

**Investigating Restricted Tournament Replacement  
in ECGA for Non-Stationary Environments**

**Claudio F. Lima, Carlos Fernandes,  
and Fernando G. Lobo  
IlliGAL Report No. 2008010  
May, 2008**

Illinois Genetic Algorithms Laboratory  
University of Illinois at Urbana-Champaign  
117 Transportation Building  
104 S. Mathews Avenue Urbana, IL 61801  
Office: (217) 333-2346  
Fax: (217) 244-5705

# Investigating Restricted Tournament Replacement in ECGA for Non-Stationary Environments

Claudio F. Lima<sup>1,2</sup>, Carlos Fernandes<sup>3,4</sup>, and Fernando G. Lobo<sup>1</sup>

<sup>1</sup>Informatics Laboratory (UALG-ILAB), DEEI-FCT  
University of Algarve, Faro, Portugal  
{clima.research,fernando.lopez}@gmail.com

<sup>2</sup>Illinois Genetic Algorithm Laboratory (IlliGAL)  
University of Illinois at Urbana-Champaign, Urbana, USA

<sup>3</sup>LaSEEB  
Technical University of Lisbon, Lisbon, Portugal  
cfernandes@laseeb.org

<sup>4</sup>Departamento ATC  
University of Granada, Granada, Spain

May 2, 2008

## Abstract

This paper investigates the incorporation of restricted tournament replacement (RTR) in the extended compact genetic algorithm (ECGA) for solving problems with non-stationary optima. RTR is a simple yet efficient niching method used to maintain diversity in a population of individuals. While the original version of RTR uses Hamming distance to quantify similarity between individuals, we propose an alternative substructural distance to enforce the niches. The ECGA that restarts the search after a change of environment is compared with the approach of maintaining diversity, using both versions of RTR. Results on several dynamic decomposable test problems demonstrate the usefulness of maintaining diversity throughout the run over the approach of restarting the search from scratch at each change. Furthermore, by maintaining diversity no additional mechanisms are required to detect the change of environment, which is typically a problem-dependent and non-trivial task.

## 1 Introduction

The extended compact genetic algorithm (ECGA) (Harik, 1999; Harik, Lobo, & Sastry, 2006) replaces the standard crossover and mutation operators by building a probabilistic model of promising solutions and sampling from the corresponding probability distribution. This feature allows ECGA and other advanced estimation of distribution algorithms (EDAs) (Larrañaga & Lozano, 2002; Pelikan, Goldberg, & Lobo, 2002) to automatically identify the problem decomposition and important problem substructures, leading to superior performance for many problems when compared with EAs that use fixed, problem-independent variation operators.

Despite the growing interest in EDAs, the application of these algorithms to dynamic optimization problems (DOPs) has been scarce (Ghosh & Muehlenbein, 2004; Abbass, Sastry, & Goldberg, 2004; Sastry, Abbass, & Goldberg, 2004; Yang, 2005c; Yang, 2005b; Yang, 2005a). Moreover, most of these works have been limited to univariate EDAs (Ghosh & Muehlenbein, 2004; Yang, 2005c; Yang, 2005b; Yang, 2005a), that use a probabilistic model of simple and fixed structure which assumes each variable to be independent from every other. The relevance of applying more powerful EDAs has been highlighted in recent work (Abbas, Sastry, & Goldberg, 2004; Sastry, Abbass, & Goldberg, 2004), where the ECGA was modified in order to tackle problems with dynamic environments. While the previous approach is based on increasing diversity after a change, in this paper we focus on maintaining diversity throughout the run, removing the need of predicting the changes of environment.

More precisely, we investigate restricted tournament replacement (RTR) (Harik, 1995) as the source of continuous diversity in the ECGA. The original version of RTR uses genotypic distance to promote diversity in the population. Besides testing the traditional RTR, we propose a substructural distance that exploits model structural information. We compare the ECGA that restarts the population with the approach of maintaining diversity, testing both versions of RTR, when solving several dynamic decomposable problems of bounded difficulty (Abbas, Sastry, & Goldberg, 2004; Yang, 2004).

The paper is structured as follows. The next section provides a brief introduction to the ECGA, while Section 3 analyzes the application of ECGA to DOPs. Section 4 presents the results obtained and corresponding discussion, and Section 5 focuses on the generality of application of the proposed approach to DOPs. The paper ends with major conclusions.

## 2 Extended Compact Genetic Algorithm

Estimation of distribution algorithms (Larrañaga & Lozano, 2002; Pelikan, Goldberg, & Lobo, 2002; Pelikan, 2005) replace traditional variation operators of genetic algorithms (GAs) by building a probabilistic model of promising solutions (that survive selection) and sampling the corresponding probability distribution to generate the offspring population. The extended compact genetic algorithm (eCGA) (Harik, 1999; Harik, Lobo, & Sastry, 2006) uses a product of marginal distributions on a disjoint partition of variables of the problem to sample new solutions,

$$p(X) = \prod_{i=1}^m p(X_{I_i}), \quad (1)$$

where  $X = (X_1, X_2, \dots, X_\ell)$  is a vector that contains all the variables of the problem and  $I_i$  is the index set that contains the index of the variables that belong to the  $i^{th}$  marginal distribution.

