\mu)$, the vectors $H(\mu)^{-1}d(\mu)$ and $H(\mu)^{-1}y(\mu)$ are needed, which can be readily computed if the Cholesky factor of $H(\mu)$ is stored while solving $\mathcal{P}(\mu)$. In fact, the first vector is already available after solving $\mathcal{P}(\mu)$.

When both the constraints (1.2) and (1.3) are binding, our algorithm can be roughly summarized as follows. By the complementarity, CDT problem reduces to the nonlinear system of two equations with two variables

$$\begin{aligned} \Psi(\mu, \lambda) &= 0, \\ \Phi(\mu, \lambda) &= 0. \end{aligned}$$

We eliminate the variable λ by solving the first equation while fixing μ . After substitution, we reduce the above system into a single equation $\Phi(\mu, \lambda(\mu)) = 0$. Fortunately, both the first equation (for fixed μ) and the second equations (after the substitution $\lambda = \lambda(\mu)$) are well-behaved monotone functions and easy to solve.

Algorithm 1 and 2 have been programmed in Matlab and run on a SUN 3/160 workstation network at Rice University with a machine epsilon about 2.22×10^{-16} .

4.2 Results for Yuan’s Problems

In his paper [15], Yuan used five small test problems in testing his algorithm (see his paper for details). We ran our algorithm on these five problems, called Y1 through Y5. We chose the stopping tolerances as $\tau_1 = 10^{-3}$ and $\tau_2 = 0.1\tau_1$, respectively, for Algorithm 1 and 2. The final solutions we obtained have at least three and often four significant digits in agreement with Yuan’s solutions. We tabulate the number of Cholesky factorizations of $H(\mu, \lambda)$, which is the leading work in both ours and Yuan’s algorithms, in Table 1. Because ours and Yuan’s results were obtained on different machines with different stopping criteria, Table 1 should not be considered as a rigorous comparison, but rather a reasonable indicator.

Table 1: Number of Factorizations for Yuan’s Problems

Test Problem	Y1	Y2	Y3	Y4	Y5
Yuan’s Algorithm	8	6	11	14	11
Our Algorithm	7	2	2	15	10

From the table, we can see that for Problems Y2 and Y3 for which only one of the two constraints is active (i.e., either $\mu = 0$ or $\lambda = 0$), our algorithm appears to be much faster than Yuan’s. This probably should not be too much of a surprise, because unlike Yuan’s algorithm, which searches in both μ and λ directions simultaneously, our algorithm searches along the two directions alternatively. For the remaining three problems for which both constraints are active, our algorithm seems to be comparable with Yuan’s algorithm in terms of number of matrix factorizations.

4.3 Results for Random Problems

Randomly generated problems are also used in our preliminary numerical experiments. We use the Matlab M-file “rand” to generate normally distributed random numbers to form the needed matrices and vectors: A, B, g and h . To ensure the positive definiteness of B , we first generate a $n \times n$ matrix C and then let $B = C^T C$. The trust-region radius δ is

also generated randomly but the absolute value is taken to ensure positivity. To ensure the constant $\theta > \theta_{\min}$ (see (1.4)), we let θ be the best decrease on $\|A^T d + h\|$ in the steepest descent direction within the trust-region $\|d\| \leq \delta$, as was suggested by Celis, Dennis and Tapia [2]. More specifically,

$$\theta = \|A^T(-\alpha Ah) + h\|$$

where

$$\alpha = \min \left(\frac{\delta}{\|Ah\|}, \frac{h^T(A^T A)h}{h^T(A^T A)^2 h} \right).$$

We mention that another existing way of choosing θ is

$$\theta = \min\{\|A^T d + h\| : \|d\| \leq r\delta\}$$

for some $r < 1$, proposed by Powell and Yuan [11].

Our test results on five random problems, R1 to R5, with various sizes are presented in Table 2. To see how long a step is and how much reduction in the objective function as well as in the linearized constraint is achieved, we include in the table the four quantities $q(d^*)$ (note $q(0) = 0$), $\|d^*\|$, $\|h\|$ and $\|A^T d^* + h\|$ where d^* is the computed solution. Listed in the table are also the number of outer iterations and the number of Cholesky factorization N_{fac} needed for Algorithm 1 to reach a solution for each problem. We note that the number N_{fac} is also equal to the accumulated total number of inner iterations taken by Algorithm 2 embedded in Algorithm 1.

Table 2: Results for Random Problems

Problem	$n : m$	$q(d^*)$	$\ d^*\ $	$\ h\ $	$\ A^T d^* + h\ $	Iter	N_{fac}
R1	4 : 2	-0.02352	0.01449	0.30307	0.29380	2	6
R2	8 : 4	-1.05896	1.05458	2.52440	1.85602	4	12
R3	12 : 6	-0.28778	0.11656	2.20569	2.11969	3	9
R4	16 : 8	-1.74549	0.43507	2.67182	2.33413	3	9
R5	20 : 10	-1.26240	0.81427	3.04088	2.32549	3	10

Many more experiments with random problems have been done. The set of problems presented in Table 2 is just a small sample. The choice of this sample is such that both constraints (1.2) and (1.3) are binding but otherwise is arbitrary. We have found that it is often easier to solve problems with only one constraint binding instead of two. We also have observed from our computational experiences that the extent of difficulty in solving a CDT problem by our algorithm is mainly dictated by, aside from the conditioning of involved matrices, the values of the constants θ and δ and their relation.

