Information Sciences 424 (2018) 159–174

Contents lists available at ScienceDirect

Information Sciences

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

Multi-Objective Simultaneous Optimistic Optimization

Abdullah Al-Dujaili, S. Suresh*

School of Computer Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798, Singapore

ARTICLE INFO

Article history:
Received 6 March 2017
Revised 16 September 2017
Accepted 29 September 2017
Available online 3 October 2017

Keywords:
Multi-objective optimization
Optimistic methods
Multi-armed bandits
Simultaneous optimistic optimization
Finite-time analysis
Asymptotic analysis

ABSTRACT

Optimistic methods have been applied with success to single-objective optimization. Here, we attempt to bridge the gap between optimistic methods and multi-objective optimization. In particular, this paper is concerned with solving black-box multi-objective problems given a finite number of function evaluations and proposes an optimistic approach, which we refer to as the Multi-Objective Simultaneous Optimistic Optimization (MO-SOO). Popularized by multi-armed bandits, MO-SOO follows the optimism in the face of uncertainty principle to recognize Pareto optimal solutions, by combining several multi-armed bandits in a hierarchical structure over the feasible decision space of a multi-objective problem. Based on three assumptions about the objective functions smoothness and hierarchical partitioning, the algorithm finite-time and asymptotic convergence behaviors are analyzed. The finite-time analysis establishes an upper bound on the Pareto-compliant unary additive epsilon indicator characterized by the objectives smoothness as well as the structure of the Pareto front with respect to its extrema. On the other hand, the asymptotic analysis indicates the consistency property of MO-SOO. Moreover, we validate the theoretical provable performance of the algorithm on a set of synthetic problems. Finally, three-hundred bi-objective benchmark problems from the literature are used to substantiate the performance of the optimistic approach and compare it with three state-of-the-art stochastic algorithms in terms of two Pareto-compliant quality indicators. Besides sound theoretical properties, MO-SOO shows a performance on a par with the top performing stochastic algorithm.

1. Introduction

Many real-world application and decision problems involve optimizing two or more objectives at the same time (see, e.g., [15]). These problems are often referred to as Multi-Objective Optimization (MOO). In the general case, MOO problems are hard because the objective functions are often conflictual, and it is difficult to design strategies that are optimal for all objectives simultaneously. Furthermore, with conflicting objectives, there does not exist a single optimal solution but a set of incomparable optimal solutions: each is inferior to the other in some objectives and superior in other objectives. This induces a partial order on the set of feasible solutions to an MOO problem. The set of optimal feasible solutions according to this partial order is referred to as the *Pareto optimal set* and its corresponding image in the objective space is commonly named as the *Pareto front* of the problem. The task of MOO algorithms therefore becomes finding the Pareto front or producing a good approximation of it (referred to as an *approximation set* of the problem).

* Corresponding author.
E-mail addresses: aldujail001@e.ntu.edu.sg (A. Al-Dujaili), ssundaram@ntu.edu.sg (S. Suresh).

https://doi.org/10.1016/j.ins.2017.09.066

Generally, certain assumptions are made about the objective functions being optimized (e.g., its continuity or differentiability). However, these assumptions are not necessarily satisfied by real-world problems. Sometimes, the only information available about the objective functions are their point-wise evaluations: computing their derivatives or other measures are either expensive, unreliable, or even impossible. Such problems are called *black-box* multi-objective optimization problems and appear very often in real-world settings [19]. In this paper, we study the problem of black-box MOO given a finite number of objective functions evaluations (often referred to as the evaluation budget).

Conventionally, solving a multi-objective optimization problem follows one of two principles, namely *preference-based* and *ideal* principles [4,13,34,40]. Following the preference-based principle, the MOO problem is transformed into a single-objective optimization problem (through an aggregation/scalarization function that exploits a priori information), which then can be solved using one of many available single-objective optimizers [20]. While preference-based algorithms converge to a single solution in each run, ideal-based algorithms search for a set of solutions at once. One example in this approach is evolutionary multi-objective algorithms [42] in which a population of solutions evolves, following a crude analogy with Darwinian evolution, towards better solutions. Recently, there has been a growing interest of formulating multi-objective problems within the framework of reinforcement learning (see, for instance, [29]).