This kind of probability distribution belongs to a class of probabilistic models known as marginal product models (MPMs). For example, the following MPM, [1,3] [2] [4], for a 4-bit problem represents that the 1<sup>st</sup> and 3<sup>rd</sup> variables are linked, and the 2<sup>nd</sup> and 4<sup>th</sup> variables are independent.

In ECGA, both the structure and the parameters of the model are searched and optimized to best fit the data (promising solutions). A greedy MPM search is performed at each generation. The greedy search algorithm starts with the simplest possible model, assuming that all variables are independent, and then keeps merging groups of variables whenever a certain score metric is improved. This process goes on until no further improvement is possible.

The score metric used in ECGA is the minimum description length (MDL) metric (Rissanen, 1978), that penalizes both inaccurate and complex models, thereby leading to a (locally) optimal distribution. More precisely, the MPM complexity is given by the sum of model complexity,  $C_m$ ,

and compressed population complexity,  $C_p$ . The model complexity,  $C_m$ , quantifies the model representation in terms of the number of bits required to store all the marginal probabilities,

$$C_m = \log_2(n + 1) \sum_{i=1}^m (2^{k_i} - 1), \quad (2)$$

where  $n$  is the population size,  $m$  is the number of linkage groups, and  $k_i$  is the size of the  $i^{\text{th}}$  group.

The compressed population complexity,  $C_p$ , which quantifies the data compression in terms of the entropy of the marginal distribution over all partitions, is given by

$$C_p = n \sum_{i=1}^m \sum_{j=1}^{2^{k_i}} -p_{ij} \log_2(p_{ij}), \quad (3)$$

where  $p_{ij}$  is the frequency of the  $j^{\text{th}}$  gene sequence of the genes belonging to the  $i^{\text{th}}$  partition.

After learning the probabilistic model, the offspring population is generated by randomly sampling substructures, according to the probabilities stored in the MPM. More details about the pseudo-code of ECGA are given in the next section.

### 3 ECGA for Non-Stationary Environments

In this section we review previous efforts on using ECGA to solve DOPs and introduce a new methodology to tackle these class of problems by maintaining diversity along the run.

#### 3.1 Previous Work

Dynamic optimization problems have attracted significant interest from the evolutionary computation community in the last two decades (Branke, 2001; Jin & Branke, 2005). More recently, EDAs have been also applied to DOPs, although a significant part of the work has been focused in simpler univariate models (Ghosh & Muehlenbein, 2004; Yang, 2005c; Yang, 2005b; Yang, 2005a).

Abbass, Sastry, and Goldberg (Abbass, Sastry, & Goldberg, 2004) were the first to introduce ECGA (Harik, 1999; Harik, Lobo, & Sastry, 2006) to problems with dynamic environments. Their approach is based on random restarts of the population at each change so that diversity in the population can be increased at the beginning of each new environment. This work demonstrated the utility of learning possible structural decompositions in changing environments.

The outline of the algorithm can be viewed in Figure 1. The algorithm is similar to the original ECGA, but in order to deal with DOPs, diversity is increased by resetting the population when a change of environment is detected. The change of environment was assumed to be known, so that the authors could focus on the ECGA response behavior to new environments when the population is restarted. Additionally, Abbass et al. proposed a slightly different approach that used the MPM from the previous generation when the population was to be restarted. Although for problems where the structure does not change with the environment this procedure is advantageous, when the environment change is related to structural changes this procedure can be harmful. Since we are interested in studying both fitness and structural changes in DOPs, we use the more general version for comparison with our approach.

The ECGA with random restart was later extended (Sastry, Abbass, & Goldberg, 2004) to include substructural niching (Sastry, Abbass, Goldberg, & Johnson, 2005). In substructural niching, niches are defined within the linkage groups rather than at the individual level. After the corresponding schema average fitness is calculated for each substructure (Sastry, Pelikan, & Goldberg,

---

***ECGA with random restart***

---

- (1) Create a random population  $P$  of  $n$  individuals.
  - (2) Evaluate population  $P$ .
  - (3) If an environment change is detected:
    - (3.1) Re-initialize population  $P$  at random.
    - (3.2) Evaluate population  $P$ .
  - (4) Select  $P'$  individuals from  $P$  by tournament selection.
  - (5) Model the selected individuals  $P'$  using the greedy MPM search procedure.
  - (6) Sample a new population  $O$  according to the MPM learned in step 4.
  - (7) Replace all individuals in population  $P$  by those from  $O$ .
  - (8) If stopping criteria is not satisfied, return to step 2.
- 

Figure 1: Steps of the extended compact genetic algorithm (ECGA) with random restart of the population for non-stationary environments.

