Trust-region interior-point method for large sparse l_1 optimization

L.Lukšan, C. Matonoha, J. Vlček [1]

Institute of Computer Science, Academy of Sciences of the Czech Republic,
Pod Vodárenskou věží 2, 182 07 Praha 8.
email: {luksan,matonoha,vlcek}@cs.cas.cz

1 Introduction

Consider the problem

$$\min F(x), \quad x \in \mathcal{R}^n,$$

where $F : \mathcal{R}^n \to \mathcal{R}$ is a twice continuously differentiable objective function. Basic optimization methods (trust-region and line-search methods) generate points $x_i \in \mathcal{R}^n$, $i \in \mathcal{N}$, in such a way that x_1 is arbitrary and

$$x_{i+1} = x_i + \alpha_i d_i, \quad i \in \mathcal{N}, \tag{1}$$

where $d_i \in \mathcal{R}^n$ are direction vectors and $\alpha_i > 0$ are step sizes.

For a description of trust-region methods we define the quadratic function

$$Q_i(d) = \frac{1}{2} d^T B_i d + g_i^T d$$

which locally approximates the difference $F(x_i + d) - F(x_i)$, the vector

$$\omega_i(d) = (B_i d + g_i)/\|g_i\|$$

for the accuracy of a computed direction, and the number

$$\rho_i(d) = \frac{F(x_i + d) - F(x_i)}{Q_i(d)}$$

for the ratio of actual and predicted decrease of the objective function. Here $g_i = g(x_i) = \nabla F(x_i)$ and $B_i \approx \nabla^2 F(x_i)$ is an approximation of the Hessian matrix at the point $x_i \in \mathcal{R}^n$.

Trust-region methods are based on approximate minimizations of $Q_i(d)$ on the balls $\|d\| \leq \Delta_i$ followed by updates of radii $\Delta_i > 0$. Direction vectors $d_i \in \mathcal{R}^n$ are chosen to satisfy the conditions

$$\|d_i\| \leq \Delta_i, \tag{2}$$

$$\|d_i\| < \Delta_i \quad \Rightarrow \quad \|\omega_i(d_i)\| \leq \overline{\omega}, \tag{3}$$

$$-Q_i(d_i) \geq \underline{\sigma} \|g_i\| \min(\|d_i\|, \|g_i\|/\|B_i\|), \tag{4}$$

[1]This work was supported by the Grant Agency of the Czech Academy of Sciences, project No. IAA1030405, the Grant Agency of the Czech Republic, project No. 201/06/P397, and the institutional research plan No. AV0Z10300504.

where $0 \leq \overline{\omega} < 1$ and $0 < \underline{\sigma} < 1$. Step sizes $\alpha_i \geq 0$ are selected so that

$$\rho_i(d_i) \leq 0 \quad \Rightarrow \quad \alpha_i = 0, \tag{5}$$

$$\rho_i(d_i) > 0 \quad \Rightarrow \quad \alpha_i = 1. \tag{6}$$

Trust-region radii $0 < \Delta_i \leq \overline{\Delta}$ are chosen in such a way that $0 < \Delta_1 \leq \overline{\Delta}$ is arbitrary and

$$\rho_i(d_i) < \underline{\rho} \quad \Rightarrow \quad \underline{\beta}\|d_i\| \leq \Delta_{i+1} \leq \overline{\beta}\|d_i\|, \tag{7}$$

$$\rho_i(d_i) \geq \underline{\rho} \quad \Rightarrow \quad \Delta_i \leq \Delta_{i+1} \leq \overline{\Delta}, \tag{8}$$

where $0 < \underline{\beta} \leq \overline{\beta} < 1$ and $0 < \underline{\rho} < 1$.

2 Survey of trust-region methods

A crucial part of each trust-region method is a direction determination. There are various commonly known methods for computing direction vectors satisfying conditions (2)-(4) which we now mention briefly. To simplify the notation we omit the major index i, use the inner index k, and write the symbol $\succeq 0$ to indicate that the matrix is positive semidefinite.

2.1 Moré-Sorensen 1983

The most sophisticated method is based on a computation of the optimal locally constrained step. In this case, the vector $d \in \mathcal{R}^n$ is obtained by solving the subproblem

$$\text{minimize} \quad Q(d) = \frac{1}{2} d^T B d + g^T d \quad \text{subject to} \quad \|d\| \leq \Delta. \tag{9}$$

Necessary and sufficient conditions for this solution are

$$\|d\| \leq \Delta, \quad (B + \lambda I)d = -g, \quad B + \lambda I \succeq 0, \quad \lambda \geq 0, \quad \lambda(\Delta - \|d\|) = 0. \tag{10}$$

The Moré-Sorensen method [13] is based on solving the nonlinear equation

$$\frac{1}{\|d(\lambda)\|} = \frac{1}{\Delta} \quad \text{with} \quad (B + \lambda I)d(\lambda) + g = 0$$

by the Newton's method, possibly the modified Newton's method [19] using the Choleski decomposition of $B + \lambda I$. This method is very robust but requires 2-3 Choleski decompositions for one direction determination on the average.

