Problem Difficulty in Real-Valued Dynamic Problems

Karsten Weicker

University of Stuttgart, Institute of Computer Science,
Breitwiesenstr. 20–22, 70565 Stuttgart, Germany,
email: Karsten.Weicker@informatik.uni-stuttgart.de

Abstract. The article at hand identifies two kinds of problem difficulty in dynamic environments: discontinuities in the fitness landscape caused by moving optima and the discrepancy between tracking and optimization. These problem difficulties are supported by theoretical considerations and experimental findings. Various methods of resolution are discussed and a new adaptation method, namely the correction of strategy variables, is proposed.

1 Introduction

Recently there is an increasing interest in the application of evolutionary algorithms on dynamic, non-stationary problems. Despite of a growing number of published papers presenting new techniques for tackling those problems there is hardly any work available concerning the special requirements posed by dynamic behavior and the arising problem difficulty for certain evolutionary algorithms. Concerning the standard evolutionary algorithms, Bäck (1998) argues that self-adaptive evolution strategies (Rechenberg, 1994; Schwefel, 1995) are a particular good means to cope with dynamic problems. This work discusses the circumstances under which evolution strategies cannot cope with the dynamics and examines those circumstances empirically.

The various evolutionary standard algorithms have been extended by a wide range of techniques to tackle non-stationary problems. Examples are the addition of an external memory to remember past solution candidates (Branke, 1999b), a hypermutation operator which increases the diversity in the population (Cobb, 1990), also diploid concepts are used as an internal model to memorize partial previous solution candidates (Goldberg & Smith, 1987), ideas of thermodynamics to control selection (Mori, Kita, & Nishikawa, 1996), special local search operators to keep track of moving local optima in genetic algorithms Vavak, Fogarty, and Jukes (1996), and cultural algorithms to control the operators (Saleem & Reynolds, 2000). A more extensive overview on the whole field may be drawn from the survey by Branke (1999a). However, those techniques are out of the scope of this work which examines standard evolution strategies instead.

Preliminary experimental investigations on problem difficulties for evolution strategies in dynamic environments are contained in Weicker and Weicker (1999, 2000). There, specific problems are found to be rather difficult for evolution strategies. Those observations led to the more rigorous investigations in this paper.

2 Problem difficulty in dynamic environments

As Bäck (1998) has shown empirically, evolution strategies (ES) are useful for tackling dynamic problems. Especially the self-adaptation enables the algorithm to reproduce slight dynamic modifications of the problem or even, in the case of bigger changes, to re-optimize the problem. However, the present article examines which characteristics of dynamic problems might hinder the optimization with evolution strategies and, thus, identifies a certain kind of ES-hardness. In particular, this paper is concerned with drifting landscapes where a static landscape moves slightly over time. As an exemplary application De Jong (2000) discusses the control of chemical production processes where a gradual movement of the optimal control point is caused by aging equipment or varying quality of raw material. Within such a problem framework, the following two subsections introduce two hypotheses which reflect properties of dynamic ES-hard problems.

2.1 Discontinuity caused by dynamics

In any evolutionary algorithm the operators have a certain underlying conception of how the next individuals should be created using previous individuals. These conceptions can be made explicit by operator-defined fitness landscapes (Jones, 1995) in the case of discrete search spaces and by probability density functions (pdf) in the case of continuous search spaces, which is the case with evolution strategies. In order to get good results there should be a high correlation between the characteristics of the problem and this conception of the operator. One example for a good operator-problem interaction are the Gaussian pdf of the evolution strategy mutation applied to real-valued, smooth, stationary problems. A smooth, partially monotonous search space (e.g. sphere model) guarantees that even a small step in the right direction is a good step. This fits perfectly to the internal model of the zero-mean Gaussian pdf. The additional self-adaptation mechanisms of the evolution strategies enable a quick adaptation to quite different problem spaces.

However, by introducing dynamic, time-varying aspects into a problem this simple correlation between operator and problem is disturbed. This may be explained for the evolution strategy mutation using Figure 1. The upper row shows three different one-dimensional fitness landscapes where the circled cross marks the current position of a candidate solution. The middle row shows how the fitness landscape is shifted from one generation to the next. And the lower row shows the arising discontinuity from this shift as the difference between the fitness values we could expect if the problem was stationary and the fitness values we encounter in the next generation. The smoothness we could actually expect gets disarranged from generation t to generation $t+1$. In the landscape in the middle row, a step to the right is necessary to reach the optimum, the mutation cannot do wrong at time t if a small step is chosen instead of a bigger step. But this is not true at time $t + 1$ anymore since the small step reaches a worse fitness than at time t. Here a bigger step is desirable. Thus dynamics introduce a new difficulty in the optimization which is probably not met by a mutation preferring small steps as is summarized in the following hypothesis.