Among the several lessons learned from the aforementioned MOO solvers over the past decades is that, in order to generate a dense and good approximation set, one must maintain the set diversity. Furthermore, one must not discard inferior solutions too easily, as some of them may pave the way towards rarely-visited regions of the Pareto front [25]. In other words, the exploration-vs.-exploitation trade-off in search for the Pareto optimal set should be thought carefully about, at the *algorithmic design level*. With this regard, in this paper, we are motivated to address the problem of multi-objective optimization within the framework of *optimistic sequential decision-making methods*. i.e., methods that implement the *optimism in the face of uncertainty* principle. Such principle finds its foundations in the machine learning field addressing the exploration-vs.-exploitation dilemma, known as the *multi-armed bandit problem* [35].

Within the context of single-objective optimization, optimistic sequential decision-making approaches formulate the complex problem of global optimization over the decision space $\mathcal{X}$ as a hierarchy of simple bandit problems over subspaces of $\mathcal{X}$ and look for the optimal solution through $\mathcal{X}$-partitioning search trees: each leaf corresponds to a subspace of $\mathcal{X}$, with the root corresponding to $\mathcal{X}$ and nodes at depth $h \in \mathbb{N}_0$ represent a partition of $\mathcal{X}$ at scale h. At step t, such algorithms optimistically expand a leaf node (i.e., partition the corresponding subspace) that may contain the optimum. In other words, optimistic algorithms consider partitions of the search space at multiple scales in search for the optimal solution [9,28]. Recently, the optimistic optimization algorithm, Naive Multi-scale Search Optimization [3], has been shown to be a viable alternative to solve black-box optimization problems – see the results of the Black-Box Optimization Competition (BBComp) within the Genetic and Evolutionary Computation Conference (GECCO'2015) [26].

On the other hand, two observations can be made about optimistic methods within the context of multi-objective optimization. First, there has been very little/limited yet slowly growing research reported on optimistic methods for multi-objective optimization. For instance, the focus of multi-objective multi-armed bandit problems has been distinctly on a discrete set of arms [16], or solving a subproblem (e.g., selecting a genetic operator in evolutionary multi-objective algorithms [24]). Second, the algorithmic development and validation have been dominantly empirical (see, for instance, [2]).

Being one of the simplest single-objective optimistic methods with a theoretically provable performance, this paper is inspired by the Simultaneous Optimistic Optimization (SOO) [28] to develop an optimistic algorithm for multi-objective problems. We refer to this algorithm as the Multi-Objective Simultaneous Optimistic Optimization (MO-SOO). In order to find a good approximation set of the Pareto front, MO-SOO employs – similar to optimistic methods – hierarchical bandits over the decision space. Represented by a divide-and-conquer tree structure, the hierarchical bandits are realized by partitioning the decision space over multiple scales. At each step, MO-SOO expands leaf nodes (partitions the corresponding subspaces) that may optimistically contain *Pareto optimal* solutions. Based on three assumptions about the function smoothness and partitioning strategy, we analyze the finite-time and asymptotic convergence behaviors of MO-SOO. The finite-time study is based on quantifying how much exploration is required to achieve near-optimal objective-wise solutions. As a result, we are able to upper bound the loss of the obtained solutions with respect to the objective-wise optimal solutions. Using this objective-wise loss bound, an upper bound on the Pareto-compliant unary additive epsilon indicator [43] is established as a function of the number of iterations. The bound is characterized by the objectives smoothness as well as the structure of the Pareto front with respect to its extrema. First time in the literature, a deterministic upper bound on a Pareto-compliant indicator is presented for a solver of continuous MOO problems. However, the presented bound holds down to a problem-dependent constant. Furthermore, the systematic sampling nature of the decision space in MO-SOO helps in analyzing the asymptotic behavior, which indicates its consistency, viz. optimality in the limit. Using symbolic maths, the theoretical provable performance of the algorithm has been validated on a synthetic problem.

Complementing the theoretical results, an empirical validation study has been conducted using 300 bi-objective benchmark problems from the literature [8]. The test suite considers problems with various objective functions categories reflecting real-world scenarios such as separability and multi-modality. It can also be used to validate the algorithms scalability with the decision space dimension. To test the performance of our proposition in multiple objectives setting, we benchmark the algorithm on 100 multi-objective problems collected from the literature.

Furthermore, MO-SOO has been compared with 4 state-of-the-art stochastic algorithms, namely MOEA/D [41], MO-CMA-ES [37], NSGA-II [14], and SMS-EMOA [7] in terms of two Pareto-compliant quality indicators [21]: the hypervolume (I_H^-) and the unary additive ϵ-indicator ($I_{\epsilon+}^1$). The results are presented in form of data profiles, which adequately