2.2 Powell 1970, Dennis-Mei 1975

Simpler methods are based on minimization of $Q(d)$ on the two-dimensional subspace containing the Cauchy and Newton steps

$$d_C = -\frac{g^T g}{g^T B g}\, g, \quad d_N = -B^{-1} g.$$

The most popular is the dogleg method [2], [15], where $d = d_N$ if $\|d_N\| \leq \Delta$ and $d = (\Delta/\|d_C\|)d_C$ if $\|d_C\| \geq \Delta$. In the remaining case, d is a combination of d_C and d_N such that $\|d\| = \Delta$. This method requires only one Choleski decomposition for one direction determination.

2.3 Steihaug 1983, Toint 1981

If B is not sufficiently small or sparse or explicitly available, then it is either too expensive or not possible to compute its Choleski factorization. In this case, methods based on matrix-vector multiplications are more convenient.

Steihaug [20] and Toint [21] proposed a technique for finding an approximate solution of (9) that does not require the exact solution of a linear system but still produce an improvement on the Cauchy point. This implementation is based on the conjugate gradient algorithm [14] for solving the linear system $Bd = -g$. We either obtain an unconstrained solution with a sufficient precision or stop on the trust-region boundary. The latter possibility occurs if either a negative curvature is encountered or the constraint is violated. This method is based on the fact that $Q(d_{k+1}) < Q(d_k)$ and $\|d_{k+1}\| > \|d_k\|$ hold in the subsequent CG iterations if the CG coefficients are positive and no preconditioning is used. Note that the inequality $\|d_{k+1}\| > \|d_k\|$ is not satisfied in general if a preconditioner C (symmetric and positive definite) is used. In this case we have $\|d_{k+1}\|_C > \|d_k\|_C$, where $\|d_k\|_C^2 = d_k^T C d_k$.

The CG steps can be combined with the Newton step $d_N = -B^{-1}g$ in the multiple dogleg method [20]. Let $k \ll n$ (usually $k = 5$) and d_k be a vector obtained after k CG steps of the Steihaug-Toint method. If $\|d_k\| < \Delta$, we use d_k instead of $d_C = d_1$ in the dogleg method.

2.4 Preconditioned Steihaug-Toint

There are two possibilities how the Steihaug-Toint method can be preconditioned. The first way uses the norms $\|d_i\|_{C_i}$ (instead of $\|d_i\|$) in (2)–(8), where C_i are preconditioners chosen. This possibility has been tested in [5] and showed that such a way is not always efficient. This is caused by the fact that the norms $\|d_i\|_{C_i}$, $i \in \mathcal{N}$, vary considerably in the major iterations and the preconditioners C_i, $i \in \mathcal{N}$, can be ill-conditioned. The second way uses the Euclidean norms in (2)–(8) even if arbitrary preconditioners C_i, $i \in \mathcal{N}$, are used. In this case, the trust-region can be leaved prematurely and the direction vector obtained can be farther from the optimal locally constrained step than that obtained without preconditioning. This shortcoming is usually compensated by the rapid convergence of the preconditioned CG method. Our computational experiments indicated that the second way is more efficient in general.

2.5 Gould-Lucidi-Roma-Toint 1997

Although the Steihaug-Toint method is certainly the most commonly used in trust-region methods, the resulting direction vector may be rather far from the optimal solution even in the unpreconditioned case. This drawback can be overcome by using the Lanczos process [5], as we now explain. Initially, the conjugate gradient algorithm is used as in the Steihaug-Toint method. At the same time, the Lanczos tridiagonal matrix is constructed from the CG coefficients. If a negative curvature is encountered or the constraint is violated, we switch to the Lanczos process. In this case, $d = Z\tilde{d}$, where $\tilde{d}$ is obtained by

minimizing the quadratic function

$$\frac{1}{2}\tilde{d}^T T\tilde{d} + \|g\| e_1^T \tilde{d} \tag{11}$$

subject to $\|\tilde{d}\| \leq \Delta$. Here $T = Z^T BZ$ (with $Z^T Z = I$) is the Lanczos tridiagonal matrix and e_1 is the first column of the unit matrix. Using a preconditioner C, the preconditioned Lanczos method generates basis such that $Z^T CZ = I$. Thus we have to use the norms $\|d_i\|_{C_i}$ in (2)–(8), i.e., the first way of preconditioning, which can be inefficient when $C_i,\ i \in \mathcal{N}$, vary considerably in the trust-region iterations or are ill-conditioned.

2.6 Shifted Steihaug-Toint

This method applies the Steihaug-Toint method to the shifted subproblem

$$\text{minimize} \quad \tilde{Q}(d) = Q_{\tilde{\lambda}}(d) = \frac{1}{2} d^T (B + \tilde{\lambda} I) d + g^T d \quad \text{s.t.} \quad \|d\| \leq \Delta. \tag{12}$$

The number $\tilde{\lambda} \geq 0$, which approximates λ in (10), is found by solving a small-size subproblem of type (11) with the tridiagonal matrix T obtained by using a small number of Lanczos steps. This method, like method [5], combines good properties of the Moré-Sorensen and the Steihaug-Toint methods. Moreover, it can be successfully preconditioned by the second way. The point on the trust-region boundary obtained by this method is usually closer to the optimal solution in comparison with the point obtained by the original Steihaug-Toint method.

The shifted Steihaug-Toint method [7] consists of the three major steps.