2004), the sampling probabilities are changed according with their associated fitness. While this methodology can be used to maintain diversity in the population, Abbass et al. used as a way to accelerate the growth ratio of highly-fit substructures, using the random restart of the population as the main source of diversity.

Our approach is to maintain diversity throughout the run so that ECGA can be able to respond to a change without explicitly needing additional methods to detect new environments.

### 3.2 Restricted Tournament Replacement

When handling DOPs with evolutionary algorithms (EAs), the main approaches to deal with premature convergence can be roughly grouped into four categories (Jin & Branke, 2005): increase diversity after an environment change, maintain diversity throughout the run, memory-based approaches, and multi-population approaches. While previous efforts with ECGA (Abbas, Sastry, & Goldberg, 2004; Sastry, Abbass, & Goldberg, 2004) have been done in the first category, this paper follows the second approach by maintaining diversity in the population with RTR.

The restricted tournament replacement (RTR) (Harik, 1995) is a niching method that has been used successfully in the hierarchical Bayesian optimization algorithm (hBOA) (Pelikan, 2005). In RTR, each new solution  $X$  is incorporated in the original population using the following procedure:

1. Select a random subset of individuals  $W$  with size  $w$  from the original population.
2. Let  $Y$  be the solution from  $W$  that is most similar to  $X$ , in terms of genotypic distance.
3. Replace  $Y$  with  $X$ , if  $X$  is better, otherwise discard  $X$ .

The window size  $w$  is set to the problem size, as suggested elsewhere (Pelikan, 2005). The ECGA that uses RTR to address DOPs is outlined in Figure 2. After sampling the offspring population,

- (1) Create a random population  $P$  of  $n$  individuals.
  - (2) Evaluate population  $P$ .
  - (3) Select  $P'$  individuals from  $P$  by tournament selection.
  - (4) Model the selected individuals  $P'$  using the greedy MPM search procedure.
  - (5) Sample a new population  $O$  according to the MPM learned in step 4.
  - (6) If an environment change is detected, reevaluate population  $P$ .
  - (7) Insert individuals from population  $O$  back into  $P$  using the restricted tournament replacement.
  - (8) If stopping criteria is not satisfied, return to step 3.
- 

Figure 2: Steps of the extended compact genetic algorithm (ECGA) with restricted tournament replacement (RTR) for non-stationary environments.

the new individuals are incorporated one by one back into the original population using RTR. To make a fair comparison with the restart approach, we first assume that we know when the change of environment occurs, but later on (Section 5) we remove this assumption to further demonstrate the usefulness of maintaining diversity in the population. At each change, the original population is reevaluated before RTR takes place, so that fitness values are updated and subsequent comparisons are made correctly.

As mentioned before, RTR uses genotypic distance to measure similarity between individuals, which for the case of binary-coded strings is simply the Hamming distance. We propose an alternative similarity metric based on substructural distance. The substructural distance between two individuals is defined as the number of substructures in which they differ, according to the MPM structure.

Consider the following example, with the MPM [1,2,3,4] [5,6,7,8] [9,10,11,12]:

<b>Offspring</b>	1111 1111 1111		
<b>Parent1</b>	0111 1011 1110	$d_1 = 3$	$d_2 = 3$
<b>Parent2</b>	0000 1111 1111	$d_1 = 4$	$d_2 = 1$

The individual **Offspring** is compared with two parents ( $w = 2$ ) to decide to which to compete for a place in the population. Using the traditional similarity measure of RTR, the Hamming distance ( $d_1$ ), the closest parent is **Parent1**. However, if we use substructural distance ( $d_2$ ), **Parent2** is the closest individual to **Offspring**. While **Parent2** is  $d_1 = 4$  bits away from **Offspring**, the different bits belong all to the same substructure, meaning that these individuals only differ in one substructure. Therefore, and substructural-wise, these two individuals belong to the same niche.

By using substructural distance in RTR, the niches are maintained at the substructural level rather than at the individual level. In this way, the problem decomposition capability of EDAs is also exploited for diversity maintenance. This method is based on the same principle of substructural niching (Sastry, Abbass, Goldberg, & Johnson, 2005). However, in that work explicit




Figure 3: (a) Trap, (b) modified trap, and (c) switching trap subfunctions.

substructural fitness information is used, while in our case the choice between good and poor substructures is implicitly made through the individual’s overall fitness.

In the next section, we present the results obtained for several test functions with both versions of RTR in ECGA and compare them to the restart approach (Abbass, Sastry, & Goldberg, 2004).

