



UNIVERSITÉ
University of Namur
DE NAMUR

Institutional Repository - Research Portal Dépôt Institutionnel - Portail de la Recherche

researchportal.unamur.be

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Sensitivity of trust-region algorithms to their parameters

Gould, Nick; Orban, D.; Sartenaer, A.; Toint, Philippe

Published in:
4OR

DOI:
[10.1007/s10288-005-0065-y](https://doi.org/10.1007/s10288-005-0065-y)

Publication date:
2005

Document Version
Peer reviewed version

[Link to publication](#)

Citation for published version (HARVARD):
Gould, N, Orban, D, Sartenaer, A & Toint, P 2005, 'Sensitivity of trust-region algorithms to their parameters', 4OR, vol. 3, no. 3, pp. 227-241. <https://doi.org/10.1007/s10288-005-0065-y>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

4OR manuscript No.
(will be inserted by the editor)

Sensitivity of trust-region algorithms to their parameters

N. Gould¹, D. Orban², A. Sartenaer³ and Ph. L. Toint³

¹ Computational Science and Engineering Departement, Rutherford Appleton Laboratory, Chilton, Oxfordshire, England (email: n.gould@rl.ac.uk)

² Département de Mathématiques et Génie Industriel, Ecole Polytechnique de Montréal, 2900, Bd E. Montpetit, H3T 1J4 Montréal, Canada (email: dominique.orban@polymtl.ca)

³ Department of Mathematics, University of Namur, 61, rue de Bruxelles, B-5000 Namur, Belgium (email: annick.sartenaer@fundp.ac.be, philippe.toint@fundp.ac.be)

Received: date / Revised version: date

Abstract In this paper, we examine the sensitivity of trust-region algorithms on the parameters related to the step acceptance and update of the trust region. We show, in the context of unconstrained programming, that the numerical efficiency of these algorithms can easily be improved by choosing appropriate parameters. Recommended ranges of values for these parameters are exhibited on the basis of extensive numerical tests.

Keywords: unconstrained programming – trust-region methods – algorithmic parameters

1 Introduction

Trust-region methods form a popular class of iterative optimization algorithms, in which the objective function is approximated by a model and this model is minimized in a neighbourhood—the trust region—of the current iterate. Originally proposed in the context of nonlinear least-squares fitting (see Levenberg (1944), Marquardt (1963) and Morrison (1960)), this class of methods was then developed to form a robust theoretical and practical framework containing a number of locally convergent algorithms for smooth unconstrained minimization problems, and, in particular, Newton method (see Goldfeldt et al. (1966), Powell (1970) and Winfield (1973)). The reader is referred to Conn et al. (2000), Moré (1983) and Moré and Sorensen (1984) for additional motivation and convergence analysis.

Like many other algorithms, trust-region methods depend on the choice of a set of parameters, which specify when a trial step is deemed successful and how the trust-region size is updated as the iterations proceed. The purpose of this paper is to present an experimental study of the sensitivity, measured in terms of efficiency, of a trust-region algorithm for smooth unconstrained optimization as a function of these parameters.

We first present the trust-region algorithm in §2 and discuss the values of its parameters. We next describe the setting of our numerical experiments in §3, report on its results from the point of view of efficiency in §4. Some conclusions and perspectives are finally presented in §5.

2 The problem and algorithm

We consider the problem

$$\begin{aligned} & \text{minimize } f(x), \\ & x \in \Re^n \end{aligned} \tag{1}$$

where f is a twice continuously differentiable function from \Re^n into \Re . We assume that the problem is well defined in that f is bounded below. The philosophy of trust-region methods is to calculate, at iterate x_k , a model m_k of the objective function in the *trust region*

$$\mathcal{B}_k = \{x_k + s \mid \|s\| \leq \Delta_k\},$$

where $\|\cdot\|$ is the Euclidean norm on \Re^n and $\Delta_k > 0$ is the *trust-region radius*. A step s_k is then computed that approximately minimizes this model within the trust region. If the value of the objective function computed at the trial point $x_k + s_k$ produces a decrease in the objective function which is comparable to that predicted by the model, the trial point is accepted as the next iterate and the trust-region radius is possibly increased. Otherwise, the trial step is rejected and the radius decreased. More formally, our algorithm is defined as Algorithm 1 on page 3.