1. Carry out $k \ll n$ steps of the unpreconditioned Lanczos method (described, e.g., in [5]) to obtain the tridiagonal matrix $T = T_k = Z_k^T BZ_k$.

2. Solve the subproblem

$$\text{minimize} \quad \frac{1}{2}\tilde{d}^T T\tilde{d} + \|g\| e_1^T \tilde{d} \quad \text{subject to} \quad \|\tilde{d}\| \leq \Delta, \tag{13}$$

   using the Moré-Sorensen method [13] to obtain the Lagrange multiplier $\tilde{\lambda}$.

3. Apply the (preconditioned) Steihaug-Toint method [20] to subproblem (12) to obtain the direction vector $d = d(\tilde{\lambda})$.

2.7 Hager 2001

There are several recently developed techniques for large scale trust-region subproblems that are not based on conjugate gradients. Hager [6] developed a method that solves (9) with the additional constraint that d is contained in a low-dimensional subspace. The subspaces are modified in successive iterations to obtain quadratic convergence to the optimum and they are designed to contain both the prior iterate and the iterate that is generated by applying one step of the sequential quadratic programming algorithm [1]

to (9). At first, the Lanczos method is used to generate an orthonormal basis for the $k-$dimensional Krylov subspace (usually $k = 10$). Then problem (9) is reduced to the $k-$dimensional one to obtain an initial iterate. The main loop consists in seeking vectors $d \in \mathcal{S}$ where $\mathcal{S}$ contains the following four vectors:

- The previous iterate. This causes that the value of the objective function can only decrease in consecutive iterations.
- The vector $Bd + g$. It ensures descent if the current iterate does not satisfy the first-order optimality conditions.
- An estimate for an eigenvector of B associated with the smallest eigenvalue. It will dislodge the iterates from a nonoptimal stationary point.
- The SQP iterate. The convergence is locally quadratic if the subspace $\mathcal{S}$ contains the iterate generated by one step of the sequential quadratic programming algorithm applied to (9).

An orthonormal basis for the subspace $\mathcal{S}$ is constructed, original problem (9) is reduced to the four-dimensional one, and a new iterate d is found via this small subproblem. The iteration is finished as soon as $\|(B + \lambda I)d + g\|$ with a Lagrange multiplier λ is smaller than some sufficiently small tolerance. The SQP method is equivalent to the Newton's method applied to the nonlinear system

$$(B + \lambda I)d + g = 0, \quad \frac{1}{2}d^T d - \frac{1}{2}\Delta^2 = 0.$$

The Newton iterate can be expressed in the following way:

$$d_{SQP} = d + z, \quad \lambda_{SQP} = \lambda + \nu,$$

where z and ν are solutions of the linear system

$$\begin{aligned} (B + \lambda I)z + d\,\nu &= -((B + \lambda I)d + g), \\ d^T z &= 0, \end{aligned}$$

which can be solved by preconditioned MINRES or CG methods. The latter case with the incomplete Choleski-type decomposition of the matrix $B + \lambda I$ has shown to be more efficient in practice.

2.8 Rojas-Santos-Sorensen 1997, 2000

Another approach for finding the direction vector d is based on the idea of Sorensen [17],[18]. Consider the bordered matrix

$$B_\alpha = \begin{pmatrix} \alpha & g^T \\ g & B \end{pmatrix}$$

where α is a real number and observe that

$$\frac{\alpha}{2} + Q(d) = \frac{1}{2}(1, d^T)B_\alpha \begin{pmatrix} 1 \\ d \end{pmatrix}.$$

Thus there exists a value of the parameter α such that we can rewrite problem (9) as

$$\text{minimize} \quad \frac{1}{2} d_\alpha^T B_\alpha d_\alpha \quad \text{subject to} \quad \|d_\alpha\|^2 \leq 1 + \Delta^2, \quad e_1^T d_\alpha = 1, \tag{14}$$

where $d_\alpha = (1, d^T)^T$ and e_1 is the first canonical unit vector in $\mathcal{R}^{n+1}$. This formulation suggests that we can find the desired solution in terms of an eigenpair of B_α. The resulting algorithm is superlinearly convergent.

3 Comparison of methods

We present a numerical comparison of nine methods for computing direction vectors satisfying conditions (2)-(4):

- MS – the method of Moré and Sorensen [13] for computing the optimal locally constrained step.
- DL – the dogleg strategy of Powell [15] or Dennis and Mei [2].
- MDL – the multiple dogleg strategy mentioned in [20].
- ST – the basic (unpreconditioned) Steihaug [20] and Toint [21] method.
- SST – the basic (unpreconditioned) shifted Steihaug-Toint method [7].
- GLRT – the method of Gould, Lucidi, Roma and Toint [5] which combines CG method with the Lanczos process to give a good approximation of the optimal locally constrained step.
- PH – the preconditioned Hager method described in [6]. The incomplete Choleski preconditioner is used.
- PST – the preconditioned Steihaug-Toint method. The incomplete Choleski preconditioner is used.
- PSST – the preconditioned shifted Steihaug-Toint method. The incomplete Choleski preconditioner is used.