Hypothesis 1 *The Gaussian pdf does not match the character and the demands of tracking problems.*

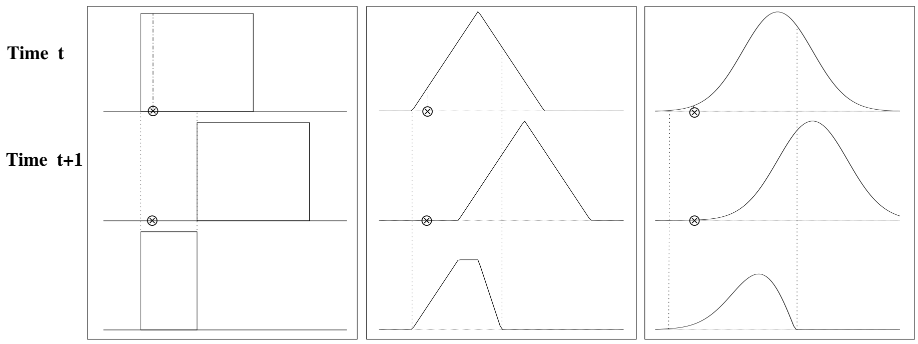

Fig. 1. This figure shows three different landscapes at time t (upper row) and $t + 1$ (middle row). By the change of the landscape there arises a region of discontinuity shown in the lower row where the fitness should be equal or better to a point situated to the left but where the fitness decreases caused by the dynamics.

2.2 Combination of tracking and optimization

Another prerequisite in most research articles on dynamic optimization is the fact that usually only the tracking of an optimum is considered, i.e. the underlying optimization problem is assumed to be easy. However, if this is not the case, the question arises whether both the optimization task and the tracking task fit to the operator's pdf to generate new individuals. In particular, any self-adaptation mechanism has to support both tasks. The tracking of the current region of interest is a necessary requirement for the optimization – however, as it is explained in detail in Section 5, successful steps in the optimization task may disturb the settings of the self-adaptive strategy parameters in such a way that no tracking is possible anymore. This leads to the second hypothesis.

Hypothesis 2 *The ES self-adaptation mechanisms might fail in problem settings which require optimization within a tracking task.*

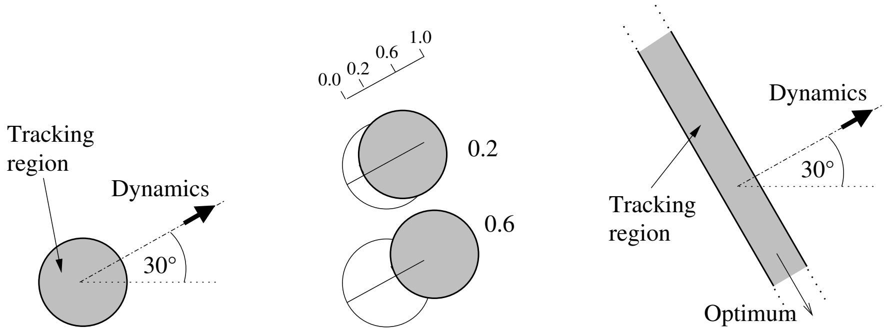

Fig. 2. On the left, the moving circle problem is shown, the meaning of the dynamics parameter is shown exemplary for the moving circle problem in the middle, and the moving corridor problem is shown on the right.

Such a combination of tracking and optimization may occur in a constrained optimization problem where the constraints are changing over time. In Section 4 a simple problem model for such a combination is given. A more realistic example can be found in (Weicker & Weicker, 1999).

After a short introduction to the considered test problems in the next section, the hypotheses are supported by experiments in Sections 4 and 5. Also in both sections methods for the resolution of the problem are discussed and tested partially.

3 Problem Framework

In order to examine the hypotheses two very simple problems are defined resembling the underlying necessary characteristics of the hypotheses. For the sake of simplicity, a two-dimensional search space $\Sigma = \mathbb{R} \times \mathbb{R}$ is considered in which a tracking region with "good" fitness is defined. All other points are assigned a "bad" fitness. In more realistic optimization tasks the tracking region may be defined by points with near-optimal fitness or by dynamic constraints. This underlines the relevance of those simplified problems for dynamic optimization in general.