A number of choices are possible for the model m_k . Let $\langle \cdot, \cdot \rangle$ denote the usual inner product in \Re^n . In what follows, we focus on a quadratic model of the form

$$m_k(x_k + s) = f(x_k) + \langle \nabla_x f(x_k), s \rangle + \frac{1}{2} \langle s, \nabla_{xx} f(x_k) s \rangle, \tag{4}$$

which is sometimes known as *Newton's model*. Note that the model and objective function coincide up to second order at x_k , i.e., $m_k(x_k) = f(x_k)$, $\nabla_x m_k(x_k) = \nabla_x f(x_k)$ and $\nabla_{xx} m_k(x_k + s) = \nabla_{xx} f(x_k)$ for all s .

We will not expand in full detail on what we mean by “sufficient model reduction” in Step 2. If we define the *Cauchy point* x_k^C to be a model minimizer along the intersection of the steepest descent direction and the trust region, we say that a step s_k produces sufficient decrease if and only if

$$m_k(x_k) - m_k(x_k + s_k) \geq \kappa_{\text{red}} [m_k(x_k) - m_k(x_k^C)], \tag{5}$$

Algorithm 1: The Basic Trust-Region Algorithm

Step 0. [Initialization] An initial point x_0 and an initial trust-region radius Δ_0 are given, as well as the parameters η_1, η_2, α_1 and α_2 that satisfy

$$0 \leq \eta_1 < \eta_2 < 1 \text{ and } 0 < \alpha_1 < 1 < \alpha_2. \quad (2)$$

Compute $f(x_0)$ and set $k = 0$.

Step 1. [Model definition] Define a model $m_k(x_k + s)$ of $f(x_k + s)$ in \mathcal{B}_k .

Step 2. [Step calculation] Compute a step s_k that “sufficiently reduces the model” m_k and such that $x_k + s_k \in \mathcal{B}_k$.

Step 3. [Acceptance of the trial point] Compute $f(x_k + s_k)$ and

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{m_k(x_k) - m_k(x_k + s_k)}. \quad (3)$$

If $\rho_k \geq \eta_1$, then set $x_{k+1} = x_k + s_k$; otherwise, set $x_{k+1} = x_k$.

Step 4. [Trust-region radius update] Set

$$\Delta_{k+1} = \begin{cases} \alpha_1 \|s_k\| & \text{if } \rho_k < \eta_1 \\ \Delta_k & \text{if } \eta_1 \leq \rho_k < \eta_2 \\ \max[\alpha_2 \|s_k\|, \Delta_k] & \text{if } \rho_k \geq \eta_2. \end{cases}$$

Increment k by one, and loop back to Step 1.

for some $\kappa_{\text{red}} \in (0, 1)$. The quantity $m_k(x_k) - m_k(x_k + s_k)$ is referred to as the *predicted decrease*. Condition (5) is known to ensure global convergence of the algorithm to first-order critical points of problem (1) under suitable assumptions (see Conn et al. (2000) and Moré (1983)). In the next section, we return to the choice of an algorithm ensuring (5) when the model (4) is considered.

The algorithm depends on the constants η_1, η_2, α_1 and α_2 , whose values are only restricted to satisfy (2). The values

$$\eta_1 = 0.25, \quad \eta_2 = 0.75 \quad (6)$$

(see, for instance, Coleman and Li (1996), Conn et al. (1988), Lin and Moré (1999), Sartenaer (1993), Sebusdandi and Toint (1993) and Shahabuddin (1996)) and

$$\alpha_1 = 0.5 \quad \text{and} \quad \alpha_2 = 2 \quad (7)$$

(see, for instance, Conn et al. (1988), Dennis et al. (1998), Dennis and Vicente (1996), Sartenaer (1993) and Shahabuddin (1996)) have been used in practical implementations. Is the behaviour of the algorithm relatively insensitive to variations in these values? The purpose of this note is to examine this question by analyzing the performance of the algorithm for a set of alternative parameter values over a reasonable set of test problems.

The choice of η_1 also has a theoretical implication: choosing $\eta_1 > 0$ is known to be necessary in order to guarantee that all limit points of

the sequence of iterates satisfy the first-order optimality conditions (see Moré (1983), Thomas (1975) and Yuan (1998)). It is thus also of interest to investigate the effect of this condition on numerical efficiency.

The authors are aware that, at least the choice of α_1 and α_2 may be made less crucial if an interpolation scheme is used to determine the value of Δ_{k+1} from that of $f(x_k)$, $f(x_k + s_k)$ and $\nabla_x f(x_k)$ (see Conn et al. (1992), Dennis and Schnabel (1996) and Lin and Moré (1999)). However, good values for these parameters remain helpful in simpler implementations of the trust-region method.