These methods are implemented in the interactive system for universal functional optimization UFO [10] as subroutines for solving trust-region subproblems. They were tested by using two collections of 22 sparse test problems with 1000 and 5000 variables (subroutines `TEST14` and `TEST15` described in [11], which can be downloaded from `www.cs.cas.cz/~luksan/test.html`). The results are given in Tables 1 and 2, where `NIT` is the total number of iterations, `NFV` is the total number of function evaluations, `NFG` is the total number of gradient evaluations, `NDC` is the total number of Choleski-type decompositions (complete for methods MS, DL, MDL and incomplete for methods PH, PST, PSST), `NMV` is the total number of matrix-vector multiplications, and `Time` is the total computational time in seconds. Note that `NFG` is much greater than `NFV` in Table 1 since the Hessian matrices are computed by using gradient differences. At the same time, the problems referred in Table 2 are the sums of squares having the form $F = (1/2) f^T(x) f(x)$ and `NFV` denotes the total number of the vector $f(x)$ evaluations. Since $f(x)$ is used in the expression $g(x) = J^T(x) f(x)$, where $J(x)$ is the Jacobian matrix of $f(x)$, `NFG` is comparable with `NFV` in this case.

N	Method	NIT	NFV	NFG	NDC	NMV	Time
1000	MS	1911	1952	8724	3331	1952	3.13
	DL	2272	2409	10653	2195	2347	2.94
	MDL	2132	2232	9998	1721	21670	3.17
	ST	3475	4021	17242	0	63016	5.44
	SST	3149	3430	15607	0	75044	5.97
	GLRT	3283	3688	16250	0	64166	5.40
	PH	1958	2002	8975	3930	57887	5.86
	PST	2608	2806	12802	2609	5608	3.30
	PSST	2007	2077	9239	2055	14440	2.97
5000	MS	8177	8273	34781	13861	8272	49.02
	DL	9666	10146	42283	9398	9936	43.37
	MDL	8913	9244	38846	7587	91784	48.05
	ST	16933	19138	84434	0	376576	134.52
	SST	14470	15875	70444	0	444142	146.34
	GLRT	14917	16664	72972	0	377588	132.00
	PH	8657	8869	37372	19652	277547	127.25
	PST	11056	11786	53057	11057	23574	65.82
	PSST	8320	8454	35629	8432	59100	45.57

Table 1: Comparison of methods using TEST14.

N	Method	NIT	NFV	NFG	NDC	NMV	Time
1000	MS	1946	9094	9038	3669	2023	5.86
	DL	2420	12291	12106	2274	2573	9.00
	MDL	2204	10586	10420	1844	23139	7.86
	ST	2738	13374	13030	0	53717	11.11
	SST	2676	13024	12755	0	69501	11.39
	GLRT	2645	12831	12547	0	61232	11.30
	PH	1987	9491	9444	6861	84563	11.11
	PST	3277	16484	16118	3278	31234	11.69
	PSST	2269	10791	10613	2446	37528	8.41
5000	MS	7915	33607	33495	14099	8047	89.69
	DL	9607	42498	41958	9299	9963	128.92
	MDL	8660	37668	37308	7689	91054	111.89
	ST	11827	54699	53400	0	307328	232.70
	SST	11228	51497	50333	0	366599	231.94
	GLRT	10897	49463	48508	0	300580	214.74
	PH	8455	36434	36236	20538	281736	182.45
	PST	9360	41524	41130	9361	179166	144.40
	PSST	8634	37163	36881	8915	219801	140.44

Table 2: Comparison of methods using TEST15.

The results in Tables 1 and 2 require several comments. All problems are sparse with a simple sparsity pattern. For this reason, the methods based on complete Choleski-type decompositions (CD) are very efficient, much better than unpreconditioned methods based on matrix-vector multiplications (MV). Since `TEST14` contains reasonably conditioned problems, the preconditioned MV methods are competitive with the CD methods. On the contrary, `TEST15` contains several very ill-conditioned problems (one of them had to be removed) and thus the CD methods work better than the MV methods. In general, the CD methods are very efficient for ill-conditioned but reasonably sparse problems but if the problems do not have sufficiently sparse Hessian matrices, then the CD methods can be much worse than the MV methods. The efficiency of the MV methods also strongly depends on a suitable preconditioner.

4 Large scale l_1 optimization

Consider the l_1 optimization problem: Minimize the function

$$F(x) = \sum_{i=1}^{m} |f_i(x)|, \tag{15}$$

where $f_i : \mathcal{R}^n \to \mathcal{R}$, $0 \le i \le m$, are smooth functions depending on n_i variables. We assume that the function $F(x)$ is partially separable, which means that n, $m = \mathcal{O}(n)$ are large and $n_i = \mathcal{O}(1)$, $0 \le i \le m$, are small.

The minimization of F is equivalent to the sparse nonlinear programming problem with $n + m$ variables $x \in \mathcal{R}^n$, $z \in \mathcal{R}^m$:

$$\text{minimize} \quad \sum_{i=1}^{m} z_i \quad \text{subject to} \quad -z_i \le f_i(x) \le z_i, \quad 1 \le i \le m. \tag{16}$$

This problem satisfies the Mangasarian-Fromowitz constraint qualification conditions and the necessary first-order (KKT) conditions have the form

$$\sum_{i=1}^{m} u_i \nabla f_i(x) = 0, \quad z_i = |f_i(x)|, \quad |u_i| \le 1, \quad \text{and} \quad u_i = \frac{f_i(x)}{|f_i(x)|} \quad \text{if} \quad |f_i(x)| > 0. \tag{17}$$

Problem (16) can be solved by an arbitrary nonlinear programming method utilizing sparsity (sequential linear programming, sequential quadratic programming, interior-point, and nonsmooth equation). Original problem (15) can also be solved by the trust-region methods described in [16] and [22].

