Computers and Operations Research 95 (2018) 1–11

Contents lists available at ScienceDirect

Computers and Operations Research

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

Convex preference cone-based approach for many objective optimization problems

Ankur Sinha[a,*], Pekka Malo[b], Markku Kallio[b]

[a] Production and Quantitative Methods, Indian Institute of Management Ahmedabad, Ahmedabad 380015, India
[b] Department of Information and Service Economy, Aalto University School of Business, PO Box 21220, 00076 Aalto, Helsinki, Finland

ARTICLE INFO

Article history:
Received 29 June 2017
Revised 21 February 2018
Accepted 23 February 2018

Keywords:
Evolutionary algorithms
Many-objective optimization
Interactive
Cone contraction
Preference-based method

ABSTRACT

Many objective optimization problems have turned out to be a considerable challenge for evolutionary algorithms due to the difficulty of finding and visualizing high-dimensional Pareto frontiers. Fortunately, however, the task can be simplified whenever an interaction with a human decision maker is possible. Instead of finding the entire Pareto frontier, the evolutionary search can be guided to the parts of the space that are most relevant for the decision maker. In this paper, we propose an interactive method for solving many objective optimization problems. Drawing on the recent developments in multiple criteria decision making, we introduce an effective strategy for leveraging polyhedral preference cones within an evolutionary algorithm. The approach is mathematically motivated and is applicable to situations, where the user's preferences can be assumed to follow an unknown quasi-concave and increasing utility function. In addition to considering the preference cones as a tool for eliminating non-preferred solution candidates, we also present how the the cones can be leveraged in approximating the directions of steepest ascent to guide the subsequent search done by the evolutionary algorithm through a proposed merit function. To evaluate the performance of the algorithm, we consider well known test problems as well as a practical facility location problem.

1. Introduction

Large part of the research on evolutionary computation has focused on solving multi-objective optimization problems with two or three objectives (Deb, 2001). However, the recent past has shown an increasing need for algorithms that are able to handle a larger number of objectives (Chikumbo et al., 2012; Ishibuchi et al., 2008). Such problems are commonly referred to as many-objective optimization problems by the evolutionary optimization community. To match with the demand, several attempts have been made to scale the existing evolutionary multi-objective (EMO) algorithms to deal with more than three objectives. However, the task has turned out to be considerably more challenging than what may have been anticipated. A major hurdle for the functioning of Pareto dominance based algorithms has been that the proportion of non-dominated solutions in a randomly chosen set of objective vectors becomes large with an increased number of objectives (Deb and Jain, 2014). Consequently, the search ability of EMOs quickly deteriorates to the extent that their usefulness for solving such problems can be questioned (Deb and Jain, 2014; Deb and Saxena, 1997). Even visualization of the resulting frontiers and evaluation of performance metrics, such as hyper-volume, become difficult tasks.

While considering the many-objective optimization challenge, it is worthwhile to note that most of the use-cases are in context of decision-making problems, where the users are interested in evaluating only a very limited part of the search space. Hence looking at the problems from the decision maker's perspective offers a quite different direction for solving them. Instead of striving to find the entire Pareto optimal front at a large computational cost, it is possible to leverage the user's preference information to guide the search into the most relevant areas of the search space. Such approaches have been common in the area of Multi-Criterion Decision Making (Steuer, 1986) and are referred to as interactive approaches. However, these approaches have been relatively new in the context of EMOs. Most of the existing evolutionary methodologies for many objective optimization are not designed to be interactive; preference information from the decision maker is utilized either before the beginning of the search process (a priori approach) or at the end of the search process (a posteriori approach) to produce the optimal solutions, but not in the middle. Some studies in this direction are biased niching based EMO (Branke and Deb, 2004), reference point based EMO approach

* Corresponding author.
E-mail addresses: asinha@iima.ac.in, asinha@iimahd.ernet.in (A. Sinha), Pekka.Malo@aalto.fi (P. Malo), markku.kallio@aalto.fi (M. Kallio).

https://doi.org/10.1016/j.cor.2018.02.015
0305-0548/

(Deb et al., 2006), the light beam approach and reference direction based EMO (Thiele et al., 2009). Once the optimization process is started, there is no interaction involved with the decision maker until a set of representative Pareto solutions is found. This is a reasonable strategy, when there are ways to visualize the frontier geometrically. But in the case of many objective problems, a more convenient strategy may be obtained by allowing interactions to collect preference information from the decision maker after every few generations (Deb et al., 2010; Sinha et al., 2014). Some approaches interact with the decision maker and iterate the process of elicitation and search until a satisfactory solution is found. However, only a limited number of studies have been performed where preference information is elicited during the search process (Branke and Deb, 2004; Fowler et al., 2009; Jaszkiewicz, 2007; Purshouse et al., 2014).