## 4 Experiments

### 4.1 Experimental Setup

In this paper, we use the same experimental setup as used elsewhere (Abbass, Sastry, & Goldberg, 2004) to be able to compare the different approaches. Abbass et al. (Abbass, Sastry, & Goldberg, 2004) used dynamic versions of adversarial problems with bounded difficulty (Goldberg, 2002; Yang, 2004), where the identification of the underlying problem substructure is crucial for GA success. Obviously, the solver should not have any previous knowledge about the problem structure. One such class of problems are the  $m$ - $k$  additively decomposable problems (Goldberg, 2002), which consists of  $m$  additively separable functions of size  $k$  each, resulting in a total problem size of  $\ell = m \cdot k$ . The fitness is simply given by the sum of the fitness contribution of all subfunctions. Like in (Abbass, Sastry, & Goldberg, 2004), experiments are performed for three different subfunctions. The subfunctions are defined over unitation  $u$  (number of ones) and can be visualized in Figure 3. The changes in the environment are assumed to be cyclic, where two different cycles of length 5 and 10 generations are tested.

The first function is the well-know deceptive trap function (Ackley, 1987; Deb & Goldberg, 1993) adapted to dynamic environments by switching between the two versions plotted in Figure 3 (a). At the beginning of the run and for odd cycles, the fitness is calculated according to the standard trap function, with the optimal solution at 1111 and the local optimum at 0000. For even cycles, the trap is inverted so that the optimal and local optima switch their location. While the function plotted is for  $k = 4$ , we test this function with different subfunction sizes,  $k = 3, 4, 5$ , to analyze the behavior of the algorithms for increasing problem difficulty.

The second test function is a modified version of the previous trap function to break the symmetry in the attractors (Abbass, Sastry, & Goldberg, 2004). This function can be visualized in Figure 3 (b). In this case, the location of both local and optimal solutions for odd cycles does not tell anything about the optima location for the even cycles and vice-versa, therefore the algorithms cannot take advantage from previous local optima information.

The third function, called switching trap (Abbass, Sastry, & Goldberg, 2004), promotes a more severe change between environments by changing both optima location and subfunction size. As