We introduce a trust-region interior-point method [9] that utilizes a special structure of the l_1 optimization problem. Constrained problem (16) is replaced by a sequence of unconstrained problems

$$\begin{aligned} \text{minimize} \quad B(x, z; \mu) &= \sum_{i=1}^{m} z_i - \mu \sum_{i=1}^{m} \log(z_i - f_i(x)) - \mu \sum_{i=1}^{m} \log(z_i + f_i(x)) \\ &= \sum_{i=1}^{m} z_i - \mu \sum_{i=1}^{m} \log(z_i^2 - f_i^2(x)) \end{aligned} \tag{18}$$

with a barrier parameter $0 < \mu \leq \overline{\mu}$, where we assume that $z_i > |f_i(x)|$, $1 \leq i \leq m$, and $\mu \to 0$ monotonically. Here $B(x, z; \mu) : \mathcal{R}^{n+m} \to \mathcal{R}$ is a function of $n + m$ variables $x \in \mathcal{R}^n$, $z \in \mathcal{R}^m$.

The interior-point method is a trust-region modification of the Newton method and is iterative, so it generates a sequence of points $x_k \in \mathcal{R}^n,\ k \in \mathcal{N}$. An approximation of the Hessian matrix is computed by gradient differences which can be carried out efficiently if the Hessian matrix is sparse. Since the Hessian matrix need not be positive definite in a non-convex case, a standard line-search realization cannot be used. There are two basic possibilities, either a trust-region approach or a line-search strategy with suitable restarts, which eliminate this insufficiency. We have implemented and tested both these possibilities and our tests have shown that the first possibility is more efficient.

5 Description of the method

Differentiating $B(x, z; \mu)$ given by (18) we obtain necessary conditions for a minimum in the form

$$A(x)u(x, z; \mu) = 0, \quad u(x, z; \mu) = Z^{-1} f(x), \tag{19}$$

where

$$A(x) = [g_1(x), \ldots, g_m(x)], \quad g_i(x) = \nabla f_i(x), \quad 1 \leq i \leq m, \quad Z = \mathrm{diag}(z_1, \ldots, z_m),$$

$$u(x, z; \mu) = [u_1(x, z_1; \mu), \ldots, u_m(x, z_m; \mu)]^T, \quad u_i(x, z_i; \mu) = \frac{2\mu f_i(x)}{z_i^2 - f_i^2(x)}, \quad 1 \leq i \leq m.$$

System of $n + m$ nonlinear equations (19) can be solved by the Newton method, which uses second-order derivatives. In every step of the Newton method we solve a set of $n + m$ linear equations to obtain increments Δx and Δz of x and z, respectively. These increments can be used for obtaining new quantities

$$x^+ = x + \alpha \Delta x, \quad z^+ = z + \alpha \Delta z,$$

where $\alpha > 0$ is a suitable step size. This is a standard way for solving general nonlinear programming problems.

The structure of $B(x, z; \mu)$ for special nonlinear programming problem (16) allows us to obtain a minimizer $z(x; \mu) \in \mathcal{R}$ of the function $B(x, z; \mu)$ for a given $x \in \mathcal{R}^n$. It can be proved that the function $B(x, z; \mu)$ (with x fixed) has a unique stationary point which is its global minimizer. This stationary point is characterized by the equations which have the solutions

$$z_i(x; \mu) = \mu + \sqrt{\mu^2 + f_i^2(x)}, \quad 1 \leq i \leq m. \tag{20}$$

Assuming $z = z(x; \mu)$ we denote $B(x; \mu) = B(x, z(x; \mu); \mu)$ and

$$u_i(x; \mu) = u_i(x, z(x; \mu); \mu) = \frac{f_i(x)}{z_i(x; \mu)} = \frac{f_i(x)}{\mu + \sqrt{\mu^2 + f_i^2(x)}}, \quad 1 \leq i \leq m. \tag{21}$$

In this case, the barrier function $B(x; \mu)$ depends only on x. In order to obtain a minimizer $(x, z) \in \mathcal{R}^{n+m}$ of $B(x, z; \mu)$, it suffices to minimize $B(x; \mu)$ over $\mathcal{R}^n$.