4.4 Accuracy vs. Computational work

From the nature of the trust-region strategy as a device of enforcing global convergence and also from theoretical results developed for trust-region algorithms in unconstrained optimization, it is safe to say that in general one need not solve the CDT problem to a high accuracy. A reasonably good approximate solution is all we are asking for at least from the practical point of view. Therefore, we believe that our stopping criteria can be further relaxed without affecting the practical performance of the trust-region algorithm. In doing so, the effort required to obtain a satisfactory step can generally be significantly reduced. To illustrate this, we ran our algorithm on a random problem for stopping tolerances $\tau_1 = 10^{-3}, 10^{-2}, 10^{-1}$ and $\tau_2 = 0.1\tau_1$. The data for this problem are $n = 10, m = 5, \delta = 1, \theta = 3$ and $\|h\| = 3.704085$. The relevant quantities obtained from the algorithm at termination are given in Table 3, where again N_{fac} denotes the number of Cholesky factorization.

Table 3: Effects of Stopping Tolerances

Tolerance	$q(d)$	$\ d\ $	$\ A^T d + h\ $	μ	λ	N_{fac}
$\tau_1 = 10^{-3}$	-1.17216	1.00000	3.00003	0.52841	0.05215	12
$\tau_1 = 10^{-2}$	-1.15070	1.00019	2.98652	0.53393	0.05168	10
$\tau_1 = 10^{-1}$	-0.84760	1.00042	2.80405	0.61749	0.07847	6

5 Concluding remarks

Compared with Yuan's algorithm [15], the only existing one for computing an ℓ_2 -norm CDT trust-region step in all of \mathbf{R}^n that we are currently aware of, our algorithm is conceptually simpler because zero finding for a monotone function of a single variable is better understood than locating a maximizer in the first quadrant of a concave function of two variables. The limited numerical experiments do seem to confirm that our algorithm is indeed reliable and robust as expected.

As far as efficiency is concerned, we believe that methods of any kind, including ours, for solving the n -dimensional ℓ_2 -norm CDT problem are not likely to be cheap in the context of computing an iterative step for a trust-region algorithm for which the CDT problem is merely a subproblem that has to be solved at each iteration. On the other hand, solving the CDT problem supposedly produces better iterative steps and therefore enhances robustness of global convergence. From our computational experiments, we have been led to believe that our approach can produce good steps for small-sized problems (say, $n \leq 20$) at very affordable costs, given the fact that high accuracy in computed steps is not needed in trust-region algorithms.

To summarize, we conclude that our method has the potential to be used as a building block, along with a BFGS-type (reduced or full) Hessian approximation technique, to construct reliable and robust trust-region algorithms for small-sized problems in constrained optimization.

Acknowledgment

The author would like to thank Professors John Dennis and Richard Tapia for their careful reading of the manuscript and many valuable suggestions that have brought significant improvements to this paper.

References

- [1] M.R. Celis, J.E. Dennis, J.M. Martínez, R.A. Tapia, and K. Williamson. An algorithm based on a convenient trust-region subproblem for nonlinear programming. Technical report, Department of Mathematical Sciences, Rice University, In preparation, 1989.
- [2] M.R. Celis, J.E. Dennis, and R.A. Tapia. A trust region strategy for nonlinear equality constrained optimization. In P.T. Boggs, R.H. Byrd, and R.B.Schnabel, editors, *Numerical Optimization*, pages 71–82. SIAM, Philadelphia, 1985.
- [3] T. F. Coleman and A. R. Conn. On the local convergence of quasi-Newton methods for the nonlinear programming problems. *SIAM J. Num. Anal.*, 21:755–769, 1984.
- [4] D. M. Gay. Computing optimal locally constrained steps. *SIAM J. Sci. Stat. Comput.*, 2:186–197, 1981.
- [5] M. D. Hebden. An algorithm for minimization using exact second derivatives. Technical Report TP 515, A.E.R.E. Harwell, England, 1973.
- [6] J. J. Moré. Recent developments in algorithms and software for trust region methods. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming Study, the State of the Art*, pages 258–287. Springer-Verlag, Berlin, 1983.
- [7] J. J. Moré and D. C. Sorensen. Newton’s method. In Gene H. Golub, editor, *Studies in Numerical Analysis, (MAA Studies in Mathematics, Volume 24)*, pages 29–82. The Mathematical Association of America, 1984.
- [8] J. Nocedal and M. Overton. Projected Hessian updating algorithms for nonlinear constrained optimization. *SIAM J. Num. Anal.*, 22:821–850, 1985.
- [9] J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press, New York, 1970.

- [10] M. J. D. Powell. Variable metric methods for constrained optimization. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming Study, the State of the Art*, pages 288–311. Springer-Verlag, Berlin, 1983.
- [11] M.J.D. Powell and Y. Yuan. A trust region algorithm for equality constrained optimization. Technical Report NA2, DAMPT, University of Cambridge, 1986, to appear in *Math. Prog.*.
- [12] C. H. Reinsch. Smoothing by spline functions, II. *Numer. Math.*, 16:451–454, 1971.
- [13] D. C. Sorensen. Newton’s method with a model trust region modification. *SIAM J. Num. Anal.*, 19:409–426, 1982.
- [14] Y. Yuan. On a subproblem of trust region algorithms for constrained optimization. Technical Report NA10, DAMPT, University of Cambridge, 1987.
- [15] Y. Yuan. A dual algorithm for minimizing a quadratic function with two quadratic constraints. Technical Report NA3, DAMPT, University of Cambridge, 1988.