Journal of Computational and Applied Mathematics 325 (2017) 97–112

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Hybrid genetic deflated Newton method for global optimisation

Marcus M. Noack[a,b,c,*], Simon W. Funke[b,c]

[a] Kalkulo AS, P.O. Box 134, 1325 Lysaker, Norway

[b] Simula Research Laboratory, P.O. Box 134, 1325 Lysaker, Norway

[c] Department of Informatics, University of Oslo, Gaustadalleen 23 B, 0373 Oslo, Norway

HIGHLIGHTS

- New hybrid optimisation scheme.
- Optima are removed from objective function by deflation.
- Resulting algorithm combines strength of local and global optimisation schemes.
- Algorithm finds multiple solutions if they exist.
- Algorithm outperforms genetic algorithms and traditional hybrid methods.

ARTICLE INFO

Article history:
Received 18 August 2016
Received in revised form 23 January 2017

Keywords:
Hybrid optimisation
Deflated Newton
Genetic algorithm

ABSTRACT

Optimisation is a basic principle of nature and has a vast variety of applications in research and industry. There is a plurality of different optimisation procedures which exhibit different strengths and weaknesses in computational efficiency and probability of finding the global optimum. Most methods offer a trade-off between these two aspects. This paper proposes a hybrid genetic deflated Newton (HGDN) method to find local and global optima more efficiently than competing methods. The proposed method is a hybrid algorithm which uses a genetic algorithm to explore the parameter domain for regions containing local minima, and a deflated Newton algorithm to efficiently find their exact locations. In each iteration, identified minima are removed using deflation, so that they cannot be found again. The genetic algorithm is adapted as follows: every individual of every generation of offspring searches its adjacent space for optima using Newton's method; when found, the optimum is removed by deflation, and the individual returns to its starting position. This procedure is repeated until no more optima can be found. The deflation step ensures that the same optimum is not found twice. In the subsequent genetic step, a new generation of offspring is created, using procreation of the fittest. We demonstrate that the proposed method converges to the global optimum, even for small populations. Furthermore, the numerical results show that the HGDN method can improve the number of necessary function and derivative evaluations by orders of magnitude.

* Corresponding author at: Kalkulo AS, P.O. Box 134, 1325 Lysaker, Norway.
E-mail addresses: mcn@simula.no (M.M. Noack), simon@simula.no (S.W. Funke).

http://dx.doi.org/10.1016/j.cam.2017.04.047

1. Introduction

Optimisation is one of the most fundamental principles of nature. Most physical principles can be formulated in the structure of an optimisation problem. Additionally, inversions, like seismic tomography and weather predictions, are typical optimisation problems. It is therefore important to develop efficient methods to optimise functions. This paper is concerned with the problem of finding the local and global minimisers $\{\mathbf{x}^*\}$ of a real-valued function $f : \mathbb{R}^n \to \mathbb{R}$. More precisely, we are seeking points $\mathbf{x}^* \in \mathbb{R}^n$ for which the optimality condition

$$f(\mathbf{x}^*) \leq f(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{x}^*\| < r \tag{1}$$

holds for a sufficiently small $r > 0$. The search for maxima is analogous and will be treated accordingly. The objective function f can be highly non-linear, but we assume that it is continuous and at least twice differentiable. In many practical applications, evaluating f or its derivatives involves computationally expensive operations, such as the solution of a discretised partial differential equation. Therefore, it is crucial to solve the problem with as few functional and derivative evaluations as possible.

Solving problem (1) is numerically challenging, because f can have multiple local and global optima. Local knowledge about the function, such as evaluations and derivatives, is therefore not sufficient to find the global solution or to identify whether an optimum is a local or a global optimum [1]. Hence, existing local optimisation methods cannot be applied directly. Instead, a solution strategy must explore the global parameter space. Genetic algorithms and simulated annealing are popular methods that use randomised search strategies motivated from natural processes. They are robust, find the proximity of the optimal solutions in a reasonable time for a small number of dimensions, are parallelisable and easy to implement [2,3]. Furthermore, they have little assumptions on the objective function f. However, they require many function evaluations, especially in high dimensional spaces. To improve the efficiency, hybrid schemes have been proposed which combine the efficiency of local optimisation methods with the generality of global methods [4–7]. Renders and Flasse [7] in particular showed that hybrid methods can offer a significant improvement compared to genetic algorithms. We are referring to these hybrid methods as traditional hybrid methods in the course of the paper.

This paper presents a new hybrid optimisation method that combines a genetic algorithm with a fast, local optimisation method. The algorithm is based on two basic components: a global search method based on the genetic algorithm, and a local search method. For the local search, we employ a deflated Newton scheme [8]. The deflated Newton method efficiently identifies multiple local minima or maxima in proximity of the starting point, and deflates the function accordingly. As a result, a smaller population size is sufficient to efficiently map the local and global optima of f, which we show, can result in a significant performance increase of the overall algorithm. The key to the success of the deflated Newton method is that the found optima are "removed", meaning that a deflation is placed where the optimum was located. A subsequent Newton search will not converge to the same point, but will find another optimum or diverge, meaning that there are no optima in the vicinity of the individual. This leads to a performance gain of the overall algorithm compared to traditional hybrid methods. The following genetic algorithm will relocate the individuals by using procreation of the fittest. In the new locations, all offspring individuals will again start the search for optima. The proposed method is easy to implement because existing implementations of genetic and local optimisation methods can be reused. The overall goal is to minimise the required function and derivative evaluations to find local and global optima of a function.

The remainder of the paper is organised as follows. In Section 2, the ingredients of the proposed method are mentioned and explained briefly. The following sections will give some information about global and local optimisation schemes. Next, it is explained how these methods work together to form the basis of the proposed hybrid method. Afterwards, the method of deflation is described and used to improve the existing hybrid methods. The proposed method was applied in several standard problems and benchmarked against genetic and traditional hybrid optimisation methods. The numerical results are shown in Section 3.

2. Methodology and theory

Two classes of methods exist when it comes to optimising functions: local methods and global methods. In most fields, the use of either local or global methods means a trade-off between computing time and probability of finding the global optimum. This section will give an introduction to local and global optimisation schemes and will use them to draw the path to the proposed method.

2.1. Global methods

Global methods, like the Monte-Carlo method or the genetic algorithm are randomised algorithms and can guarantee to find the global optimum [9]. For our purposes the genetic algorithm is particularly interesting. This algorithm is inspired by the natural selection in biological evolution and works as follows. The core of a genetic algorithm is called recombination and is shown in Algorithm 1. A random population is created and placed in the search space. We refer to a population as a plurality of chosen points (individuals) in the search space. The fittest individuals have the best chance to procreate and produce offspring. The fitness in this context is the function value at the point that is associated with a certain individual.