If a vector $d \in \mathcal{R}^n$ solves the equation

$$\nabla^2 B(x;\mu)d = -g(x;\mu), \tag{22}$$

where $g(x;\mu) = \nabla B(x;\mu) \neq 0$, and if the matrix

$$G(x;\mu) = \sum_{i=1}^{m} u_i(x;\mu)\nabla^2 f_i(x)$$

is positive definite, then the direction vector d is descent for $B(x;\mu)$, i.e. $d^T g(x;\mu) < 0$. Unfortunately, the positive definiteness of this matrix is not assured, which causes that the standard line-search methods cannot be used. For this reason, the trust-region methods (see Section 2) were developed. These methods use a direction vector obtained as an approximate minimizer of the quadratic subproblem

$$\text{minimize} \quad Q(d) = \frac{1}{2}d^T\nabla^2 B(x;\mu)d + g^T(x;\mu)d \quad \text{subject to} \quad \|d\| \leq \Delta, \tag{23}$$

where Δ is a trust-region radius. The direction vector d serves for obtaining a new point $x^+ \in \mathcal{R}^n$. Denoting

$$\rho(d) = \frac{B(x+d;\mu) - B(x;\mu)}{Q(d)} \tag{24}$$

we set

$$x^+ = x \quad \text{if} \quad \rho(d) < \underline{\rho} \quad \text{or} \quad x^+ = x + d \quad \text{if} \quad \rho(d) \geq \underline{\rho} \tag{25}$$

and update the trust-region radius in such a way that

$$\underline{\beta}\|d\| \leq \Delta^+ \leq \overline{\beta}\|d\| \quad \text{if} \quad \rho(d) < \overline{\rho} \quad \text{or} \quad \Delta \leq \Delta^+ \leq \overline{\gamma}\Delta \quad \text{if} \quad \rho(d) \geq \overline{\rho}, \tag{26}$$

where $0 < \underline{\rho} < \overline{\rho} < 1$ and $0 < \underline{\beta} \leq \overline{\beta} < 1 < \overline{\gamma}$.

We also use the maximum step length $\overline{\Delta}$ which has no theoretical significance but is very useful for practical computations. First, the problem functions can sometimes be evaluated only in a relatively small region (if they contain exponentials) so that the maximum step length is necessary. Secondly, the problem can be very ill-conditioned far from the solution point, thus large steps are unsuitable. Finally, if the problem has more local solutions, a suitably chosen maximum step length can cause a local solution with a lower value of F to be reached. Therefore, the maximum step length $\overline{\Delta}$ is a parameter, which is most frequently tuned.

A very important part is the update of the barrier parameter μ. There are two requirements which play opposite roles. First, $\mu \to 0$ should hold since this is the main property of every interior-point method. On the other hand, $\nabla^2 B(x;\mu)$ can be ill-conditioned if μ is too small. Thus the lower bound $\underline{\mu}$ for μ is used. We have tested various possibilities for the barrier parameter update including simple geometric sequences, which were proved to be unsuitable. Better results were obtained by setting

$$\mu^+ = \max(\underline{\mu}, \|g\|^2) \quad \text{if} \quad \rho(d) \geq \underline{\rho} \quad \text{and} \quad \|g\|^2 \leq \tau\mu \tag{27}$$

(where $0 < \tau < 1$) and $\mu^+ = \mu$ otherwise.

6 Implementation details

We have pointed out that the direction vector $d \in \mathcal{R}^n$ should be a solution of quadratic subproblem (23). Usually, an inexact approximate solution suffices. From the list of methods described in Section 2 we have used two approaches based on direct decompositions of the matrix $\nabla^2 B$.

The first strategy, the dogleg method, described in [2], [15], seeks d as a linear combination of the Cauchy step $d_C = -(g^T g / g^T \nabla^2 B g) g$ and the Newton step $d_N = -(\nabla^2 B)^{-1} g$. The Newton step is computed by using either the sparse Gill-Murray decomposition [4] or the sparse Bunch-Parlett decomposition [3]. The sparse Gill-Murray decomposition has the form $\nabla^2 B + E = LDL^T = R^T R$, where E is a positive semidefinite diagonal matrix (which is equal to zero when $\nabla^2 B$ is positive definite), L is a lower triangular matrix, D is a positive definite diagonal matrix and R is an upper triangular matrix. The sparse Bunch-Parlett decomposition has the form $\nabla^2 B = PLML^T P^T$, where P is a permutation matrix, L is a lower triangular matrix and M is a block-diagonal matrix with 1×1 or 2×2 blocks (which is indefinite when $\nabla^2 B$ is indefinite).

The second strategy, the optimum step method, computes a more accurate solution of (23) by using the Newton method applied to the nonlinear equation

$$\frac{1}{\|d(\lambda)\|} - \frac{1}{\Delta} = 0 \tag{28}$$

where $(\nabla^2 B + \lambda I) d(\lambda) = -g$. This way, described in [13] in more details, follows from the KKT conditions for (23). Since the Newton method applied to (28) can be unstable, the safeguards (lower and upper bounds to λ) are usually used.

7 Computational experiments

The primal interior-point method was tested by using two collections of 22 relatively difficult problems with an optional dimension chosen from [11], which can be downloaded from `www.cs.cas.cz/~luksan/test.html` as Test 14 and Test 15. The functions $f_i(x)$, $1 \leq i \leq m$, given in [11] serve for defining the objective function

$$F(x) = \sum_{1 \leq i \leq m} |f_i(x)|. \tag{29}$$

The first set of the tests concerns a comparison of interior-point methods with various trust-region and line-search strategies and the bundle variable metric method proposed in [12]. Medium-size test problems with 200 variables are used. The results of computational experiments are reported in Tables 3 and 4 where only summary results (over all 22 test problems) are given. Here `M` is the method used: T1 – the dogleg method with the Gill-Murray decomposition, T2 – the dogleg method with the Bunch-Parlett decomposition, T3 – the optimum step method with the Gill-Murray decomposition, L – the line-search method with restarts described in [8], B – the bundle variable metric method described in [12]; `NIT` is the total number of iterations, `NFV` is the total number of function evaluations, `NFG` is the total number of gradient evaluations, `NR` is the total number of restarts, `NL` is

the number of problems for which the best known local minimizer was not found (even if the parameter $\overline{\Delta}$ was tuned), NF is the number of problems for which no local minimizer was found (either a premature termination occurred or the number of function evaluations exceeded the upper bound), NT is the number of problems for which the parameter $\overline{\Delta}$ was tuned (for removing overflows and obtaining the best known local minimum), and Time is the total computational time in seconds.