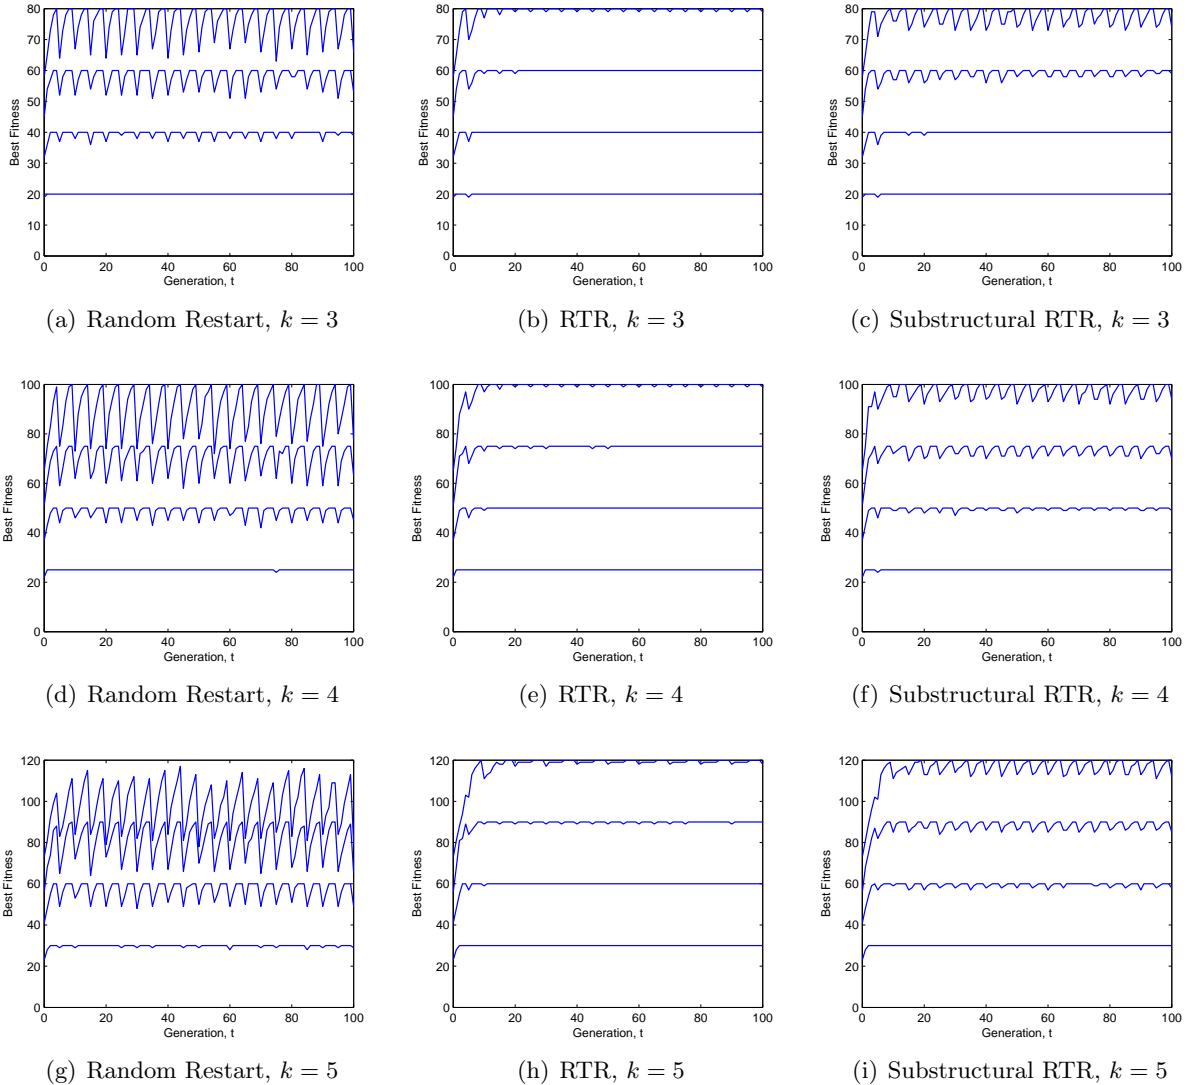



Figure 4: Results for the trap problem with  $k = 3, 4, 5$ ,  $m = 5, 10, 15, 20$ , and  $t = 5$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).

can be seen in Figure 3 (c), the environment changes between a modified trap  $k = 3$  and an inverted modified trap with  $k = 4$ .


For all experiments, the population size is set to  $n = 5000$  and the tournament selection to  $s = 16$  (Abbass, Sastry, & Goldberg, 2004). Each run terminates when the algorithms reach  $t = 100$  generations. All the results presented through the paper are averaged over 30 independent runs.

## 4.2 Results and Discussion


Figures 4 and 5 present the results for the dynamic version of the trap function (Figure 3(a)). Clearly, ECGA that maintains diversity performs much better than the restart approach. Although for  $k = 3, 4$ , the restart approach is able to find the optimum for all problem sizes, with  $k = 5$  it cannot reach the global solution for the larger problem. Even for  $t = 10$ , where enough time is given,




(a) Random Restart,  $k = 3$




(b) RTR,  $k = 3$




(c) Substructural RTR,  $k = 3$




(d) Random Restart,  $k = 4$




(e) RTR,  $k = 4$




(f) Substructural RTR,  $k = 4$



(g) Random Restart,  $k = 5$



(h) RTR,  $k = 5$



(i) Substructural RTR,  $k = 5$

Figure 5: Results for the trap problem with  $k = 3, 4, 5$ ,  $m = 5, 10, 15, 20$ , and  $t = 10$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).






Figure 6: Results for the modified trap problem with  $k = 4$ ,  $m = 5, 10, 15, 20$ , and  $t = 5$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).




Figure 7: Results for the modified trap problem with  $k = 4$ ,  $m = 5, 10, 15, 20$ , and  $t = 10$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).




Figure 8: Results for the switching trap problem with  $k = 3 \leftrightarrow 4$ ,  $\ell = 24, 48, 72, 96$ , and  $t = 5$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).




Figure 9: Results for the switching trap problem with  $k = 3 \leftrightarrow 4$ ,  $\ell = 24, 48, 72, 96$ , and  $t = 10$ . Comparison between ECGA with random restart (left), RTR (middle), and substructural RTR (right).

the algorithm gets stuck in local optima. Eventually, ECGA with the restart approach could have solved this problem, but would require a larger population size to do so. On the contrary, ECGA with RTR benefits from diversity to require smaller population sizes for correct model building.

However, different behavior can be observed for the two RTR approaches. While the standard RTR responds quite fast to the environment change, keeping most of the times the global solution for both environments in the population, the substructural RTR takes slightly more time to respond. Nevertheless, the standard RTR does not scale as well as its substructural counterpart regarding problem difficulty, adjusted through the subfunction size  $k$ . This can be seen for  $k = 5$  and  $m = 20$ , where it fails to obtain the best solution for the first environment. Indeed, allowing more time between environment changes ( $t = 10$ ) deteriorate the performance.

While convergence is delayed by using standard RTR, not all local optima can be preserved. Note that for a problem with  $m$  subfunctions there are a total of  $2^m$  local optima. On the contrary, when using substructural RTR the maintenance of substructures is more tractable since only  $2 \cdot m$  different locally optimal substructures need to be preserved.

