ARTICLE IN PRESS

Swarm and Evolutionary Computation (xxxx) xxxx–xxxx

Contents lists available at ScienceDirect

Swarm and Evolutionary Computation

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

Regular Paper

A fast hypervolume driven selection mechanism for many-objective optimisation problems

Shahin Rostami[a,*], Ferrante Neri[b]

[a] Department of Computing and Informatics, Bournemouth University, Bournemouth BH12 5BB, United Kingdom
[b] Centre for Computational Intelligence, De Montfort University, Leicester LE1 9BH, United Kingdom

ARTICLE INFO

Keywords:
Multi-objective optimisation
Many-objective optimisation
Hypervolume indicator
Selection mechanism
Evolutionary optimisation

ABSTRACT

Solutions to real-world problems often require the simultaneous optimisation of multiple conflicting objectives. In the presence of four or more objectives, the problem is referred to as a "many-objective optimisation problem". A problem of this category introduces many challenges, one of which is the effective and efficient selection of optimal solutions.

The hypervolume indicator (or *s*-metric), i.e. the size of dominated objective space, is an effective selection criterion for many-objective optimisation. The indicator is used to measure the quality of a non-dominated set, and can be used to sort solutions for selection as part of the contributing hypervolume indicator. However, hypervolume based selection methods can have a very high, if not infeasible, computational cost.

The present study proposes a novel hypervolume driven selection mechanism for many-objective problems, whilst maintaining a feasible computational cost. This approach, named the Hypervolume Adaptive Grid Algorithm (HAGA), uses two-phases (narrow and broad) to prevent population-wide calculation of the contributing hypervolume indicator. Instead, HAGA only calculates the contributing hypervolume indicator for grid populations, i.e. for a few solutions, which are close in proximity (in the objective space) to a candidate solution when in competition for survival. The result is a trade-off between complete accuracy in selecting the fittest individuals in regards to hypervolume quality, and a feasible computational time in many-objective space. The real-world efficiency of the proposed selection mechanism is demonstrated within the optimisation of a classifier for concealed weapon detection.

1. Introduction

Optimisation metaheuristics are composed of two phases: search to generate a new candidate solution and selection to choose the solutions to retain for the following iteration, e.g. see [26,12]. In multi-objective optimisation, the most critical operation is the selection since the fitness based comparisons must take into account the fact that a solution can be better performing than another in terms of one objective and not another. Candidate solutions in this situation are said to not dominate each other. The theoretical set of solutions which are not dominated by any other solution is referred to as Pareto-optimal (or simply Pareto) set [21,66]. Metaheuristics designed to solve multi-objective problems aim to detect an approximation of the Pareto set (approximation set) [103]. The term approximation set is used to refer to "the set of all non-dominated points found during the run" [52], that is, the population at each iteration/generation of a multi-objective optimisation algorithm [53].

Many applications, such as engineering design, require that one solution (or in some cases a few alternatives) rather than a large set is ultimately selected. The process of performing this selection is named Decision Making while the criterion or algorithm that leads to the decision making is said to be the Decision Maker (DM). In other words, the DM implicitly classifies "interesting and uninteresting" solutions. The area of the objective space where the interesting solutions fall within is named the Region Of Interest (ROI). It must be noted that the multi-objective optimisation algorithm that detects the set of non-dominated solutions and the DM are related entities that perform different phases of the same task.

A good representation of a Pareto-optimal set, in terms of DM action, is characterised in three key areas, see [73]. These are illustrated graphically in Fig. 1 and listed in the following:

- *Proximity*: This tells the DM how close the approximation set is to the true Pareto-optimal front. An ideal approximation set should be as close as possible in proximity to the true Pareto-optimal front. In practise, proximity cannot be used as a measure of quality of the

* Corresponding author.
E-mail addresses: srostami@bournemouth.ac.uk (S. Rostami), fneri@dmu.ac.uk (F. Neri).

http://dx.doi.org/10.1016/j.swevo.2016.12.002
Received 13 July 2016; Received in revised form 23 November 2016; Accepted 11 December 2016
2210-6502/

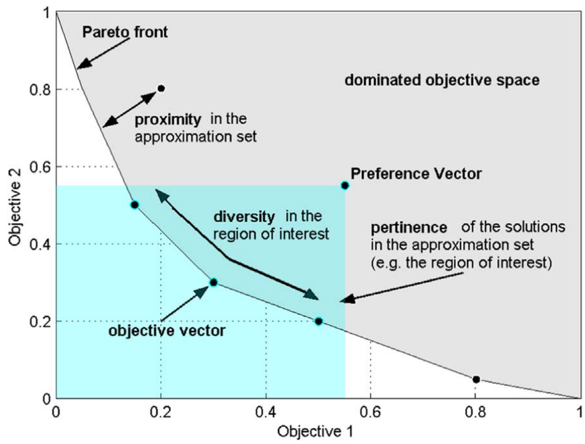

Fig. 1. Proximity, diversity, and pertinence characteristics in an approximation set in two-objective space.

approximation set during the optimisation process, because the true Pareto set is not known.