3 The framework for the numerical experiments

We now turn to the discussion of the framework in which our numerical experiments are conducted. Obviously, such an experimental investigation is never perfect, and several of the choices made here, although reasonable from our point of view, are not the only ones that one could consider. We are well aware of the limitations of our approach. We briefly discuss them and propose directions for additional research in §5.

A first decision involves selecting a suitable performance measure for the algorithm. We have chosen to measure algorithm efficiency by CPU time and number of iterations to obtain convergence, which is declared as soon as

$$\|\nabla_x f(x_k)\| \leq 10^{-5}. \quad (8)$$

(We are aware that relative/weighted tests might be more useful for badly-scaled/nonlinear problems.) We had intended to declare failure when (8) was not met in the first 1000 iterations of the algorithm, but this situation never happened in our tests. Note that since we are only considering unconstrained minimization, this measure in terms of iteration counts is equivalent to considering the number of function evaluations and is justified in the frequent case where objective function evaluations are computationally costly and dominate the internal work of the algorithm. If evaluating the objective function is relatively cheap compared to the algorithm's internal work, which might be the case when the dimension of the problem is large and the linear algebra therefore more expensive, then overall CPU time is the obvious choice for measuring performance.

The second important choice is that of the procedure to compute, at each iteration, a trial step s_k that approximately minimizes the model within the trust region. We have chosen to use the truncated conjugate-gradient iteration, or Steihaug-Toint algorithm (see Steihaug (1983) and Toint (1981)), as implemented in the GLTR module of the GALAHAD library (see Gould et al. (1999) and Gould et al. (2003b)). In this method, the iterates generated by the conjugate-gradient algorithm are used until either they leave the trust

region or negative curvature is discovered—in either of these cases the last conjugate-gradient step is truncated on the trust-region boundary. Thus, the Steihaug-Toint procedure terminates at $x_k + s$ that satisfies

$$\|\nabla_x f(x_k + s)\| \leq \min \left[\frac{1}{10}, \|\nabla_x f(x_k)\|^{1/2} \right] \quad \text{or} \quad \|s\| = \Delta_k.$$

This procedure is known to guarantee (5).

We also decided to follow LANCELOT (see Conn et al. (1992)) in using the value $\Delta_0 = 0.1\|\nabla_x f(x_0)\|$ for the initial trust-region radius. We did not perform a sensitivity analysis on this parameter because this issue has already been considered in Sartenaer (1997).

We have chosen to test the algorithm on a set of 24 problems from the CUTER collection of Gould et al. (2003a), whose average dimension is around 1000. All problems are unconstrained, with some of them being highly nonlinear. These problems have proved to be reasonably hard¹ in the past and are, in our opinion and despite their fairly uniform dimensions, reasonably representative of the unconstrained part of the collection as a whole. Their name and size can be found in Table 1. The number of problems in our test set had to be kept relatively low in order to make the test for a large number of parameter values tractable.

problem name	dimension	problem name	dimension
BIGGSB1	1000	NCB20B	1000
CURLY10	1000	NONCVXU2	1000
CURLY20	1000	NONDIA	1000
CURLY30	1000	NONDQUAR	1000
EDENSCH	2000	PENALTY1	1000
EIGENBLS	1056	POWER	1000
FREUROTH	1000	QUARTC	1000
GENROSE	1000	SINQUAD	1000
LINVERSE	999	SPARSINE	1000
MSQRTALS	1024	SPMSRTLS	1000
MSQRTBLS	1024	VAREIGVL	999
NCB20	1000	WOODS	1000

Table 1. The test problems set from the CUTER collection

Finally, we decided on a set of parameter values to be experimented with. In a first stage, we chose to let (η_1, η_2) vary in the (coarse) uniform grid

$$G_\eta^C = \left\{ (\eta_1, \eta_2) \left| \begin{array}{l} \eta_1 \in \{0, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45\}, \\ \eta_2 \in \{0.5, 0.6, 0.7, 0.8, 0.9\} \end{array} \right. \right\},$$

¹ On easy problems, the influence of the parameters is less prominent.

while (α_1, α_2) varies in the grid

$$G_\alpha = \left\{ (\alpha_1, \alpha_2) \left| \begin{array}{l} \alpha_1 \in \{0.25, 0.33, 0.5, 0.66, 0.75\}, \\ \alpha_2 \in \{1.5, 2.0, 2.5, 3.0, 3.5, 5\} \end{array} \right. \right\}.$$