To study the first hypothesis, the *moving circle problem* as a pure tracking task is introduced. Here the tracking region consists of a circle which moves within the search space (cf. left part of Figure 2). The fitness within the tracking region is determined by the distance to the center of the circle where the optimum is located. All other points are assigned a constant bad fitness value. This problem is very similar to the moving hill or the moving sphere function in various previous examinations (e.g. Cobb, 1990; Collard, Escazut, & Gaspar, 1996; Branke, 1999b; Liles & De Jong, 1999; Morrison & De Jong, 1999). In order to avoid pathological behavior, the direction of the dynamics was chosen to be in a 30° angle from one axis. This prevents dependence on one axis only as well as the other symmetric extreme, an angle of 45°. The strength of the dynamics may vary in the range $[0.0, 1.0]$ which corresponds to the fraction of new, non-overlapping parts of the tracking region (cf. middle part of Figure 2). After each generation the tracking region is moved according to the strength. Strength 0.0 means that the problem is stationary and strength 1.0 imply that the circles do not overlap from one generation to the next but are still adjacent to each other.

In the examination of the second hypothesis, the focus is on a combination of tracking and optimization which might be unfavorable for the evolutionary algorithm. Thus a scenario was chosen where the direction of the tracking tasks is arranged orthogonal to the direction of the optimization task, the *moving corridor problem* (cf. the right part of Figure 2). The fitness value within the corridor is determined by how far we have moved towards the optimum. The strength of the dynamics is determined similarly to the moving circle model. Also, the corridor is arranged in a 30° angle from one axis. Such a model seems to be pathological in a two-dimensional problem setting, however, those orthogonal effects may occur easily in multi-dimensional search spaces with moving constraints.

In the following two sections various evolution strategies are applied to the two problems. As standard step-size adaptations, uniform self-adaptation with one strategy variable as standard deviation for all search space dimensions and separate self-adaptation with a different strategy variable for each search space dimension are compared. Also, two extensions are introduced in the next two sections. For each algorithm, problem, and strength of dynamics, 100 independent experiments are executed using different random seeds. All investigated algorithms use a $(1, 40)$-strategy without recombination.

4 Experiments: Moving Circle

The moving circle problem is used to examine Hypothesis 1 which states that the Gaussian pdf does not match the character and demands of tracking problems. Figure 1 sketches the underlying thoughts that a small step means worsening if the tracking region moves. This is empirically examined using the average best fitness over all experiments, the percentage of invalid individuals created, and the percentage of experiments which got completely lost, i.e. they could not track the tracking region.

The experiments are shown in Figures 3, 4, and 5. In uniform step-size adaptation the individual is extended by one strategy parameter for the standard deviation of the mutation, in separate step-size adaptation a distinct strategy parameter is used for each search space dimension. For the sake of completeness, the figures contain the performance of a technique, the offset mutation, introduced in the remainder of this section and a correction strategy which is combined with offset and uniform mutation and introduced in the next section.

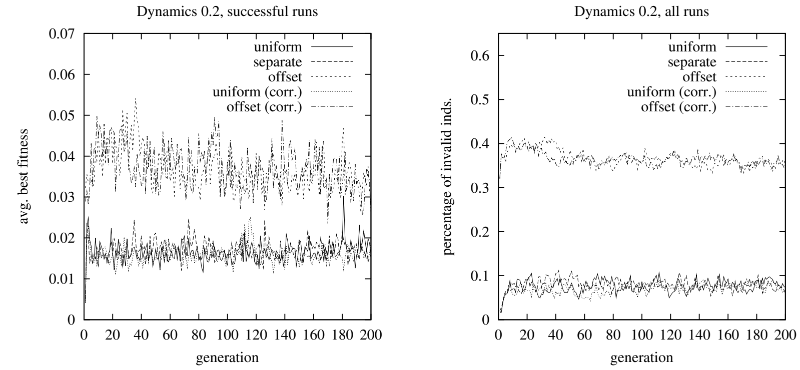

Fig. 3. Results for the moving circle problem with static strength of dynamics 0.2. The average best fitness of all successful experiments (left) shows better performance of all variants of standard step-size adaptation compared to both offset techniques. This behavior is also reflected in the percentage of invalid created individuals (right).

