Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 195 (2009) 31–45
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Continuous Optimization

Finding local optima of high-dimensional functions using direct search methods

Lars Magnus Hvattum [a,*], Fred Glover [b]

[a] Norwegian University of Science and Technology, Trondheim, Norway
[b] University of Colorado, Boulder, CO, USA

Received 11 June 2007; accepted 27 January 2008
Available online 5 February 2008

Abstract

This paper focuses on a subclass of box-constrained, non-linear optimization problems. We are particularly concerned with settings where gradient information is unreliable, or too costly to calculate, and the function evaluations themselves are very costly. This encourages the use of derivative free optimization methods, and especially a subclass of these referred to as direct search methods. The thrust of our investigation is twofold. First, we implement and evaluate a number of traditional direct search methods according to the premise that they should be suitable as local optimizers when used in a metaheuristic framework. Second, we introduce a new direct search method, based on Scatter Search, designed to remedy the lack of a good derivative free method for solving problems of high dimensions. Our new direct search method has convergence properties comparable to those of existing methods in addition to being able to solve larger problems more effectively.

Keywords: Non-linear optimization; Local minimum; Derivative free; Direct search; Scatter Search

1. Introduction

The important problem of minimizing non-linear functions can be handled in a large number of ways. Typically, techniques based on Newton's method can be applied successfully, by using gradient and Hessian information to calculate a good step and gradually move towards a (local) optimum of the function being minimized ([27] is a good introduction to these methods). However, sometimes Newton based approaches can be the wrong choice. This can be the case if: (1) the function evaluations are inaccurate, (2) the derivatives of the function are unavailable or unreliable, or (3) the function is not smooth [36,39]. In these cases, a better choice can be to rely on so-called derivative free methods, i.e., methods that do not explicitly use derivatives of the function being optimized.

Derivative free methods can be divided in two groups. One group includes methods that, instead of using the gradient directly, approximates the derivatives by building a model of the function based on function evaluations [6]. This requires that numerical function values are available, and the task of building a representative model can be difficult if the function evaluations are noisy. The second group of methods includes direct search methods, which are methods that do not try to use gradient information at all (i.e., they do not make any attempt at approximating the gradient) and that only require ordinal information about function values [22,36]. Direct search methods are hence deemed suitable for problems involving simulation-based optimization or optimizing non-numerical functions, as well as, in practice, problems involving non-smooth or discontinuous functions [19].

* Corresponding author. Tel.: +47 73593187.
E-mail addresses: lars.m.hvattum@iot.ntnu.no (L.M. Hvattum), fred.glover@colorado.edu (F. Glover).

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2008.01.039

In this paper, the focus is on the use of direct search methods to find local optima. Many direct search methods are simple in nature, and only undertake to find a local optimum (close to the starting point of the search). Thus, they are often used in hybrid methods for global optimization, e.g., as in [4,5,13–15,30]. However, quite often the choice of which direct search method to use is left unjustified, or, else only a few alternative methods are examined. At the same time, one of the most well-known and extensively used methods, the Nelder–Mead [26], can perform remarkably poorly both from a theoretical and practical point of view [12,22].

One issue that is often mentioned in discussions about direct search methods (as well as other derivative free methods), is their tendency to perform well only on problems with few dimensions and to falter when the number of dimensions grows [1,6,17,19,36]. This motivates the work presented here, in which different direct search methods (as well as two other derivative free methods) are evaluated based on their ability to quickly find local optima of high-dimensional functions.

All test problems used in the computational experiments presented later in this paper can be seen as instances of the following model:

$$\min_{x \in [\ell,u] \subseteq \Re^n} f(x), \tag{1}$$

where $f : \Re^n \rightarrow \Re$, and, $\ell, u \in \Re^n$. Strictly speaking, however, in accordance with the definition of direct search methods proposed in [22], several variations of this problem can be tackled. For example, the objective function need not be numerical (i.e., one could have $f(x) : \Re^n \rightarrow \mathrm{R}$, seeking to determine x to minimize $|\{y : f(y) \leqslant f(x)\}|$, where $\leqslant$ defines an order on R) In addition, the functions that are optimized can contain noise. None of the solution methods examined in this work are specifically adapted to exploit box-constraints ($x \in [\ell, u]$), although we make reference to such constraints as a convenient foundation for generating an initial solution point. For the purpose of this paper, the formulation (1) will thus suffice, even though, unless otherwise stated, the methods discussed are applicable in the wider context mentioned above.

The remainder of this paper is as follows. In Section 2 we examine several traditional and some lesser known direct search methods (as well as two other derivative free methods – one mainly because it is claimed to be effective for problems with many variables). Then, in Section 3 we develop a direct search based on ideas from Scatter Search [9,20] and conduct computational tests that demonstrate the efficacy of our method in this setting. Finally, conclusions and suggestions for future research are discussed in Section 4.

2. Traditional direct search methods

Our goal in this section is to examine and evaluate eight different methods (see Table 1) according to their ability to find local optima of functions having many dimensions. The motivation for this is that these methods are frequently used as subsolvers in various global optimization methods, and their efficiency in finding local optima quickly is of interest to practitioners that implement such global methods. Table 2 gives some references for each method (their original source and/or other references where the method is used or discussed), and their classification within the realm of direct search methods. See [22] for one suggested taxonomy of direct search methods. The parameters defined for these methods are given the standard values found in the literature. Note that two of the methods, HPS and SPSA, are not labeled as direct search methods since they require a numerical objective function, and are thus less general than the other methods.

For some of these solution methods there exists a theory of convergence. Early work can be found in [38], where the author proves that the limit of the infimum of the norm of the gradient at the best vertex at iteration k converges to 0 as $k \rightarrow \infty$ for MDS, CS, and HJ, given that the function is continuously differentiable. Although this does not apply to the functions for which direct search methods are the preferred approach [19], the assumption of continuous differentiability sometimes holds in restricted areas of the function domain at hand. In addition, convergence analysis of these methods go much further, and a hierarchy of results based on the degree of smoothness can be found in [2]. At the other end of the scale some of the direct search methods have negative convergence results, such as NM, [12]. In this paper, the emphasis

Table 1
Direct search and other methods from the literature

Method	Full name
NM	Nelder–Mead
MDS	Multi-directional search
CS	Compass search
HJ	Hooke and Jeeves
ROS	Rosenbrock's algorithm (with improvements by Palmer)
SW	Solis and Wets' algorithm
HPS	Heuristic pattern search
SPSA	Simultaneous perturbation stochastic approximation