Preliminary experiments using the grids G_η^C and G_α seemed to indicate that the region

$$\mathcal{Z} = \{0 \leq \eta_1 \leq 0.1, \quad 0.7 \leq \eta_2 < 1\}$$

was worth being discretized more finely as it appeared to be where the best efficiency would occur. As a result, the coarse grid G_η^C was replaced by the finer, but no longer uniform, grid

$$G_\eta^F = \left\{ (\eta_1, \eta_2) \left| \begin{array}{l} \eta_1 \in \{0, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 0.1, 0.15, 0.2, \\ 0.25, 0.3, 0.4\} \\ \eta_2 \in \{0.5, 0.6, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 0.99, 0.995, \\ 0.999\} \end{array} \right. \right\}.$$

The grid G_α remained unchanged. Altogether, this gives a set of 3,960 different values for $(\eta_1, \eta_2, \alpha_1, \alpha_2)$. We then applied the algorithm using these values on the 24 test problems, which resulted in a grand total of 95,040 test runs. All tests were performed in double precision Fortran 90 on an 1.6GHz Intel Pentium IV Linux PC with 512 MBytes of RAM.

4 Numerical results and analysis

4.1 Combined performance

Remarkably, optimizing the trust-region algorithm parameters seems to depend only marginally on whether one aims for a minimum number of iterations (and thus of function/derivatives calculations) or for a minimum CPU time. The performance measured in these two quantities seems indeed very related, as is clear from Figure 1, where the couples consisting of the average iteration count and the average CPU time (the averages being taken on all test problems) are plotted for each of the 3,960 tested variants. In this figure, the performance of the “standard” parameter choice, as defined by (6) and (7), is identified as the point at the intersection of the vertical and horizontal cross lines.

The “comet-like” structure of the cloud of performance couples is very elongated for our test set. Moreover, the densest “core” part of the cloud appears to correspond to relatively efficient variants, the worst ones being quite far in the sparser “tail”. The standard variant belongs to the core, but




Fig. 1. Combined performance of all 3,960 tested variants (standard parameter choice marked by crosshair)




Fig. 2. Combined performance of the more efficient tested variants

just barely, which shows that many significantly better parameter choices exist.

Figure 2 is obtained by zooming on the tip of the core, i.e., the part of the figure containing the most efficient variants. This figure shows several

vertical clusters of points corresponding to variants using exactly the same average number of iterations and slightly different CPU times. These clusters reflect the inaccuracy of the timing routine used (a few percent), and should therefore be interpreted as a single combined performance measure.

We now explore in more details the clusters of the efficient frontier of the above figures. This frontier is defined as the set of clusters that are not dominated by any other cluster, that is the clusters such that there is no other cluster in the figure giving better average number of iterations *and* better average CPU time.

The first cluster on the left corresponds to the best variant in terms of iterations. The associated performance is presented in Table 2, which shows the associated parameter values, the total number of iterations and the CPU time for this cluster, for the standard choice of the parameters, and for the worst choice. Note that the CPU time in brackets is a rounded average of the times associated with all the variants in the entire cluster.

	η_1	η_2	α_1	α_2	# its	CPU
Parameter choices for first cluster	$[0, 10^{-2}]$	0.99	0.25	3.5	14.750	(9.45)
Standard choice	0.25	0.75	0.5	2	20.625	17.13
Worst choice	0.4	0.5	0.75	5	47.000	63.90

Table 2. Average iteration counts and CPU times (s) for the best, standard and worst parameter choices in terms of iterations

We see that the best parameter choice results in an algorithm that requires, on average, $14.750/20.625 = 0.715$ as many iterations as the standard one. On the other hand, the worst choice produces a method that needs, on average, $47/14.75 = 3.19$ as many iterations as the best, and $47/20.625 = 2.28$ as many as the standard. The proportions in time reinforce these trends.

Another outcome of our tests is that the performance of the variants cannot be distinguished if they differ only by the choice of η_1 in the range 0 to 10^{-4} . For most choices of the other parameters, this is also the case for η_1 up to 10^{-2} . This is interesting because it seems to imply that a small strictly positive value of η_1 is as good as possible from the numerical efficiency point of view, while, at the same time, ensuring the best theoretical convergence properties. Conversely, this indicates that choosing $\eta_1 = 0$ may not have a numerically detrimental effect despite being theoretically less satisfactory. This behaviour is also apparent in Table 2.