The results for the modified trap with  $k = 4$  are shown in Figures 6 and 7. The function setup is now adequate to break the symmetry of the attractors (Abbass, Sastry, & Goldberg, 2004). Again, ECGAs that use RTR outperform the restart approach. However, for  $t = 10$ , the standard RTR fails to obtain the optimal solution for the second environment. Contrary to the previous problem, substructural RTR is now able to maintain the global optima after reaching it, but does not suffer from the difficulties of the standard RTR, when the time between changes increases to  $t = 10$ .

The last function considered switches between two modified trap functions of different size, the first with  $k = 3$  and the second with  $k = 4$ . The problem sizes  $\ell = 24, 48, 72, 96$  are chosen to be compatible with both subfunction sizes. Given that for the same total problem size, a switch in the environment means a change in the number of subfunctions, the fitness of the optimal solution for each environment differs, as can be easily verified in Figures 8 and 9. For the ECGA with restart, the outcome is similar to previous problems, the algorithm being slower than diversity-based approaches, and failing to get the optimal solution for  $t = 5$  and  $\ell = 96$ . The behavior of both RTR approaches is somewhat similar, responding relatively fast to the change of environment.

## 5 Black-Box Optimization in Non-Stationary Environments

Black-box optimization algorithms need little or no information about a problem when solving it. These type of algorithms has often been regarded as incompatible with no free lunch (NFL) theorems (Wolpert & Macready, 1997). However, in practice we are not interest in solving all possible problems, but a much narrow subset that focus on a particular domain of interest (or problems under certain bounds of difficulty). In fact, evolutionary algorithms (EAs) themselves are more adequate for problems where some sort of structure can be exploited, whether by recombining different good solutions or through small variations of good solutions (mutation). EDAs are no exception and take advantage of their mechanisms in problems where identifying the problem decomposition and important problem substructures is crucial to solve them efficiently. For example, consider a problem with size  $\ell = 100$  solely composed by a trap function with  $k = 100$ . For this problem, random search will be on average as good as any EDA or EA in finding the global optimum (obviously without previous knowledge on the problem), simply because there is no decomposition or substructure that can be exploited.

The same argument can be made for environment changes in DOPs. If the number of environments or the period between changes varies unboundedly, on average no method (without previous knowledge) will outperform the random restart of the population at each change. Therefore, we

recognize that diversity maintenance approaches are more suitable if these changes are bounded in a certain way. On the other hand, using a restart approach to ensure diversity requires the solver to have an efficient method to detect changes, which often depends on the problem itself and might not be a trivial task. On one way or the other, there is some cost associated with each approach, and in some sense a particular domain of application.

In the previous section we have assumed to know the occurrence of environment changes, as was done in previous related work (Abbass, Sastry, & Goldberg, 2004; Sastry, Abbass, & Goldberg, 2004). Here, we remove this assumption to further demonstrate the utility of using RTR to maintain diversity. Specifically, we make a simple change to the ECGA described in Figure 2. Step (6) is removed, and a new step is performed between steps (7) and (8). The step consists simply in evaluating again all individuals in population  $P$  that were not replaced by offspring at step (7). By doing this we keep the fitness of the preserved solutions updated and only need to reevaluate a proportion of the population. However, the replacement operation done before updating the fitness of the parent population will be performed with old fitness values. This is the price to pay if one doesn't want to reevaluate the entire parent population at each generation. Alternatively, if the fitness function is too sensitive to this procedure (possibly leading to the loss of previous good solutions), one may opt to evaluate the entire parent population before proceeding to RTR.

We have performed experiments for the modified and switching trap problems with both versions of RTR. Figures 10, 11, 12, and 13 illustrate similar performance to the case where the environment change was known. In some cases, even slightly better performance can be observed when the environment change is not known. This is due to the particular fitness function definition itself and the effect caused by using old fitness information when the environment changes. In other cases, this might be equally prejudicial.

The corresponding proportion of reevaluated individuals in the population is also plotted. The proportion of reevaluations increases when the niches stabilized around the substructural optima for both environments. Despite the extra cost of reevaluating a proportion of the population, the diversity maintenance reduces the population size required to solve the problems when compared with the ECGA without niching.

## 6 Conclusions

This paper addressed restricted tournament replacement (RTR) in the extended compact genetic algorithm (ECGA) to tackle dynamic optimization problems. Apart from the original version of RTR that uses Hamming distance to quantify similarity between individuals, we have proposed an alternative metric based on substructural niching (Sastry, Abbass, Goldberg, & Johnson, 2005). These approaches were compared with the ECGA that randomly restarts the population after an environment change. Results on several dynamic decomposable test problems demonstrated the usefulness of maintaining diversity in the population over the approach of restarting the search from the beginning at each change. Additionally, substructural RTR was shown to be more robust than the original version. Finally, by maintaining diversity along the run, no additional mechanisms are required to detect a change of environment, which is known to be a demanding task.

Although results are encouraging, further investigation needs to be done on the scalability of the RTR approach for increasing number of different environments. A combination of substructural niching with a restart mechanism based on model information may achieve a more robust performance. This is currently topic of ongoing work.