From the experiments we can conclude that self-adaptive evolution strategies with either uniform step-size for all dimensions or separate step-sizes exhibit a substantially increasing number of invalid individuals with increasing dynamics. As a consequence the tracking accuracy measured by the average best fitness decreases with increasing dynamics. Also at the end of 200 generations, we see in Table 1 that for both uniform and separate step-size adaptation the fraction of complete lost runs increases considerable with increasing dynamics.

Experiments with the 1/5–success rule have been executed but are not included in the the results for reasons of very inferior performance. Covariance matrix adaptation was not used in the context of this work since preliminary studies (Weicker & Weicker, 1999) have shown that those adaptation mechanisms are often too inertial.

One approach to generate more effective new individuals in the face of dynamics is to shift the mean of the generating distribution away from zero, resulting in a directed mutation. One approach to reach such a behavior is the skewing of the underlying pdf as it was done by Hildebrand, Reusch, and Fathi (1999). Another approach by Ghozeil and Fogel (1996) uses an offset direction vector which determines the new position where a non-zero Gaussian mutation takes place. The second method was modified slightly for this work. In general the strategy parameters are an offset vector represented as an n-dimensional unit vector in Cartesian coordinates $\boldsymbol{v} \in \mathbb{R}^n$ and the standard deviation as $\sigma \in \mathbb{R}_+$. In this examination $n = 2$. The mutation on an individual $\boldsymbol{x}$ works as follows.

$$\begin{aligned}
\sigma &= \sigma \exp\left(\frac{1.0}{\sqrt{1.0n}}\mathcal{N}(0,1)\right) \\
v_i &= v_i + \mathcal{N}(0,1) \\
x_i &= x_i + v_i + \mathcal{N}(0,\sigma)
\end{aligned}$$

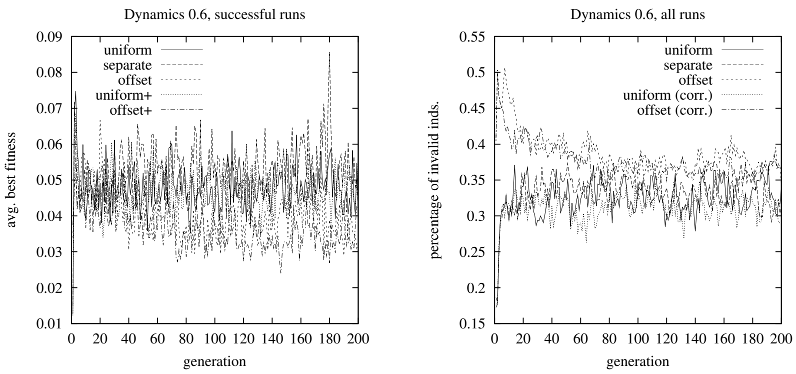

Fig. 4. Results for the moving circle problem with static strength of dynamics 0.6. Compared to the behavior with strength 0.2 in Figure 3 the performance of all techniques adjusts to each other. In case of the average best fitness there are even slight insignificant advantages for the offset techniques.

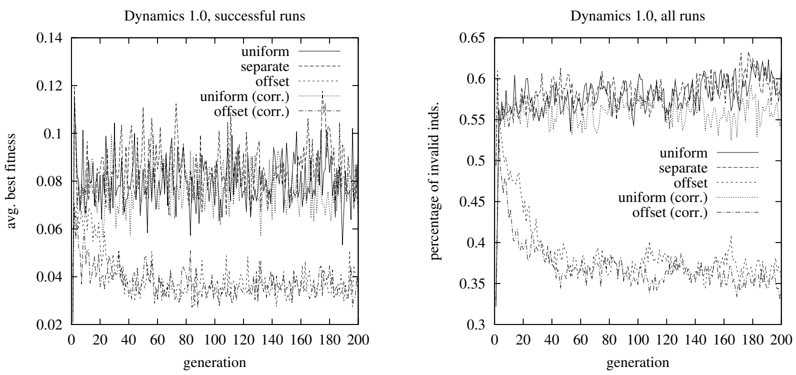

Fig. 5. Results for the moving circle problem with static strength of dynamics 1.0. The trend of Figures 3 and 4 continues. The offset technique preserves the same level of average fitness and percentage of invalid individuals where the standard step-size adaptation performs significantly worse.

	strength of dynamics		
	0.2	0.6	1.0
uniform	0.00	0.02	0.10
separate	0.00	0.02	0.09
offset	0.01	0.03	0.02

Table 1. Fraction of completely lost experiments at the end of 200 generations for the moving circle problem.