M	NIT	NFV	NFG	NR	NL	NF	NT	Time
T1	2784	3329	23741	1	-	-	4	3.70
T2	2392	2755	19912	2	-	1	8	3.19
T3	3655	4161	32421	4	1	1	7	6.52
L	5093	12659	30350	1	1	-	6	4.58
B	34079	34111	34111	22	1	1	11	25.72

Table 3: Test 14 – 22 problems with 200 variables

M	NIT	NFV	NFG	NR	NL	NF	NT	Time
T1	3331	4213	18989	17	-	-	6	3.74
T2	3170	4027	17452	17	-	1	12	3.68
T3	5424	6503	31722	11	1	1	10	7.83
L	8183	20245	52200	36	2	-	9	10.90
B	34499	34745	34745	22	1	-	11	13.14

Table 4: Test 15 – 22 problems with 200 variables

The results introduced in Tables 3 and 4 indicate that the trust-region strategies are more efficient than the restarted line-search strategies in connection with the interior-point method for l_1 optimization. The trust-region interior-point method T1 is less sensitive to the choice of parameters and requires a lower number of iterations and a shorter computational time in comparison with the bundle variable metric method B proposed in [12]. Method T1 also finds the best known local minimum (if l_1 problems have several local solutions) more frequently (see the column NL in the tables).

The second set of the tests concerns a comparison of the interior-point method, realized as the dogleg method with the Gill-Murray decomposition, with the bundle variable metric method described in [12]. Large-scale test problems with 1000 variables are used. The results of computational experiments are given in Tables 5 and 6 where P is the problem number, NIT is the number of iterations, NFV is the number of function evaluations, NFG is the number of gradient evaluations, and F is the function value reached. The last row of every table contains the summary results including the total computational time in seconds. The bundle variable metric method was chosen for the comparison since it is based on a quite different principle and can also be used for the large sparse l_1 optimization.

	Trust-region interior-point method				Bundle variable metric method			
P	NIT	NFV	NFG	F	NIT	NFV	NFG	F
1	1594	1598	6380	0.166502E-09	7819	7842	7842	0.174023E-20
2	415	516	2912	0.106432E-08	127	130	130	0.735523E-17
3	32	33	231	0.604855E-07	89	89	89	0.359364E-14
4	27	39	196	269.499	81	81	81	269.499
5	30	31	186	0.107950E-06	39	39	39	0.122456E-14
6	32	33	462	0.611870E-07	100	100	100	0.110358E-12
7	18	20	171	336.937	211	211	211	336.937
8	18	19	342	761774.	36	39	39	761774.
9	212	259	3834	327.680	6181	6181	6181	327.682
10	970	1176	17460	0.386416E-01	14369	14369	14369	0.740271E-01
11	82	90	498	10.7765	319	319	319	10.7765
12	35	36	144	982.273	115	117	117	982.273
13	27	28	112	0.277182E-06	16	17	17	0.139178E-18
14	1	12	6	0.129382E-08	3	3	3	0.129382E-08
15	202	246	812	1.96106	3948	3957	3957	1.97013
16	161	169	972	0.435729E-15	4505	4556	4556	0.475529E-03
17	484	564	2910	0.165706E-11	441	443	443	0.857271E-06
18	2093	2538	12564	0.105340E-05	1206	1216	1216	0.129694E-03
19	15	16	96	59.5986	182	182	182	59.5986
20	1226	1529	7362	0.154869E-11	7828	7830	7830	0.102202E-04
21	21	22	132	2.13866	29	30	30	2.13866
22	1423	1770	8544	1.00000	337	341	341	1.00000
Σ	9118	10774	66332	Time=42.56	47981	48092	48092	Time=155.67

Table 5: Test 14 – 22 problems with 1000 variables

	Trust-region interior-point method				Bundle variable metric method			
P	NIT	NFV	NFG	F	NIT	NFV	NFG	F
1	1464	1477	5860	0.123345E-12	359	540	540	0.815757E-08
2	121	181	605	4.00000	453	473	473	0.153343E-07
3	27	31	168	0.775716E-09	114	114	114	0.374913E-08
4	65	76	264	648.232	53	54	54	648.232
5	6	7	42	0.655031E-14	285	285	285	0.422724E-05
6	8	9	126	0.754396E-13	560	560	560	0.649530E-08
7	73	111	296	12029.9	542	650	650	12029.9
8	83	100	252	0.230723E-06	939	942	942	0.380433E-03
9	532	609	3731	2777.75	4428	4429	4429	2780.11
10	103	148	618	658.048	1389	1389	1389	658.048
11	3452	3674	13812	0.821565E-14	411	454	454	0.838373E-09
12	652	773	3918	3117.36	1879	1882	1882	3125.85
13	165	212	996	14808.8	727	728	728	14808.8
14	162	201	1134	566.112	514	514	514	566.112
15	67	93	476	181.926	654	654	654	181.926
16	268	328	1883	66.5333	1376	1376	1376	66.5333
17	122	147	1107	0.146536E-13	9092	9092	9092	0.337978E-08
18	78	89	474	0.619504E-13	3160	3160	3160	0.754900
19	29	31	330	0.382360E-12	15933	15944	15944	0.239244E-08
20	69	86	420	0.131734E-10	1509	1699	1699	0.756975E-08
21	118	195	708	1326.92	425	426	426	1327.95
22	80	112	486	2993.36	9875	9875	9875	2993.37
Σ	7744	8690	37706	Time=30.03	54677	55240	55240	Time=155.90