Figure 10: ECGA with (left) RTR and (right) substructural RTR on the modified trap problem with  $k = 4$ ,  $\ell = 20, 40, 60, 80$ , and  $t = 5$ , when the change of environment is unknown. The proportion of reevaluated individuals in the population is also shown.




Figure 11: ECGA with (left) RTR and (right) substructural RTR on the modified trap problem with  $k = 4$ ,  $\ell = 20, 40, 60, 80$ , and  $t = 10$ , when the change of environment is unknown. The proportion of reevaluated individuals in the population is also shown.

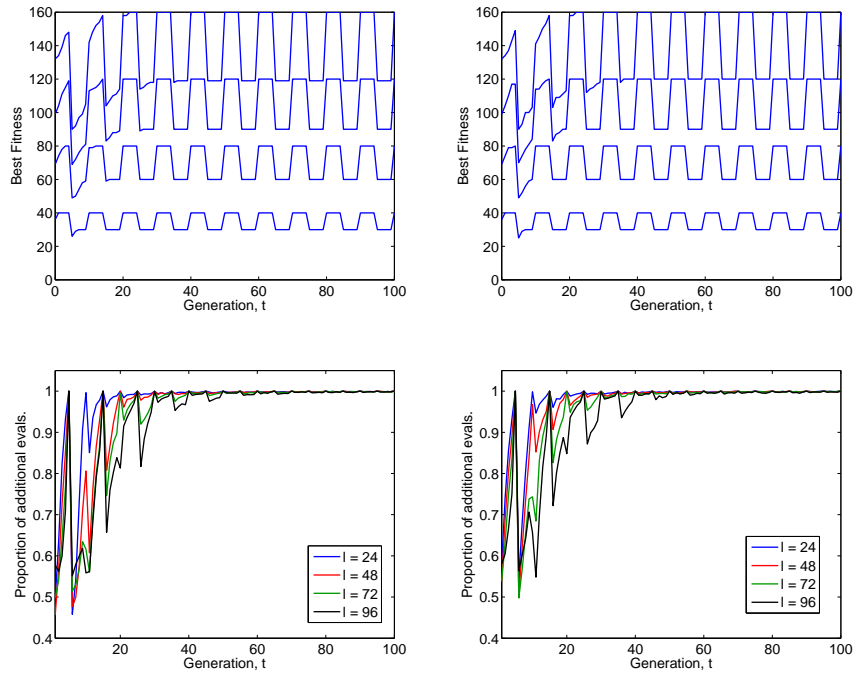


Figure 12: ECGA with (left) RTR and (right) substructural RTR on the switching trap problem with  $k = 3 \leftrightarrow 4$ ,  $\ell = 24, 48, 72, 96$ , and  $t = 5$ , when the change of environment is unknown. The proportion of reevaluated individuals in the population is also shown.

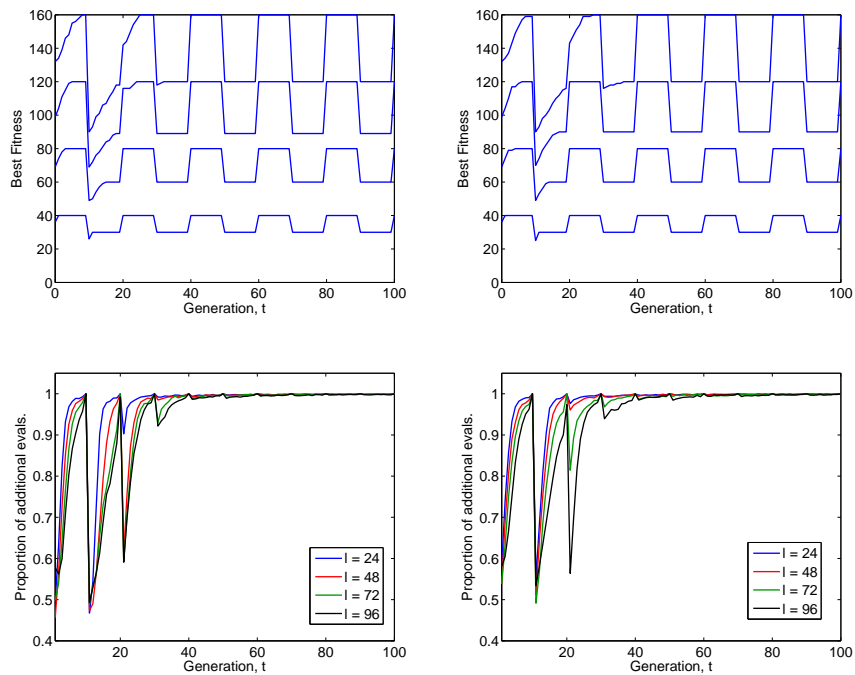


Figure 13: ECGA with (left) RTR and (right) substructural RTR on the switching trap problem with  $k = 3 \leftrightarrow 4$ ,  $\ell = 24, 48, 72, 96$ , and  $t = 10$ , when the change of environment is unknown. The proportion of reevaluated individuals in the population is also shown.

