European Journal of Operational Research 253 (2016) 337–355

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Discrete Optimization

Modified Differential Evolution with Locality induced Genetic Operators for dynamic optimization

Rohan Mukherjee[a], Shantanab Debchoudhury[b], Swagatam Das[c,*]

[a] Department of Computer Science, Rice University, USA
[b] Electrical Engineering, Virginia Polytechnic Institute and State University, USA
[c] Electronics and Communication Sciences Unit, Indian Statistical Institute, Kolkata 700108, India

ARTICLE INFO

Article history:
Received 4 December 2014
Accepted 27 February 2016
Available online 4 March 2016

Keywords:
Continuous optimization
Dynamic optimization
Differential Evolution
Self adaptation
Genetic operators

ABSTRACT

This article presents a modified version of the Differential Evolution (DE) algorithm for solving Dynamic Optimization Problems (DOPs) efficiently. The algorithm, referred as Modified DE with Locality induced Genetic Operators (MDE-LiGO) incorporates changes in the three basic stages of a standard DE framework. The mutation phase has been entrusted to a locality-induced operation that retains traits of Euclidean distance-based closest individuals around a potential solution. Diversity maintenance is further enhanced by inclusion of a local-best crossover operation that empowers the algorithm with an explorative ability without directional bias. An exhaustive dynamic detection technique has been introduced to effectively sense the changes in the landscape. An even distribution of solutions over different regions of the landscape calls for a solution retention technique that adapts this algorithm to dynamism by using the previously stored information in diverse search domains. MDE-LiGO has been compared with seven state-of-the-art evolutionary dynamic optimizers on a set of benchmarks known as the Generalized Dynamic Benchmark Generator (GDBG) used in competition on evolutionary computation in dynamic and uncertain environments held under the 2009 IEEE Congress on Evolutionary Computation (CEC). The experimental results clearly indicate that MDE-LiGO can outperform other algorithms for most of the tested DOP instances in a statistically meaningful way.

1. Introduction

Optimization problems in the real world are very often dynamic in nature. For these Dynamic Optimization Problems (DOPs), the function landscape changes temporally i.e. optima of the problem to be solved change their locations over time and thus, the optimizer should be able to track the optima continually by responding to the dynamic environment (Jin & Branke, 2005; Nguyen, Yang, & Branke, 2012). Variations in market price, probabilistic arrival of a new job in a scheduling problem, uncertainty on the demand and unit costs in a network in context to the facility location problem are some of the common instances of a dynamic environment (Gabrel, Murat, & Thiele, 2014). Under such situations, converging tendency of a conventional EA can impose severe limitations on the performance of the EA. If the population members of the EA converge rapidly, they will not succeed in effectively responding to the environmental changes. Therefore, in case of DOPs the foremost challenge is maintenance of a diverse population and simultaneous production of highly accurate solutions by monitoring the moving optima. At this point we would like to mention that there are also DOP instances where the optimal solution does not need to be tracked. For example, the work of Allmendinger and Knowles (Allmendinger & Knowles, 2010) investigates DOPs where the constraints (Ephemeral Resource Constraints (ERCs)) change over time but not the landscape and thus, also not the optimal solutions. In this paper we focus on the real-parameter bound-constrained DOPs where the objective function landscape explicitly changes with time and not on the problems with ERCs. However, the proposed algorithm is not designed for DOPs with other (more practical) equality and inequality constraints on the decision variables. The benchmark suite used for testing the algorithm also comprises of bound-constrained problems only (i.e. each decision variable is bounded from above and below).

Differential Evolution (DE) (Das & Suganthan, 2011; Storm & Price, 1997) has been used effectively as one of the most powerful optimizing tool for continuous search spaces. DE implements similar computational steps to that of standard Evolutionary Algorithms (EAs). However, unlike traditional EAs, DE-variants perturb the current-generation population members with the scaled differences of randomly selected and distinct population

* Corresponding author. Tel.: +91 3325752323; fax: +91 3325752323.
E-mail addresses: rm38@rice.edu (R. Mukherjee), sdch10@vt.edu (S. Debchoudhury), swagatamdas19@yahoo.co.in, swagatam.das@isical.ac.in (S. Das).

http://dx.doi.org/10.1016/j.ejor.2016.02.042
0377-2217/

members. Therefore, no separate probability distribution (like the Gaussian distributions used in Evolutionary Programming (EP) and Evolution Strategies (ES) or the Cauchy distributions used in case of the Fast EPs) is used to generate offspring.

Classical DE suffers from some difficulties in its application in DOPs owing to two main factors. First, in many cases, local basins of attraction covering areas in and around the local and global optima result in a premature convergence. Thereafter, explorative power is compromised due to similarity of the minimally perturbed new optimum in a changed environment. Second, DE may occasionally stop proceeding toward the global optimum even though the population has not converged to a local optimum or any other point (Lampinen & Zelinka, 2000). Ongoing research has been directed to introduce modifications in DE algorithms to locate the changing optima for dynamic landscapes. A short summary of the relevant literature has been presented in Section 2.

