Quantifying the Effects of Objective Space Dimension in Evolutionary Multiobjective Optimization

Joshua Knowles[1] and David Corne[2]

[1] School of Computer Science, Kilburn Building, University of Manchester, Manchester M13 9PL, UK
j.knowles@manchester.ac.uk
[2] School of Mathematical and Computer Sciences (MACS), Heriot-Watt University, UK

Abstract. The scalability of EMO algorithms is an issue of significant concern for both algorithm developers and users. A key aspect of the issue is scalability to objective space dimension, other things being equal. Here, we make some observations about the efficiency of search in discrete spaces as a function of the number of objectives, considering both uncorrelated and correlated objective values. Efficiency is expressed in terms of a cardinality-based (scaling-independent) performance indicator. Considering random sampling of the search space, we measure, empirically, the fraction of the true PF covered after p iterations, as the number of objectives grows, and for different correlations. A general analytical expression for the expected performance of random search is derived, and is shown to agree with the empirical results. We postulate that for even moderately large numbers of objectives, random search will be competitive with an EMO algorithm and show that this is the case empirically: on a function where each objective is relatively easy for an EA to optimize (an NK-landscape with K=2), random search compares favourably to a well-known EMO algorithm when objective space dimension is ten, for a range of inter-objective correlation values. The analytical methods presented here may be useful for benchmarking of other EMO algorithms.

Keywords: multiobjective optimization, nondominated sorting, nondominated ranking, random search, coverage indicator, inter-objective correlation, many objectives.

1 Introduction

The past two decades have seen the development of more and more effective and efficient evolutionary multiobjective optimization (EMO) algorithms [4,6]. These methods are often run with the goal of approximating the whole Pareto front and most EMO algorithms are designed to do this on problems of arbitrary parameter space and objective space dimension. Yet, the scalability of these methods, in practice, remains an issue of concern for the field.

Empirical testing of EMO algorithms relies on both test functions (see [11] for a review) and performance assessment methods [10,17,21,24]. Today, some test

S. Obayashi et al. (Eds.): EMO 2007, LNCS 4403, pp. 757–771, 2007.

functions are scalable in both parameter and objective dimension [11]; and some performance indicators are also suitable for many objective problems (notably those based on counting, i.e. cardinality-based indicators). These advances have made it possible to compare performance of EMO algorithms when the number of objectives is scaled up beyond the typical two or three. Thus, recently, researchers have shown empirically that some EMO algorithms (especially those based on dominance ranking for selection) perform poorly when the number of objectives d is greater than three [7,8,12,13,18], some suggesting alternative approaches.

However, it would be useful to know to what extent EMO algorithms are really performing poorly, relative to some absolute level of performance. In other words, it would be good to get some idea of the intrinsic difficulty of search as a function of objective space dimension. More precisely, we would like to know how particular performance statistics change as a function of objective space dimension, for a baseline method on a baseline/generic problem.

In this paper, we consider the performance of random search as an informative baseline, which is in line with a suggestion in [14]. Further, it is possible to be independent of the specifics of an objective function: if a one-to-one mapping from parameter to objective space is assumed, and sampling is random, points can be chosen from the objective space rather than the parameter space, without affecting the outcome, and hence a parameter space is not needed at all. Thus, discrete data sets consisting of objective vectors only are used, and two parameters are varied: the objective dimension, and a covariance term which influences the degree of inter-objective correlation. Since we are considering only cardinal, scaling-independent performance indicators, our conclusions are also independent of the data distribution in each objective (Gaussians are used for convenience). We also develop an analytic equation for predicting the expected performance of random search, and show that on these data it works.

The rest of the paper is organized as follows. Section 2 sets out some definitions and methods used in the remainder of the paper: it recalls the dominance relation; the performance indicator, *coverage*; ranking methods used in EMO fitness-assignment; and the data generation method and data sets used here. Section 3 presents empirical distributions for the number of nondominated points in our data sets and the distribution of coverage values obtained from runs of random search. In Section 4, we derive analytical expressions for the expected coverage indicator value for random sampling and show this agrees with the empirical results. Section 5 presents a case study where PESA-II is compared with random search on NK landscapes of varying dimension and inter-objective correlation. PESA-II compares unfavourably with the analytic performance of random search for 10 objectives, and this is confirmed empirically. Section 6 discusses the findings and suggests directions for further investigation.

2 Definitions and Methods

The standard definition of Pareto dominance in the objective space is used. Assuming mininimization, without loss of generality, x dominates y is written as $x \prec y$ and has the following meaning: