Artificial Intelligence 228 (2015) 45–65

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Bi-goal evolution for many-objective optimization problems

Miqing Li [a], Shengxiang Yang [b,*], Xiaohui Liu [a]

[a] *Department of Computer Science, Brunel University, London UB8 3PH, UK*
[b] *Centre for Computational Intelligence (CCI), School of Computer Science and Informatics, De Montfort University, Leicester LE1 9BH, UK*

ARTICLE INFO

Article history:
Received 22 August 2014
Received in revised form 14 June 2015
Accepted 20 June 2015
Available online 3 July 2015

Keywords:
Evolutionary multi-objective optimization
Many-objective optimization
Proximity
Diversity
Bi-goal evolution

ABSTRACT

This paper presents a meta-objective optimization approach, called Bi-Goal Evolution (BiGE), to deal with multi-objective optimization problems with many objectives. In multi-objective optimization, it is generally observed that 1) the conflict between the proximity and diversity requirements is aggravated with the increase of the number of objectives and 2) the Pareto dominance loses its effectiveness for a high-dimensional space but works well on a low-dimensional space. Inspired by these two observations, BiGE converts a given multi-objective optimization problem into a bi-goal (objective) optimization problem regarding proximity and diversity, and then handles it using the Pareto dominance relation in this bi-goal domain. Implemented with estimation methods of individuals' performance and the classic Pareto nondominated sorting procedure, BiGE divides individuals into different nondominated layers and attempts to put well-converged and well-distributed individuals into the first few layers. From a series of extensive experiments on four groups of well-defined continuous and combinatorial optimization problems with 5, 10 and 15 objectives, BiGE has been found to be very competitive against five state-of-the-art algorithms in balancing proximity and diversity. The proposed approach is the first step towards a new way of addressing many-objective problems as well as indicating several important issues for future development of this type of algorithms.

1. Introduction

Real-world problems commonly involve multiple objectives/criteria which are required to be optimized simultaneously. For example, an individual would like to maximize the chance of being healthy and wealthy while still having fun and time for family and friends. A software engineer would be interested in finding the cheapest test suite while achieving full coverage (e.g., statement coverage, branch coverage and decision coverage). When prescribing radiotherapy to a cancer patient, a doctor would have to balance the attack on tumor, potential impact on healthy organs, and the overall condition of the patient. These multi-objective optimization problems (MOPs) can be seen in many fields, including engineering, science, medicine and logistics. They share the same issue of pursuing several objectives at the same time, and have long been regarded as a substantial challenge in artificial intelligence (AI) [73,25].

There have been a variety of approaches for MOPs, including traditional mathematical programming methods, local search techniques, and evolutionary algorithms (EAs). Inspired by biological evolution mechanisms, EAs have been demonstrated to be successful in diverse AI applications [73,10]. For example, an EA-based AI planner, Divide and Evolutionary

* Corresponding author.
E-mail address: syang@dmu.ac.uk (S. Yang).

http://dx.doi.org/10.1016/j.artint.2015.06.007

(DaE) [8], won the Deterministic Temporal Satisficing track during the International Planning Competition (IPC7) at the 21st International Conference on Automated Planning and Scheduling (ICAPS 2011).[1] Recently, DaE has been successfully applied to multi-objective AI planning (called MO-DaE) [58]. MO-DaE, working with a well-known multi-objective EA, i.e., the indicator-based EA (IBEA) [99], has shown clear advantage over the metric-based approach using LPG metric sensitive planner [58].

A key strength of EAs for MOPs is their population-based feature which allows individuals to simultaneously approximate different parts of the Pareto front within a single execution [19,97]. Intuitively, the search process of an EA has two basic goals:

- minimizing the distance of the population to the Pareto front (i.e., proximity) and
- maximizing the distribution of the population along the Pareto front (i.e., diversity).

Since the optimal outcome of an MOP is a set of Pareto optimal solutions, the Pareto dominance relation naturally becomes a criterion to distinguish between solutions. Given two solutions p and q for an MOP, p is said to *Pareto dominate* q, if and only if p is better than q for at least one objective and is not worse for any of the others. The Pareto dominance reflects the weakest assumption about the preferred structure of the decision-maker.