- *Diversity*: This characterises the distribution of the approximation set both in the extent and uniformity of that distribution. The ideal approximation set should be uniformly distributed across the trade-off surface of the problem.
- *Pertinence*: This criteria measures the relevance of the approximation set to the DM. Ideally the approximation set should contain a number of solutions which satisfy the DM's expressed preferences.

Conventional multi-objective optimisation techniques often fail to satisfy these criteria. For example, the goal-attainment method [34] and the weighted-sum method [39] both only provide single solutions to the optimisation problem—thus failing to provide a diverse distribution of solutions. Population based meta-heuristics, such as Evolutionary Multi-Objective Optimisation (EMO) algorithms, are naturally more suitable to tackle multi-objective problems since they process a population of solutions which can then represent the Pareto set, see [23,45,55,65,79,95,98], also when coupled to local search components, e.g. see [78]. Furthermore, population-based algorithms for multi-objective optimisation can be easily endowed with simple and effective components to maintain a diversity of high quality solutions. Recently, in [48], it is proposed a selection mechanism which satisfies at first the diversity of the solutions and then promotes those with the highest proximity. An alternative approach would make use of a mathematical model to generate extra surrogate solutions, e.g. see [13].

A study on the effectiveness of variation in EMO algorithms is reported in [2].

1.1. Many-objective optimisation

The higher the number of objectives, the more challenging the pairwise comparison of solutions and the subsequent selection process. A multi-objective optimisation problem with more than three objectives is referred to as *many-objective* optimisation problem [29,43,57].

Analogous to the curse of dimensionality when large scale problems are considered, many-objective problems can introduce challenging difficulties. These challenges have been analysed in the literature, e.g. see [16,17,43], and summarised in the following list:

- It is likely that almost all candidate solutions found throughout the optimisation process will be non-dominated, this poses an issue for EMO algorithms which rely on Pareto-dominance for selection pressure, [30,54,74].
- The number of candidate solutions required to produce an approximation set which reliably represents the trade-off surface increases exponentially [41,50].
- The number of generations required to produce an approximation set increases, thus making the computational cost of a single run very high, if not infeasible.
- Search operators become ineffective at detecting new non-dominated solutions in the presence of many-objectives [41].
- Approaches which attempt to promote the diversity in the objective space, see [74], can cause issues in terms of convergence, see [85,59,57]. Indeed, convergence and diversity are conflicting in the many-objective case, see [74,57].
- The visualisation of candidate solutions becomes difficult, often resulting in the use of heat-maps or parallel-coordinate plots. This poses a difficulty to the DM as the selection of a final candidate solution may become non-intuitive [86].

The transition between the multi-objective problem domain and the many-objective problem domain is not straightforward, such that the methods used to optimise solutions for a multi-objective problem have little to none of the desired effect when applied to a many-objective problem. A fundamental example of this is that the selection mechanisms based dominance that perform well on multi-objective problems (two or three objectives) [18,21,80], often do not perform well when four or more problem objectives are considered as shown in [3,37,38,42,47,51,64,74,104]. Selection based on dominance is inefficient at producing a strong selection pressure toward the Pareto-optimal front in the presence of many objectives, as throughout the optimisation process it is likely that the entire population will consist of entirely non-dominated solutions.

Several alternative algorithmic solutions have been proposed to perform the selection. The following non-exhaustive classification is here proposed.

- *Selection methods that use a reference vector*: These methods are focussed on the diversity and ideally aim at achieving an approximation set equally spaced on the Pareto front. In order to achieve this aim, these methods use a reference vector (or weight) and normal distributions to select along each coordinate (in the objective space) the points that are sufficiently distant. A prominent family of algorithms based on this logic consists of the algorithms based on Multi-Objective Evolutionary Algorithm based on Decomposition (MOEA/D) introduced in [96] and initially proposed for multi-objective problems. A further development of the MOEA/D algorithm has been presented in [97] where the MOAE/D with Dynamic Resource Allocation (MOEA/D-DRA) has been introduced. MOEA/D-DRA decomposes the many-objective space into multiple single-objective spaces (sub-problems) and then assigns them different computational budgets. This algorithm, which has been a competition winner at IEEE CEC is currently one of the most effective and robust solutions at tackling complex multi-objective and at least five-objective problems. Another interesting example of this category is the Non-dominated Sorting Genetic Algorithm III (NSGA-III) [22,46]. A further feature of NSGA-III is the use of niching during recombination. This mechanism has been proposed to increase the exploitation of the algorithm. A recent example of an algorithm based on this logic is given in [94]. Another modern example is given in [49] where although a reference vector is used, a diversity-first convergence-second selection strategy.
- *Selection methods that divide/classify the population*: These methods have the same purpose of the methods based on reference vectors but achieve this goal by mapping and dividing the objective space. A mapping is then used as a reference to select the points so that they are equally spaced. A historical example in multi-objective optimisation is the Adaptive Grid Algorithm (AGA), see [53], where a grid in the objective space is used to control the population diversity. Another important example is [74] which employed the mechanism to promote diversity. The employment of a grid in the objective space has been reinterpreted and implemented in [61,92].