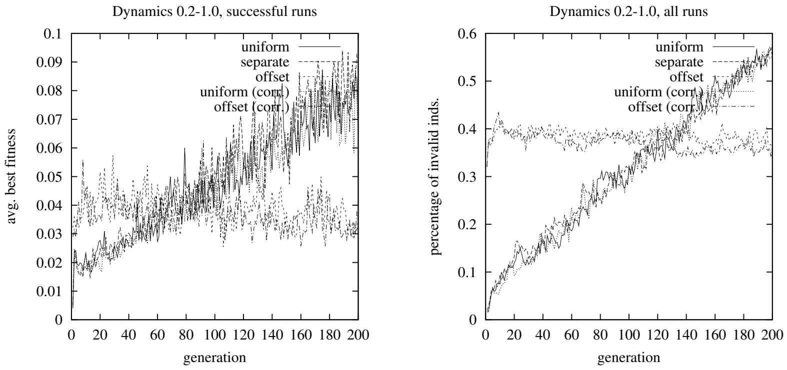

Fig. 6. The average best fitness (left) and the percentage of invalid created individuals (right) in experiments with the tracking circle model and strength of dynamics increasing linearly from 0.2 to 1.0 show an almost invariant behavior of the offset mutation and increasing problems of all other adaptation techniques.

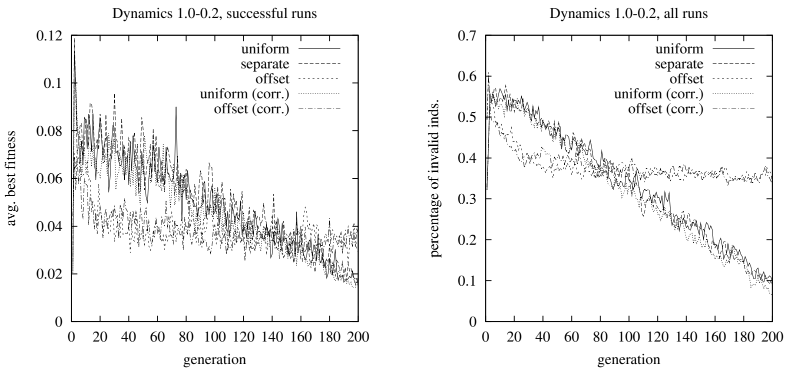

Fig. 7. Experiments with the tracking circle model and strength of dynamics decreasing linearly from 1.0 to 0.2 exhibit the same behavior as in the experiments in Figure 6.

Concerning the adaptation mechanism for the offset vector this is an ad hoc approach. No fine tuning of the mechanism was done so far.

The experiments in Figures 3, 4, and 5 and Table 1 show that this approach is almost completely unaffected by the strength of the dynamics. The tracking accuracy is constant with respect to the dynamics. However, the standard adaptation mechanisms yield a higher tracking quality with slow dynamics.

Additional experiments with increasing (resp. decreasing) dynamics between 0.2 and 1.0 reflect exactly the behavior with static strength of dynamics and are shown in Figures 6 and 7. This underlines the adaptability of the introduced directed mutation.

Thus the first hypothesis is supported by the fact that there is a high dependency of the quality of standard adaptation mechanisms on the strength of