As the primary selection criterion in the evolutionary multi-objective optimization (EMO) area, Pareto dominance is commonly used to evaluate the proximity of solutions. When Pareto dominance fails (e.g., the interested solutions are non-dominated to each other), EMO algorithms often introduce a density-based criterion to maintain diversity of the population. For example, the nondominated sorting genetic algorithm II (NSGA-II) [23] separates individuals in a population into different layers (ranks) by their Pareto dominance relation, and prefers 1) individuals in lower layers and 2) individuals with lower crowding degrees (measured by the *crowding distance* [23]) when they are located in the same layer.

An MOP with more than three objectives is called a many-objective optimization problem. Many-objective optimization is an important but very challenging topic and there has been increasing interest in the use of EAs to tackle many-objective optimization problems [14,16,26,35]. Although Pareto-based algorithms are the most popular approaches, they scale up poorly with the number of objectives [18,48,75]. When dealing with an MOP with many objectives, Pareto dominance often loses its effectiveness to differentiate individuals [57], which makes most individuals in a population become incomparable in terms of proximity (e.g., in NSGA-II most individuals fall into the first layer). Consequently, the density-based selection criterion will play a decisive role in determining the survival of individuals during the evolutionary process, leading to the individuals in the final population distributed widely over the objective space but far from the desired Pareto front [85].

A straightforward way to handle this problem (i.e., the ineffectiveness of Pareto-based algorithms in many-objective optimization) is to modify the Pareto dominance relation. Some interesting attempts include loosening the dominance condition or controlling the dominance angle, such as ϵ-dominance [22,36,61,84], α-dominance [43], ϵ-box dominance [60], and dominance area control [78]. By relaxing the area of an individual dominating, these dominance relations are able to provide sufficient selection pressure towards the Pareto front. However, how to set a proper value of the parameter(s) to determine the relaxation degree is a crucial issue in these methods, needing further studies [62,69,79].

On the other hand, the way of comparing individuals according to their quantitative difference in objectives has been found to be effective in converging towards the Pareto front. Many recent EMO algorithms originate from this motivation, introducing a variety of new criteria to distinguish between individuals, e.g., average ranking [52,70], fuzzy Pareto optimality [37,39], subspace partition [2,51], preference-inspired rank [88,87], grid-based rank [70,92], distance-based rank [32,71,91], and density adjustment strategies [1,66]. These methods provide ample alternatives to deal with many-objective optimization problems, despite some having the risk of leading the population to concentrate in one or several sub-areas of the whole Pareto front [50,67,81,65].

Recently, there has been significant interest in the use of selection criteria that involve both proximity and diversity to solve MOPs. Some such criteria, like the decomposition-based [94] and indicator-based [99] criteria, have been shown to be very promising in many-objective optimization [15,20,41,44,85]. The former uses the idea of single-objective aggregated optimization, decomposing an MOP into a number of scalar subproblems and optimizing them simultaneously. The latter defines an optimization criterion with regard to a specified performance indicator and uses this criterion to guide the search of the population. The indicator *hypervolume* is one of the most popular indicator-based criteria due to its good theoretical and empirical properties [7,13,29,42,101]. Whereas super-polynomial time complexity is required in the calculation of the hypervolume indicator (unless $P = NP$) [11], lots of effort is being made to reduce its computational cost, in terms of both the exact computation [6,12,90] and the approximate estimation [4,14,49]. Nevertheless, balancing proximity and diversity using one single criterion is not an easy task [76,38,69,68], especially for a many-objective optimization problem in which the conflict between the objectives is generally more serious than that in an MOP with two or three objectives [75,1].

In fact, evolving a population towards the optimum as well as diversifying its individuals over the whole Pareto front in many-objective optimization is, by itself, a multi-objective problem. The advance at one aspect usually comes along with the degradation at the other [33,75].

[1] http://www.sigevo.org/wiki/tiki-read_article.php?articleId=1.