Table 6: Test 15 – 22 problems with 1000 variables

The results introduced in Tables 5 and 6 confirm conclusions following from the previous tables. The trust-region interior-point method seems to be more efficient than the bundle variable metric method in all indicators. Especially, the computational time is much shorter and also the number of the best known local minima attained is greater in the case of the trust-region interior-point method.

References

[1] Boggs P.T., Tolle J.W.: *Sequential Quadratic Programming*, Acta Numerica, 1995, pp. 1-51.

[2] Dennis J.E., Mei H.H.W.: *An unconstrained optimization algorithm which uses function and gradient values*, Report No. TR 75-246, 1975.

[3] Duff I.S., Munksgaard M., Nielsen H.B., Reid J.K.: *Direct solution of sets of linear equations whose matrix is sparse and indefinite*, J. Inst. Maths. Applics., 23, 1979, pp. 235-250.

[4] Gill P.E., Murray W.: *Newton type methods for unconstrained and linearly constrained optimization*, Mathematical Programming, 7, 1974, pp. 311-350.

[5] Gould N.I.M., Lucidi S., Roma M., Toint P.L.: *Solving the trust-region subproblem using the Lanczos method*, Report No. RAL-TR-97-028, 1997.

[6] Hager W.W.: *Minimizing a quadratic over a sphere*, SIAM J. Optim., 12(1), 2001, pp. 188-208.

[7] Lukšan L., Matonoha C., Vlček J.: *A shifted Steihaug-Toint method for computing a trust-region step*, Technical Report V-914, Prague, ICS AS CR, 2004.

[8] Lukšan L., Matonoha C., Vlček J.: *Primal interior-point method for large sparse minimax optimization*, Technical Report V-941, Prague, ICS AS CR, 2005.

[9] Lukšan L., Matonoha C., Vlček J.: *Trust-region interior point method for large sparse l_1 optimization*, Technical Report V-942, Prague, ICS AS CR, 2005.

[10] Lukšan L., Tůma M., Hartman J., Vlček J., Ramešová N., Šiška M., Matonoha C.: *Interactive System for Universal Functional Optimization (UFO). Version 2006*, Report No. V977-06, Institute of Computer Science AS CS, Prague, 2006.

[11] Lukšan L., Vlček J.: *Sparse and partially separable test problems for unconstrained and equality constrained optimization*, Report No. V767-98, Institute of Computer Science AS CS, 1998.

[12] Lukšan L., Vlček J.: *Variable metric method for minimization of partially separable nonsmooth functions*, Pacific Journal on Optimization, 2, 2006, pp. 9-70.

[13] Moré J.J., Sorensen D.C.: *Computing a trust region step*, SIAM Journal on Scientific and Statistical Computations 4, 1983, pp. 553-572.

[14] Nocedal J., Wright S.J.: *Numerical Optimization*, Springer, New York, 1999.

[15] Powell M.J.D.: *A new algorithm for unconstrained optimization*, in Nonlinear Programming (Rosen J.B., Mangasarian O.L., Ritter K., eds.), Academic Press, London, 1970.

[16] Powell M.J.D.: *General Algorithms for Discrete Nonlinear Approximation Calculations*, in Approximation Theory IV (Chui G.K., Schumaker L.L., Ward J.D., eds.), Academic Press, New York, 1983.

[17] Rojas M., Santos S.A., Sorensen D.C.: *A new matrix-free algorithm for the large-scale trust-region subproblem*, SIAM J. Optim., 11(3), 2000, pp. 611-646.

[18] Sorensen D.C.: *Minimization of a large-scale quadratic function subject to a spherical constraint*, SIAM J. Optim., 7(1), 1997, pp. 141-161.

[19] Sorensen D.C.: *Newton's method with a model trust region modification*, SIAM J. Numer. Anal., 19(2), 1982, pp. 409-426.

[20] Steihaug T.: *The conjugate gradient method and trust regions in large-scale optimization*, SIAM Journal on Numerical Analysis, 20, 1983, pp. 626-637.

[21] Toint P.L.: *Towards an efficient sparsity exploiting Newton method for minimization*, in Sparse Matrices and Their Uses (Duff I.S., ed.), Academic Press, London, 1981, pp. 57-88.

[22] Yuan Y.: *On the Superlinear Convergence of a Trust Region Algorithm for Nonsmooth Optimization*, Mathematical Programming, 31, 1985, pp. 269-285.