## 7 Acknowledgments

This work was sponsored by the Portuguese Foundation for Science and Technology under grants SFRH-BD-16980-2004, SFRH-BD-18868-2004, and PTDC-EIA-67776-2006.

## References

- Abbass, H. A., Sastry, K., & Goldberg, D. E. (2004). *Oiling the wheels of change: The role of adaptive automatic problem decomposition in non-stationary environments* (IlliGAL Report No. 2004029). Urbana, IL: University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory.
- Ackley, D. H. (1987). *A connectionist machine for genetic hill climbing*. Boston: Kluwer Academic.
- Branke, J. (2001). *Evolutionary optimization in dynamic environments*. Norwell, MA: Morgan Kaufmann.
- Deb, K., & Goldberg, D. E. (1993). Analyzing deception in trap functions. *Foundations of Genetic Algorithms 2*, 93–108.
- Ghosh, A., & Muehlenbein, H. (2004). Univariate marginal distribution algorithms for non-stationary optimization problems. *Int. J. Know.-Based Intell. Eng. Syst.*, 8(3), 129–138.
- Goldberg, D. E. (2002). *The design of innovation - lessons from and for competent genetic algorithms*. Norwell, MA: Kluwer Academic Publishers.
- Harik, G., Lobo, F. G., & Sastry, K. (2006). Linkage learning via probabilistic modeling in the ECGA. In Pelikan, M., Sastry, K., & Cantu-Paz, E. (Eds.), *Scalable Optimization via Probabilistic Modeling: From Algorithms to Applications* (pp. 39–61). Springer.
- Harik, G. R. (1995). Finding multimodal solutions using restricted tournament selection. *Proceedings of the Sixth International Conference on Genetic Algorithms*, 24–31.
- Harik, G. R. (1999). *Linkage learning via probabilistic modeling in the ECGA* (IlliGAL Report No. 99010). Urbana, IL: Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign.
- Jin, Y., & Branke, J. (2005). Evolutionary optimization in uncertain environments—a survey. *IEEE Transactions on Evolutionary Computation*, 9(3), 303–317.
- Larrañaga, P., & Lozano, J. A. (Eds.) (2002). *Estimation of distribution algorithms: a new tool for evolutionary computation*. Boston, MA: Kluwer Academic Publishers.
- Pelikan, M. (2005). *Hierarchical Bayesian Optimization Algorithm: Toward a new generation of evolutionary algorithms*. Springer.
- Pelikan, M., Goldberg, D. E., & Lobo, F. (2002). A survey of optimization by building and using probabilistic models. *Computational Optimization and Applications*, 21(1), 5–20. Also IlliGAL Report No. 99018.
- Rissanen, J. J. (1978). Modelling by shortest data description. *Automatica*, 14, 465–471.
- Sastry, K., Abbass, H. A., & Goldberg, D. E. (2004). Sub-structural niching in non-stationary environments. In *Proceedings of the 17th Australian Joint Conference on Artificial Intelligence, LNAI 3339* (pp. 873–885). Springer.

- Sastry, K., Abbass, H. A., Goldberg, D. E., & Johnson, D. D. (2005). Sub-structural niching in estimation distribution algorithms. In Beyer, H., et al. (Eds.), *Proceedings of the ACM SIGEVO Genetic and Evolutionary Computation Conference (GECCO-2005)* (pp. 671–678). ACM Press.
- Sastry, K., Pelikan, M., & Goldberg, D. E. (2004). Efficiency enhancement of genetic algorithms via building-block-wise fitness estimation. In *Proceedings of the IEEE International Conference on Evolutionary Computation* (pp. 720–727).
- Wolpert, D. H., & Macready, W. G. (1997). No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1(1), 67–82.
- Yang, S. (2004). Constructing dynamic test environments for genetic algorithms based on problem difficulty. In *Proceedings of the IEEE International Conference on Evolutionary Computation* (pp. 1262–1269). IEEE Press.
- Yang, S. (2005a). Experimental study on population-based incremental learning algorithms for dynamic optimization problems. *Soft Computing*, 9(11), 815–834.
- Yang, S. (2005b). Memory-enhanced univariate marginal distribution algorithms for dynamic optimization problems. In *Proceedings of the IEEE International Conference on Evolutionary Computation* (pp. 2560–2567). IEEE Press.
- Yang, S. (2005c). Population-based incremental learning with memory scheme for changing environments. In Beyer, H., et al. (Eds.), *Proceedings of the ACM SIGEVO Genetic and Evolutionary Computation Conference (GECCO-2005)* (pp. 711–718). ACM Press.