In this paper we present an alternative approach of solving DOPs by using a modified DE algorithm with locality induced genetic operators. Our proposal is based on the fact that any dynamic change is reflected by the varying dominance of the candidate solutions over periods of change instances. Thus, attempts have been made to monitor each such solution over the entire period of optimization. The retention of the spatial traits characterizing these solutions forms an essential part of the genetic operations (mutation and crossover) in the proposed algorithm called Modified DE with Locality induced Genetic Operators (MDE-LiGO). The crossover phase of the genetic operation stage is handled by an adaptive *l*-best scheme that allows a rotation of the trial solutions and can attain a compromise between both axis parallel search and rotation-invariant search. The modification of DE is coupled with two additional features—first dynamic detection using a scheme to measure the number of unsuccessful updates and second, adapting to that change using clustering techniques. Diversity maintenance is an essential feature of MDE-LiGO.

Although a preliminary version of this article has been presented as a conference paper in Mukherjee, Debchoudhury, Kundu, Das, and Suganthan (2013), we have substantially modified and expanded it both in terms of the algorithmic features and the experimental analyses. Unlike the conference version, parameter adaptation is introduced to enable a control on the retention of traits around a promising solution. Consequently higher percentages of traits that have the ability to lead to a potential solution are identified and retained, thereby increasing the efficiency of the process. The blending rate *Br* has been selected from a pool of values sampled from a normal distribution, the parameters of which are guided by the influence of a set of successful crossover probabilities. This modification appears to make the crossover stage more functional and much more effective. A significant contribution added to this version is the introduction of an exhaustive dynamic detection stage that identifies the onset of a dynamic change. The difference in deviation of locally mutated DE individuals serves as the criterion which dictates when a change in the landscape has occurred. An analytical discussion on the evaluation of such a deviation mechanism has been presented in context to the detection of a dynamic change by MDE-LiGO. In addition extensive comparisons and experimental validations have been provided to validate the different components of the MDE-LiGO algorithm.

Organization of the rest of the paper is in order. Section 2 provides a brief description of classical DE and one of its adaptive variant. The section also presents a compact survey of the different modified EAs previously used for solving DOPs. A detailed description of the proposed algorithm with all its salient features is provided in Section 3. Section 4 describes the experimental settings and presents the results of comparing MDE-LiGO against seven state-of-the-art dynamic EAs with in-depth discussions. Section 5 experimentally investigates the effect of different strategies proposed for the MDE-LiGO framework. Finally conclusions are drawn in Section 6.

2. Related works

It was in 1966 (Fogel, Owens, & Walsh, 1966) that the earliest known attempts were made to apply EAs for solving DOPs. Since the late 1980s, the topic started to attract a lot of attention from the researchers. What followed was a gradual increase in publication of related works. Expansive survey works dedicated to the application and adaptation of EAs for tackling DOPs can be found in Jin and Branke (2005) and Nguyen et al. (2012).

A series of innovative approaches were adopted by the research fraternity in order to adapt an EA to solve DOPs. For example, the hypermutation strategy (Cobb, 1990) incorporates diversity after detection of a change in fitness landscape by escalating the rate of mutation for some generations following the dynamic change. Morrison and de Jong (2000) pointed out that when the functional landscape is changing at a high frequency with time, increasing the mutation rate more frequently can be beneficial for efficient tracking of the optima. On the contrary, the performance of lower hypermutation levels is better in case of changes that are less frequent. In variable local search (Vavak, Jukes, & Fogarty, 1997) the rate of mutation is slowly incremented. These results in randomization of the information associated with individuals that have been successful before, to maintain adequate diversity of the population. Under the random immigrant scheme (Grefenstette, 1992), a part of the population is replaced by randomly generated individuals in each generation to maintain the diversity level throughout the search process. Strategies that are adopted to maintain diversity in DOPs include fitness sharing and crowding (Cedeno & Vemuri, 1997) .Dividing the entire population to smaller subgroups often help to track multiple peaks, thereby functioning like self adaptive diverse memory. The multinational Genetic Algorithm (GA) (Ursem, 2000), the shifting balance GA (Winberg & Oppacher, 2000), and the self-organizing scouts (Branke, Kaußler, Smidt, & Schmeck, 2000; Branke & Schmeck, 2003) are examples of this technique. In a dynamic optimization scenario, it may be advantageous to store the information from a previous generation in an external memory and later recall the same. Such external memory can be coupled with GAs to solve DOPs.

Two types of memories have so far been used in conjunction with the dynamic evolutionary optimizers—*explicit memory* and *implicit memory*. GAs with explicit memory incorporate strategies for storing solutions and reintroducing them during the later stages of search (Eggermont & Lenaerts, 2000; Louis & Xu, 1996; Ramsey & Greffenstette, 1993). On the other hand, GAs with implicit memory use redundant genetic representations. The most common example is employing a diploid genetic structure (Calabretta, Galbiati, Nolfi, & Parisi, 1996; Ng & Wong, 1995). A diploid GA possesses two sets of chromosomes instead of a common single set (haploid) possessed by the regular GAs. Consequently, in this type of GAs, two genes compete for the same phenotypic feature in the same individual. To resolve this dilemma, the researchers use a dominance mapping that labels some genes as dominant and others as recessive. When a dominant gene is paired with a recessive one, only the former is expressed in the phenotype, leaving the recessive one unexpressed. This way, the dominant genes can protect their less fit, recessive counterparts from being eliminated through selection. The formerly fit genes can return back by pairing with the fitter dominant genes and then may come into expression again when the environment becomes more favorable. Through this mechanism, the GA obtains a form of implicit memory. Lewis, Hart, and Ritchie (1998) indicated that a diploid structure alone is not enough for a diploid GA to