Pursuing our exploration of the efficient frontier in Figure 2, the next dominant cluster is that corresponding to variants taking on average 15 iterations. It is in fact made up from two similar but distinct choices of parameters, as indicated in Table 3. Strictly speaking, only the first of those

is really part of the efficient frontier, since it marginally dominates the second.

	η_1	η_2	α_1	α_2	# its	CPU
Parameter choices for middle cluster	$[0, 10^{-1}]$	$\{0.995, 0.999\}$	0.25	5	15.000	(9.15)
	$[0, 10^{-4}]$	0.999	0.33	3.5	15.000	(9.35)
Standard choice	0.25	0.75	0.5	2	20.625	17.13
Worst choice	0.4	0.5	0.75	5	47.000	63.90

Table 3. Average iteration counts and CPU times (s) for the middle dominant cluster of variants, and for the standard and worst ones

We can observe that algorithms using these choices of parameters require, on average, $15/20.625=0.727$ as many iterations as the standard choice and $15/47=0.319$ as many as the worst. Again, the ratios for CPU time confirm the trends observed for iterations.

The third and last point on the efficient frontier is given by the variant which is best in terms of CPU time (i.e., is lowest in Figure 2). Table 4 shows the (unique) set of parameter values and total CPU time for this variant, and compares them again with the standard and worst choices.

	η_1	η_2	α_1	α_2	# its	CPU
Parameter choice for third cluster	0.15	0.999	0.33	5	15.25	8.76
Standard choice	0.25	0.75	0.5	2	20.625	17.13
Worst choice	0.4	0.5	0.75	5	47.000	63.90

Table 4. Average iteration counts and CPU times (s) for the best, standard and worst parameter choices in terms of CPU time

Following the same logic as above, we note that this parameter choice results in an algorithm that is, on average, $17.13/8.76 = 1.96$ as fast as the standard one. On the other hand, the worst choice produces a method that is, on average, $63.90/8.76 = 7.3$ times slower than the best, and $63.90/17.13 = 3.73$ slower than the standard. We note that this best variant in CPU time is close to those in the middle cluster, except for a surprisingly larger value of η_1 . We might thus see it as the result of a particularly lucky set of time measures.

4.2 Iterations sensitivity

We now examine the sensitivity in iterations of the above conclusions for the parameter choices of the first cluster (most efficient in terms of iterations).

η_2	0.50	0.60	0.70	0.75	0.80	0.85	0.90	0.95	0.99	0.995	0.999
η_1											
0.00	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
10^{-6}	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
10^{-5}	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
10^{-4}	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
10^{-3}	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
10^{-2}	0.75	0.79	0.80	0.78	0.75	0.76	0.82	0.83	1.00	0.89	0.94
0.10	0.69	0.78	0.75	0.80	0.80	0.78	0.86	0.84	0.91	0.83	0.87
0.15	0.72	0.72	0.75	0.77	0.80	0.74	0.76	0.78	0.95	0.86	0.90
0.20	0.74	0.70	0.74	0.75	0.73	0.77	0.81	0.85	0.92	0.90	0.91
0.25	0.70	0.72	0.70	0.78	0.72	0.76	0.84	0.77	0.90	0.86	0.87
0.30	0.70	0.66	0.72	0.78	0.71	0.76	0.76	0.78	0.91	0.90	0.89
0.40	0.71	0.72	0.72	0.74	0.73	0.77	0.80	0.80	0.88	0.87	0.86

Table 5. Relative performance in iterations as a function of η_1 and η_2 for fixed $\alpha_1 = 0.25$ and $\alpha_2 = 3.5$

Table 5 presents the relative performance, compared to the best, of all considered choices of η_1 and η_2 when α_1 and α_2 are fixed to their optimal value (see Table 2). We note in this table that, as indicated above, the results are identical for all values of η_1 that are between 0 and 0.01. This indicates that small values of this parameter are best, but also that the precise choice of a small value of η_1 is not crucial. A second observation is that a choice of η_2 larger than 0.99 is comparatively better than that of a smaller value. Globally the performance degrades nearly monotonically when η_1 grows and/or η_2 decreases.

The relative performance of all considered choices of α_1 and α_2 (for values of η_1 and η_2 fixed to one of their optimal values) are shown in Table 6. Again, the neighbourhood of the best values is clear. One notices that the optimal values for α_2 , around 3.5, are considerably larger than their standard (2). It therefore seems more efficient to increase the radius relatively unfrequently (η_2 close to 1) but to do so more decisively when it happens. This is supported by the observation that the choice $\alpha_2 = 5$ remains very satisfactory for $\alpha_1 = 0.25$ and $\alpha_1 = 0.33$ (see also Tables 3 and 4). The best value $\alpha_1 = 0.25$ shows that unsuccessful iterations ($\rho_k < \eta_1$), although not frequent, should not reduce the trust-region radius too aggressively.