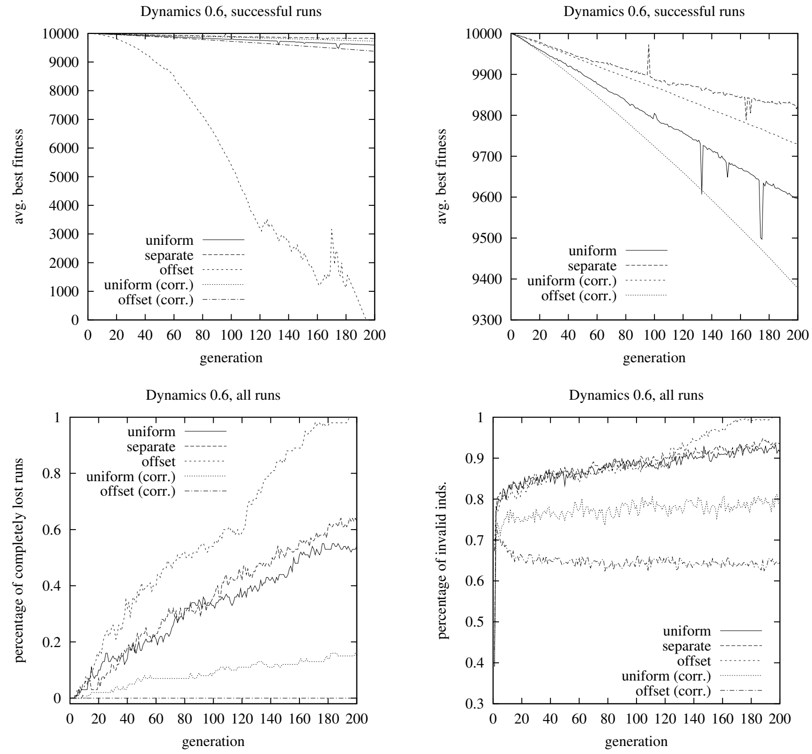

Fig. 8. Results for the moving corridor problem with strength of dynamics 0.6: the upper row shows the best fitness values of all successful experiments on average (with two differently scaled ordinates), the lower row shows the fraction of completely lost experiments (left) and the fraction of invalid generated individuals (right). Superior best fitness can be observed for offset mutation which comes along with high percentages of lost runs and invalid individuals. Only the offset mutation with strategy variable correction shows a stable performance concerning the invalid individuals and the lost runs.

dynamics and the existence of an adaptation mechanism which overcomes those limitations by breaking the principle of a zero-mean mutation favoring small steps.

5 Experiments: Moving Corridor

This section focuses on the second hypothesis which states that there are circumstances where optimization may be difficult in the presence of dynamics, i.e. if successful tracking is a necessary prerequisite for optimization. This hypothesis is examined using the introduced moving corridor problem.

The results for the standard adaptation techniques as well as the directed offset mutation are shown in Figure 8 exemplary for dynamics 0.6. Other strength

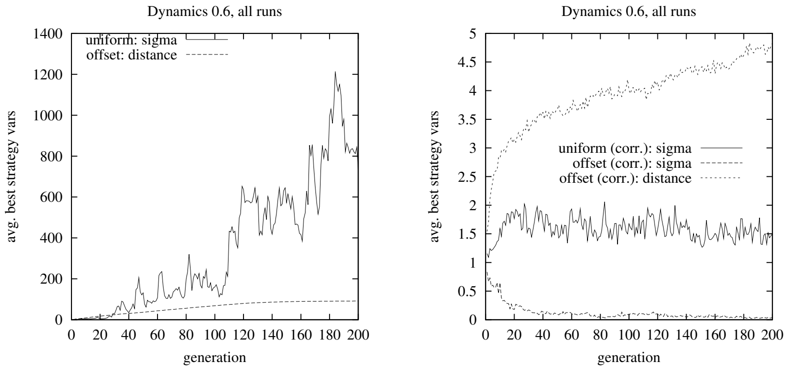

Fig. 9. Analysis of the adaptation of strategy variables for the moving corridor problem: the left picture shows the values for the uniform and the offset adaptation (the standard deviation of the offset adaptation is omitted since it stays close to 1) and the right picture shows the values for the adaptation techniques with correction of strategy variables.

of dynamics reveal a similar behavior. The offset mutation shows in its successful runs a superior behavior (as would separate adaptation do if optimization was aligned with a coordinate axis.) However this is accompanied by a huge increase in completely lost runs and invalid individuals. At the end of 200 generations almost each created individual is invalid. Uniform and separate step-size adaptation show the same tendency. The reason of this effect can be clarified by looking at the evolution of the strategy parameters shown in Figure 9. Since the adaptation mechanism always accepts those individuals (and, thus, indirectly the respective strategy values) that score the best fitness values, this results in accepting those settings which make huge jumps and hitting the tracking region by accident. However, there is only a very small probability that these strategy variables can reproduce the same effect in the next iterations. Thus the number of invalid individuals increases.

As a consequence, a technique is needed which takes care of the discrepancy between the greedy behavior concerning the optimization and the needs for successful tracking. We decided to correct the strategy variables of the accepted individual in such a way that tracking is still possible with a high success rate. The mutation of the strategy and object variables is not affected by this approach. But as soon as the surviving individual is selected (according to the best fitness) its strategy variables are modified in the following way. For all in the current generation created individuals with a higher fitness than the parent's fitness the distance $\sqrt{\sum_{i=1}^{n}(s_i^{child} - s_i^{parent})^2}$ of their strategy variables $\boldsymbol{s}^{child}$ and the parent's strategy variables $\boldsymbol{s}^{parent}$ are computed. The strategy variables of the individual with the minimal distance are used to replace the accepted individual's strategy variables. If there is none such individual no correction takes place.