Our aim in this paper is to introduce an interactive evolutionary algorithm for handling many-objective optimization problems. Throughout the construction, we assume that the underlying unknown value function, which characterizes the decision-maker's preferences, is quasi-concave and strictly increasing. Under this assumption, in each interactive step, the decision-maker is presented a set of S points which are to be ranked via pairwise comparisons. The comparison results are then used to develop convex polyhedral cones whose points cannot be better than the vertex of the cone (Kallio and Halme, 2013; Korhonen et al., 1984). The preference cones are utilized in three different ways within the algorithm: (i) first, they provide an efficient way to eliminate non-preferred candidates from the set of points to be shown to the user; (ii) second, we can extract information on the directions where the user's unknown value function is likely to increase at the fastest rates; (iii) finally, we can also construct a cone-based measure (merit function) for the quality of the solution candidates. Both merit function and the information on the steepest ascent directions have proven to be helpful in accelerating the convergence of the algorithm towards the regions that are most relevant for the decision-maker. Convergence can be guaranteed for pure integer problems with finite Pareto sets, but for more general problems the algorithm continues until a satisficing solution is found or the fixed budget of user-interactions is consumed.

The approach draws heavily on the insights from the literature on multi-criteria decision making, where cone contraction methods have been previously studied by Kallio and Halme (2013), Korhonen et al. (1984) and Kadzinski and Slowinski (2012). In this paper, we utilize the definitions proposed by Kallio and Halme (2013) due to benefit of obtaining a simple and operational test for checking if a candidate point should be considered non-preferred. However, our approach differs from Kallio and Halme (2013) in its diverse use of the preference cones. Instead of restricting the use only to the elimination of non-preferred candidates, we also leverage the cones in approximating the directions of steepest ascent to guide the subsequent search done by the evolutionary algorithm.

The use of preference cones is a relatively new idea in the context of evolutionary algorithms. Though the fields of evolutionary computation and multi-criteria decision making (MCDM) share a common goal, researchers have shown only lukewarm interest, until recently, in applying the principles of one field to the other. Our approach is perhaps closest to the ideas proposed in Deb et al. (2010) and Sinha et al. (2014), where the decision maker's preferences are approximated using polynomial value functions or by heuristically constructed polyhedral cones. Therefore, we have considered these results as baselines for our experiments. In Deb et al. (2010), preference information is used to build a strictly monotone polynomial value function, which is then utilized within a multi-objective algorithm to guide the progress towards the most preferred solution. To a degree, this resembles our approach. However, in the current paper, the value function is replaced by polyhedral cones that are constructed from a small set of pairwise comparisons. One of the benefits is that we no longer have to assume a fixed parametric form for the value function, but we can accommodate any quasi-concave monotone value function. The cones are also likely to be more tolerant towards preference uncertainty. Though polyhedral cones have been heuristically used in Sinha et al. (2014, 2010) to incorporate preference information in a multi-objective evolutionary algorithm, the current work builds on a strong mathematical foundation which lacks in the earlier study. The strengths of the proposed idea can be observed in the comparative study performed in the paper, where the cone contraction method is found to converge towards the most preferred point in very few iterations with the decision maker, thereby reducing the cognitive burden.

The rest of the article is organized as follows. In Section 2, we formulate the many-objective problem and discuss the central assumptions limiting the class of applicable problems. Section 3 presents several key definitions regarding representation of preference information using convex cones. Section 4 outlines the evolutionary algorithm and describes how the preference cones can be operationalized in practice. In Section 6, the working of the algorithm is demonstrated on three different types of test problems. One of the problems is a practical multi-criteria facility location problem, and the other two are well-known test problems that have been discussed in Deb et al. (2005). Finally, we conclude in Section 7.

2. Many-objective optimization problem

Consider a many-objective optimization problem with k objectives to be maximized with respect to continuous and/or integer variables. Let $f \in R^k$ denote the column vector of objectives. The underlying value function which the decision maker (DM) aims to maximize is $u(f)$. At the outset the value function is not known, but we assume that u is quasi-concave and strictly increasing.

Given a set F of feasible objective vectors of size k, such that $f = (f_1, \dots, f_k) \in F$, the many-objective problem is

$$\max_{f \in F} u(f). \tag{1}$$

We consider the feasible set

$$F = \{f \,|\, f \le f^U(x),\ \ h(x) \ge 0,\ \ x_j \text{ integer for } j \in J\}, \tag{2}$$

where $x \in R^n$ is a vector of decision variables, function $h \in R^m$ is concave, $f^U \in R^k$ is an upper bound that is concave in the domain $\{x \,|\, h(x) \ge 0\}$, and J is a subset of $\{1, \dots, n\}$ with $J = \{1, \dots, n\}$ for pure integer problems. For an optimal point $f' \in F$, $f' = f^U(x)$. It is noteworthy that convexity assumptions are not required for EMO approaches; however, it forms the basis for some of the theoretical discussions in the paper. The notations used in the paper have been summarized in Table 1. Dot product between two vectors a and b has been denoted as ab throughout the paper.

3. Representation of preference information using convex cones

This section begins with a review of basic properties of quasi-concave and increasing value functions. Thereafter, we introduce convex cones which are convenient for representing partial preference information accumulating over interactive iterations for optimization. Such cones identify points which are inferior to points already found, provide local information on the gradient of the value function, and are used to define a merit function. The merit function is one of the important contributions made in this paper, which is particularly suitable for assigning fitness in an evolutionary algorithm and contains all the information related to the cones.