We conclude our sensitivity analysis in terms of number of iterations by considering the performance profiles proposed by Dolan and Moré (2002) for the worst, standard and best parameter choices. Such profiles are defined as follows. Assume a certain set \mathcal{A} of competing algorithms $\mathcal{A}_1, \dots, \mathcal{A}_q$ is tested on a set \mathcal{S} of p test problems and algorithm \mathcal{A}_i reports a certain measure of performance $\pi_{i,t}$ (iteration counts, in our case) when run on test problem t such that Algorithm \mathcal{A}_i is “better” than Algorithm \mathcal{A}_j on this

α_2 α_1	1.50	2.00	2.50	3.00	3.50	5.00
0.10	0.79	0.87	0.90	0.87	0.94	0.92
0.25	0.79	0.81	0.90	0.93	1.00	0.95
0.33	0.83	0.92	0.88	0.90	0.93	0.94
0.50	0.83	0.86	0.85	0.85	0.79	0.93
0.75	0.71	0.71	0.74	0.72	0.72	0.68

Table 6. Relative performance in iterations as a function of α_1 and α_2 for fixed $\eta_1 = 10^{-5}$ and $\eta_2 = 0.99$

problem t if $\pi_{i,t} \leq \pi_{j,t}$. If we define the *indicator* function

$$\kappa(\pi_1, \pi_2; \sigma) \equiv \begin{cases} 1 & \text{if } \pi_1 \leq \sigma \pi_2, \\ 0 & \text{otherwise,} \end{cases}$$

the *performance profile* of Algorithm \mathcal{A}_i on the test set \mathcal{S} with respect to the performance measure π is defined as the function

$$\pi_i(\sigma) \equiv \frac{\sum_{t=1}^p \kappa(\pi_{i,t}, \pi_{\min,t}; \sigma)}{p}, \quad \text{for every } \sigma > 0,$$

where $\pi_{\min,t} = \min_{i=1,\dots,q} \pi_{i,t}$ is the best performance achieved on problem t . Hence, $\pi_i(\sigma)$ is a measure of the probability for the performance of Algorithm \mathcal{A}_i to be within a factor σ of the performance of the best algorithm. For instance, $\pi_i(1)$ measures the probability for the performance of Algorithm \mathcal{A}_i to be the best, and $\lim_{\sigma \rightarrow \infty} \pi_i(\sigma)$ is the probability that Algorithm \mathcal{A}_i solves a problem.

The iteration performance profiles for the algorithms corresponding to the best, standard and worst parameter choices are presented in Figure 3. Interestingly, the standard choice is less often the best ($\sigma = 1$) than the worst and best.

4.3 CPU time sensitivity

We now turn to sensitivity in CPU time and analyze the third cluster of the efficient frontier, which is the most efficient for this criterion.

Table 7 presents the relative performance in CPU time of all considered choices of η_1 and η_2 when α_1 and α_2 are fixed to their optimal value (see Table 4). Although the best performance is obtained for $\eta_1 = 0.15$ and $\eta_2 = 0.999$, we see that the choices of η_1 between 0 and 0.1 for $\eta_2 = 0.99$ and 0.995 remain highly efficient. The general conclusions in terms of CPU time are thus very consistent with those obtained when considering iteration counts: small values of η_1 are best and a choice of η_2 larger than 0.99 is comparatively better than that of a smaller value.




Fig. 3. Iteration performance profiles for the best, standard and worst parameter choices