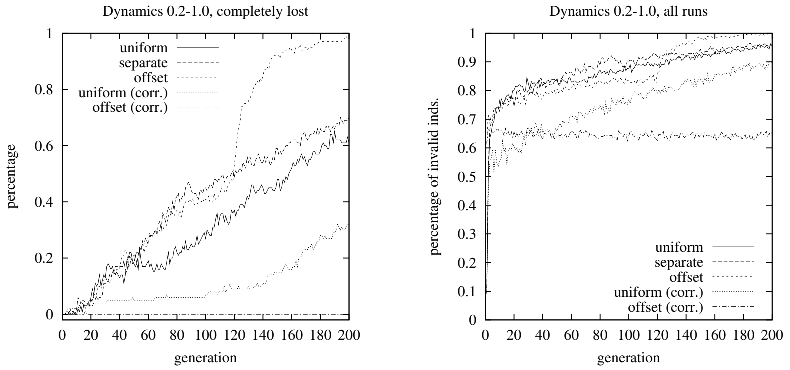

Fig. 10. Moving corridor problem with linearly increasing strength of dynamics from 0.2 to 1.0. The techniques without strategy variable correction show an almost chaotic behavior in generation 200. Uniform adaptation with correction performs better until generation 100, and, then, worsens considerably. Only offset mutation with correction shows a stable behavior and adapts without problems to the changing requirements.

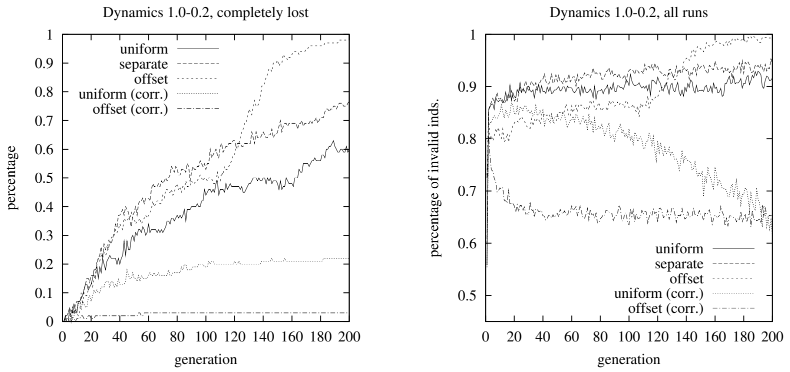

Fig. 11. Moving corridor problem with linearly decreasing strength of dynamics from 1.0 to 0.2. The results are similar to Figure 10. Only uniform adaptation with correction shows a better performance and offset mutation with correction exhibits a slightly higher number of completely lost runs.

The results for the moving corridor problem can be seen in Figures 8 and 9. The correction mechanism dampens the explosion of the strategy variables for both uniform and offset adaptation. As a consequence, the invalid created individuals can be kept on an almost constant level which leads to only few completely lost runs – in case of the offset adaptation even all runs are successful.

The results in Figures 10 and 11 with varying dynamics indicate that the correction strategy can handle non-stationary dynamics too. However diminishing dynamics are more problematic since there completely lost runs arise. In case of the uniform step-size adaptation with correction of strategy variables the increasing dynamics are more problematic since there seems to be a threshold where the percentage of invalid individuals does not suffice for the population size anymore.

The experiments with the moving circle problem in Figures 3, 4, and 5 show in addition that the mechanism does not hinder in tracking-only tasks.

All in all, the experiments show that there exist circumstances where standard adaptation techniques fail. The correction of strategy variables seems to be one possible remedy.

6 Conclusions

This work has shown empirically that two basic components of evolution strategies, the Gaussian pdf and the adaptation mechanisms, can have problems in dynamic environments that do not occur in stationary environments. By analyzing two simple dynamic problems indications to the reasons for those problems are found and guided the design of two new, proposed techniques. Directed mutations and the introduced correction of strategy variables seem to be at least feasible approaches to handle those problems. However, both extensions need further investigations on more general problems. Especially, the directed mutation must be evaluated on problems where the direction of the moving tracking region changes. Also it is clearly not a means to dynamic problems with a more chaotic behavior.

This clarifies that the goal of this work was not to replace the principles of standard ES mutation by new principles for dynamic problems. It rather illustrated a first attempt to tailor the mutation operator to a special class of dynamic problems. In the future, more theoretical considerations are necessary on the requirements posed by other classes of dynamic problems. This should lead in the long run to operators which are able to adapt best to different kinds of dynamics.

References

Bäck, T. (1998). On the behavior of evolutionary algorithms in dynamic environments. In *IEEE Int. Conf. on Evolutionary Computation* (pp. 446–451). IEEE Press.

Branke, J. (1999a). Evolutionary approaches to dynamic optimization problems: A survey. In J. Branke & T. Bäck (Eds.), *Evolutionary algorithms for dynamic optimization problems* (pp. 134–137). (part of GECCO Workshops, A. Wu (ed.))

Branke, J. (1999b). Memory enhanced evolutionary algorithms for changing optimization problems. In *1999 Congress on Evolutionary Computation* (pp. 1875–1882). Piscataway, NJ: IEEE Service Center.

Cobb, H. G. (1990). *An investigation into the use of hypermutation as an adaptive operator in genetic algorithms having continuous, time-dependent nonstationary environments* (Tech. Rep. No. 6760 (NLR Memorandum)).

Washington, D.C.: Navy Center for Applied Research in Artificial Intelligence.

Collard, P., Escazut, C., & Gaspar, A. (1996). An evolutionary approach for time dependant optimization. In *Int. Conf. on Tools for Artificial Intelligence 96* (pp. 2–9). IEEE Computer Society Press.

De Jong, K. (2000). Evolving in a changing world. In Z. Ras & A. Skowron (Eds.), *Foundation of intelligent systems* (pp. 513–519). Berlin: Springer.

Ghozeil, A., & Fogel, D. B. (1996). A preliminary investigation into directed mutations in evolutionary algorithms. In H. Voigt, W. Ebeling, & I. Rechenberg (Eds.), *Parallel Problem Solving from Nature – PPSN IV (Berlin, 1996) (Lecture Notes in Computer Science 1141)* (pp. 329–335). Berlin: Springer.

Goldberg, D. E., & Smith, R. E. (1987). Nonstationary function optimization using genetic algorithms with dominance and diploidy. In J. J. Grefenstette (Ed.), *Proc. of the Second Int. Conf. on Genetic Algorithms* (pp. 59–68). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hildebrand, L., Reusch, B., & Fathi, M. (1999). Directed mutation – A new self-adaptation for evolutionary strategies. In *1999 Congress on Evolutionary Computation* (pp. 1550–1557). Piscataway, NJ: IEEE Service Center.

Jones, T. (1995). *Evolutionary algorithms, fitness landscapes and search.* Unpublished doctoral dissertation, The University of New Mexico, Albuquerque, NM.

Liles, W., & De Jong, K. A. (1999). The usefulness of tag bits in changing environments. In *1999 Congress on Evolutionary Computation* (pp. 2054–2060). Piscataway, NJ: IEEE Service Center.

Mori, N., Kita, H., & Nishikawa, Y. (1996). Adaptation to a changing environment by means of the thermodynamical genetic algorithm. In H. Voigt, W. Ebeling, & I. Rechenberg (Eds.), *Parallel Problem Solving from Nature – PPSN IV (Berlin, 1996) (Lecture Notes in Computer Science 1141)* (pp. 513–522). Berlin: Springer.

Morrison, R. W., & De Jong, K. A. (1999). A test problem generator for non-stationary environments. In *1999 Congress on Evolutionary Computation* (pp. 2047–2053). Piscataway, NJ: IEEE Service Center.

Rechenberg, I. (1994). *Evolutionsstrategie '94.* Stuttgart: frommann-holzbog. (German)

Saleem, S., & Reynolds, R. (2000). Cultural algorithms in dynamic environments. In *Proc. of the 2000 Congress on Evolutionary Computation* (pp. 1513–1520). Piscataway, NJ: IEEE Service Center.

Schwefel, H.-P. (1995). *Evolution and optimum seeking.* New York: Wiley.

Vavak, F., Fogarty, T. C., & Jukes, K. (1996). A genetic algorithm with variable range of local search for tracking changing environments. In H.-M. Voigt, W. Ebeling, I. Rechenberg, & H.-P. Schwefel (Eds.), *Parallel problem solving from nature – PPSN IV* (pp. 376–385). Berlin: Springer.

Weicker, K., & