Applied Soft Computing 13 (2013) 1902–1921

Contents lists available at SciVerse ScienceDirect

Applied Soft Computing

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

Evolutionary algorithm characterization in real parameter optimization problems

Pilar Caamaño*, Francisco Bellas, Jose A. Becerra, Richard J. Duro

Integrated Group for Engineering Research, University of A Coruña, C/ Mendizabal s/n, 15403 Ferrol, Spain

ARTICLE INFO

Article history:
Received 4 November 2010
Received in revised form 4 October 2012
Accepted 6 January 2013
Available online 16 January 2013

Keywords:
Algorithm analysis
Real parameter optimization
Separability
Modality

ABSTRACT

This paper deals with the constant problem of establishing a usable and reliable evolutionary algorithm (EA) characterization procedure so that final users like engineers, mathematicians or physicists can have more specific information to choose the most suitable EA for a given problem. The practical goal behind this work is to provide insights into relevant features of fitness landscapes and their relationship to the performance of different algorithms. This should help users to minimize the typical initial stage in which they apply a well-known EA, or a modified version of it, to the functions they want to optimize without really taking into account its suitability to the particular features of the problem. This trial and error procedure is usually due to a lack of objective and detailed characterizations of the algorithms in the literature in terms of the types of functions or landscape characteristics they are well suited to handle and, more importantly, the types for which they are not appropriate. Specifically, the influence of separability and modality of the fitness landscapes on the behaviour of EAs is analysed in depth to conclude that the typical binary classification of the target functions into separable/non-separable and unimodal/multimodal is too general, and characterizing the EAs' response in these terms is misleading. Consequently, more detailed features of the fitness landscape in terms of separability and modality are proposed here and their relevance in the EAs' behaviour is shown through experimentation using standardized benchmark functions that are described using those features. Three different EAs, the genetic algorithm, the Covariance Matrix Adaptation Evolution Strategy and Differential Evolution, are evaluated over these benchmarks and their behaviour is explained in terms of the proposed features.

1. Introduction

When researchers from fields like engineering, applied physics or mathematics have to deal with a new optimization problem and try to address it using an evolutionary algorithm (EA), the typical procedure they follow is to apply a well-known algorithm and tune its configuration until a successful solution is obtained [66,65,48]. If this initial approach fails, the next step is to try to update it with a newer EA after consulting the literature on the topic, like the authors do in [21] or in [62], which may be useful or not, depending on the particular features of the problem to solve. This highly time-consuming trial and error stage is a consequence of the lack of a well-founded theoretical background in the evolutionary computation field that would allow analysing new algorithms in a formal way [4,9].

When a new algorithm is presented in the evolutionary computation literature, it is usually tested using problems chosen to highlight its capabilities as compared to other algorithms. For example, in [58], its authors present the Differential Evolution (DE) algorithm and compare it to five different real-parameter optimization techniques: two simulated annealing methods, two genetic algorithms and a stochastic differential equation method. They use three different test beds and the DE algorithm provides the best results in all of them. Consequently, one could think that this algorithm is a very good choice for every kind of optimization problem because there is no reference to the functions where it fails. In [28], the authors present the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) and in [2] a variant of it. They evaluate the algorithm considering a set of functions with different features, but, as a conclusion, they only highlight its invariance with respect to linear transformations. Again, one misses the inclusion of comments on the weak points and on the behaviour of the algorithm depending on the type of function to be optimized.

On the other hand, in papers where a more formal analysis of a given algorithm is performed, the comparison to other algorithms is usually too limited. For example, in [41], a comparison and analysis of different DE variants is presented, focusing on the analysis of the behaviour of the variants depending on the features of the objective function. In this case, a practical characterization of the DE algorithm is achieved, but the authors do not perform any other comparison to different algorithms that could behave

* Corresponding author at: Escuela Politécnica Superior, Universidade da Coruña, Mendizábal s/n, 15403 Ferrol, A Coruña, Spain. Tel.: +34 981 337400x3863; fax: +34 981 337410.

E-mail addresses: pcsobrino@udc.es (P. Caamaño), fran@udc.es (F. Bellas), ronin@udc.es (J.A. Becerra), richard@udc.es (R.J. Duro).

1568-4946/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.asoc.2013.01.002

better in those functions where DE and its variations fail. This kind of papers may be correct and appropriate from an EA researcher's point of view, but they are a bit frustrating for an EA user.

Recently, researchers in the field of evolutionary computation have become very interested in optimization algorithm competitions, like those organized within the IEEE CEC conference (Special Session on Real-Parameter Optimization CEC05, Evolutionary Computation in Dynamic and Uncertain Environments CEC09) or in the GECCO conference (EvoDOP-2007), where the algorithm characterization problem becomes more evident. Although in these competitions the algorithms are evaluated in a more systematic manner by specifying a common termination criterion, problem size, initialization scheme, linkages/rotations, etc. [59], they are typically focused on average performance measures. The optimization benchmarks used in these competitions classify the functions depending on their basic characteristics, that is, separable vs. non-separable, continuous vs. discontinuous, unimodal vs. multimodal. However, most authors do not provide general conclusions/limitations of the algorithms in terms of these characteristics. The practical consequence is that a researcher from a different field who needs to apply an EA to a particular problem will probably select the algorithm which obtains the best average results and is considered the "winner". As demonstrated later in this paper, this selection could prove to be completely wrong if the algorithm fails in a particular feature required for solving the specific problem the user faces.

Starting from this background, this work addresses the problem of establishing a usable and reliable evolutionary algorithm (EA) characterization procedure and defining a compact, albeit informative, set of problem features that can be related to the behaviour of the algorithms. The objective is that final users can have reliable and detailed information to choose the most suitable one for their particular problems. In this sense, it is not easy to find relevant examples in the literature that carry out a detailed EA characterization. The typical approaches consist in analysing the behaviour of particular EAs when applied to a reduced benchmark set that is taken as representative by the authors, like in [10,19,47]. An in depth analysis of the benchmark functions' fitness landscapes is not usually performed nor are the results of the algorithms' performance related to them except in very general terms. As a consequence the behaviour of the EAs cannot be understood in depth. Thus, the first step to achieve an objective and informative characterization of an EA, is to select the features of the real-parameter optimization landscapes that are most relevant in regards to its performance. The selection of this set of features and the experimental confirmation of their relevance by relating them to the performance of different algorithms is the main objective of this work.

This is not the first attempt to characterize the behaviour of an EA based on the topographical features of the landscapes. Most of the existing approaches are focused on the development of difficulty measures of the fitness landscapes [29,45]. One of the first attempts was the "Fitness Distance Correlation" (FDC) [34] measurement. It was initially developed to analyse binary-coded fitness landscapes and it was later extended for real-coded problems [69]. In spite of its first successful results, it has been demonstrated that it is not a reliable measure [1,50,35]. Related with the concept of epistasis or separability, several measures have been developed. In [30] the authors analyse the epistasis of a problem using Walsh series while in [52,51] an ANOVA analysis is carried out with the same objective. Davidor has proposed the "epistasis variance" measurement [13] to compute the hardness of a problem based on the relationships between genes but, the same author, in a later paper [14], showed the lack of reliability of this measure. Another remarkable hardness measurement based on epistasis is the epistasis correlation [53], which, as in the case of the epistasis variance, presents difficulties to detect the absence of epistasis [44]. Following this line, the work of Borenstein and Poli must be pointed out [5]. They propose a new landscape hardness measure based on the "information landscape" concept. It allows to explicitly measure the amount of information present in a landscape through a simple notion of distance and to predict the performance of a search algorithm over it through an empirical method. This method is based on the comparison between a landscape and a reference optimal one. The authors apply it only to study the behaviour of the genetic algorithm with binary coded genotype and unimodal functions, but the background methodology used to analyse the fitness landscapes is general enough to be applied to other evolutionary algorithms, which has not been done yet. A similar work is presented in [67] where the authors estimate the hardness of a fitness landscape based on temporal series obtained by "adaptive walks" over the landscape using the operators of the algorithm under analysis. The hardness of a landscape is measured by means of the regularity of the "adaptive walks".

The general conclusion of this review is that, among the commented drawbacks displayed for some of the hardness measurements, like the FDC or epistasis variance, which were developed in the field of evolutionary algorithms are based on the analysis of a specific EA without allowing the generalization of the results to other EAs. For example, in the work presented in [67], the hardness measure is calculated using the operators of the analysed EA so for each variation of the selected EA it would be necessary to calculate it again. The methodology proposed in this work is not based on a specific difficulty measure but on the topographical features of the fitness landscapes which are the same for every EA.

The remainder of the paper is structured as follows: Section 2 is devoted to the description of the fitness landscape features selected as relevant for EA performance. Section 3 contains the particular experimental setup designed to show the practical relevance of the previously selected landscape features. Section 4 is focused in the discussion of the experimental results obtained in the characterization of three well-known EAs and, finally, Section 5 includes the general conclusions of the paper and its main contributions.

2. Fitness landscape analysis

The procedure followed in this paper in order to decide on an informative characterization of EAs starts by considering the different fitness landscapes that may be found in different problems and deciding how they should be described or classified, that is, in terms of what features. In this sense, there is an extensive literature on EAs in real-parameter optimization problems, which have been thoroughly studied in the evolutionary computation field due to their appearance in most real cases. After revising the literature on this topic, there seems to be a consensus on a basic set of four general features that can be chosen as the most relevant for characterizing the fitness landscape of any problem: epistasis or separability [13], modality [49,25], ruggedness [70,71] and deceptivity [16,15]. These four features are not completely independent, in fact, the set can be reduced to only two, separability and modality, as ruggedness is closely related to modality [57] and deceptivity to separability [46]. This selection is not new, and in typical EA competitions, where the benchmark sets are carefully designed [59], the functions are classified in terms of separability and modality, typically in a binary fashion, using separable/non-separable and unimodal/multimodal categories. The main contribution of this work is to show empirically that such a classification could produce misleading conclusions about the performance of the EAs and a more detailed one is required to achieve a really useful EA characterization.