η_2	0.50	0.60	0.70	0.75	0.80	0.85	0.90	0.95	0.99	0.995	0.999
η_1	0.47	0.44	0.46	0.47	0.51	0.63	0.72	0.86	0.97	0.97	0.89
0.00	0.46	0.44	0.46	0.47	0.51	0.64	0.72	0.86	0.98	0.97	0.89
10^{-6}	0.46	0.44	0.45	0.47	0.51	0.63	0.72	0.86	0.96	0.97	0.88
10^{-5}	0.46	0.44	0.45	0.47	0.51	0.62	0.72	0.86	0.96	0.97	0.88
10^{-4}	0.45	0.43	0.45	0.47	0.51	0.62	0.72	0.85	0.96	0.96	0.87
10^{-3}	0.45	0.43	0.45	0.46	0.50	0.61	0.70	0.85	0.96	0.97	0.87
10^{-2}	0.45	0.43	0.45	0.55	0.51	0.62	0.71	0.91	0.96	0.96	0.87
0.10	0.32	0.40	0.43	0.44	0.57	0.60	0.79	0.91	0.97	0.92	0.97
0.15	0.28	0.35	0.41	0.43	0.54	0.50	0.74	0.92	0.93	0.96	1.00
0.20	0.31	0.34	0.42	0.44	0.48	0.52	0.82	0.73	0.85	0.86	0.86
0.25	0.32	0.39	0.39	0.44	0.44	0.49	0.76	0.73	0.87	0.91	0.91
0.30	0.30	0.34	0.41	0.42	0.45	0.53	0.74	0.72	0.90	0.85	0.85
0.40	0.28	0.34	0.39	0.37	0.48	0.50	0.66	0.72	0.81	0.80	0.80

Table 7. Relative performance in CPU time as a function of η_1 and η_2 for fixed $\alpha_1 = 0.33$ and $\alpha_2 = 5$

The relative CPU time performance of all considered choices of α_1 and α_2 (for values of η_1 and η_2 fixed to their optimum) are finally shown in Table 8. Again, the conclusions follow the lines stated above: values of α_1 should be moderate while values of α_2 should definitely exceed the standard choice of 2.

Finally, Figure 4 shows the CPU-time performance profiles for the best, standard and worst parameter choices. Again the profile for the best choice indicates a clear superiority of this variant on the other two.

α_2 α_1	1.50	2.00	2.50	3.00	3.50	5.00
0.10	0.67	0.77	0.87	0.92	0.74	0.90
0.25	0.79	0.78	0.87	0.79	0.82	0.94
0.33	0.75	0.89	0.88	0.78	0.79	1.00
0.50	0.72	0.78	0.67	0.71	0.73	0.69
0.75	0.58	0.66	0.56	0.47	0.49	0.54

Table 8. Relative performance in CPU time as a function of α_1 and α_2 for fixed $\eta_1 = 0.15$ and $\eta_2 = 0.999$




Fig. 4. CPU time performance profiles for the best, standard and worst parameter choices

4.4 Tentative recommendations

Recommending parameter values in general is always somewhat risky. The actual performance of an algorithm may indeed strongly depend on the class of problems on which it is applied, and the conclusions of this study rely on a particular test-problem set. The values suggested here are therefore best seen as appropriate for general purpose use and are not intended to replace application dependent choices. This said, it seems that a reasonable parameter choice could be

$$\eta_1 = 0.0001, \quad \eta_2 = 0.99, \quad \alpha_1 = 0.25, \quad \alpha_2 = 3.5.$$

If distinguishing between the two measures of efficiency is critical, the values suggested in Sections 4.2 and 4.3 might be considered, but direct experience remains of course the best source of inspiration.

5 Conclusions

We have performed simple, systematic, tests with a basic trust-region algorithm for unconstrained programming. The test problems are taken from the CUTER collection and our algorithm underwent tests for nearly 4000 values of the trust-region parameters. The commonly used “standard” values for these parameters appear not to be the best choice and alternative values have been pointed out, emphasizing the gain that these could produce over the standard values.

Only simple, multiplicative, updating rules have been used in this algorithmic framework. The authors are well aware that other rules could have been tested and that a much finer grid could have been used. The values of the α parameters could also be determined dynamically using a second or third degree polynomial interpolation of the objective function, or of the ρ function. While the conclusions drawn here therefore remain tentative and dependent on a particular set of test problems, the authors believe that they may be of interest for algorithm developers.

References

- Coleman TF, Li Y (1996) An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization* 6(2):418–445
- Conn AR, Gould N, Toint PhL (1988) Testing a class of methods for solving minimization problems with simple bounds on the variables. *Mathematics of Computation* 50:399–430
- Conn AR, Gould N, Toint PhL (1992) *LANCELOT: a Fortran package for large-scale nonlinear optimization (Release A)*. Number 17 in Springer Series in Computational Mathematics. Springer Verlag, Heidelberg, Berlin, New York
- Conn AR, Gould N, Toint PhL (2000) *Trust-Region Methods*. Number 01 in MPS-SIAM Series on Optimization. SIAM, Philadelphia, USA
- Dennis JE, Heinkenschloss M, Vicente LN (1998) Trust-region interior-point SQP algorithms for a class of nonlinear programming problems. *SIAM Journal on Control and Optimization* 36(5):1750–1794
- Dennis JE, Schnabel RB (1996) *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, Englewood Cliffs, New Jersey, USA, 1983. Reprinted as *Classics in Applied Mathematics* 16, SIAM, Philadelphia, USA
- Dennis JE, Vicente LN (1996) Trust region interior-point algorithms for minimization problems with simple bounds. In H. Fisher, B. Riedmüller, and S. Schäffler, editors, *Applied Mathematics and Parallel Computing, Festschrift for Klaus Ritter*, pages 97–107, Heidelberg, Berlin, New York. Physica-Verlag, Springer-Verlag
- Dolan ED, Moré JJ (2002) Benchmarking optimization software with performance profiles. *Mathematical Programming* 91(2):201–213
- Goldfeldt SM, Quandt RE, Trotter HF (1966) Maximization by quadratic hill-climbing. *Econometrica* 34:541–551
- Gould N, Lucidi S, Roma M, Toint PhL (1999) Solving the trust-region subproblem using the Lanczos method. *SIAM Journal on Optimization* 9(2):504–525
- Gould N, Orban D, Toint PhL (2003a) CUTER, a constrained and unconstrained testing environment, revisited. *ACM Transactions on Mathematical Software* 29(4):373–394
- Gould N, Orban D, Toint PhL (2003b) GALAHAD—a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization. *ACM Transactions on Mathematical Software* 29(4):353–372

- Levenberg K (1944) A method for the solution of certain problems in least squares. *Quarterly Journal on Applied Mathematics* 2:164–168
- Lin C, Moré JJ (1999) Newton's method for large bound-constrained optimization problems. *SIAM Journal on Optimization* 9(4):1100–1127
- Marquardt D (1963) An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics* 11:431–441
- Moré JJ (1983) Recent developments in algorithms and software for trust region methods. In A. Bachem, M. Grötschel, and B. Korte, editors, *Mathematical Programming: The State of the Art*, pages 258–287, Heidelberg, Berlin, New York, Springer Verlag
- Moré JJ, Sorensen DC (1984) Newton's method. In G. H. Golub, editor, *Studies in Numerical Analysis*, number 24 in MAA Studies in Mathematics, pages 29–82, Providence, Rhode-Island, USA, American Mathematical Society
- Morrison DD (1960) Methods for nonlinear least squares problems and convergence proofs. In J. Lorell and F. Yagi, editors, *Proceedings of the Seminar on Tracking Programs and Orbit Determination*, pages 1–9, Pasadena, USA, Jet Propulsion Laboratory
- Powell MJD (1970) A new algorithm for unconstrained optimization. In J. B. Rosen, O. L. Mangasarian, and K. Ritter, editors, *Nonlinear Programming*, pages 31–65, London, Academic Press
- Sartenaer A (1993) Armijo-type condition for the determination of a generalized Cauchy point in trust region algorithms using exact or inexact projections on convex constraints. *Belgian Journal of Operations Research, Statistics and Computer Science* 33(4):61–75
- Sartenaer A (1997) Automatic determination of an initial trust region in nonlinear programming. *SIAM Journal on Scientific Computing* 18(6):1788–1803
- Sebudandi Ch, Toint PhL (1993) Nonlinear optimization for seismic travel time tomography. *Geophysical Journal International* 115:929–940
- Shahabuddin JS (1996) *Structured trust-region algorithms for the minimization of nonlinear functions*. PhD thesis, Department of Computer Science, Cornell University, Ithaca, New York, USA
- Steihaug T (1983) The conjugate gradient method and trust regions in large scale optimization. *SIAM Journal on Numerical Analysis* 20(3):626–637
- Thomas S (1975) *Sequential estimation techniques for quasi-Newton algorithms*. PhD thesis, Cornell University, Ithaca, New York, USA
- Toint PhL (1981) Towards an efficient sparsity exploiting Newton method for minimization. In I. S. Duff, editor, *Sparse Matrices and Their Uses*, pages 57–88, London, Academic Press
- Winfield D (1973) Function minimization by interpolation in a data table. *Journal of the Institute of Mathematics and its Applications* 12:339–347
- Yuan Y (1998) An example of non-convergence of trust region algorithms. In Y. Yuan, editor, *Advances in Nonlinear Programming*, pages 205–218, Dordrecht, The Netherlands, Kluwer Academic Publishers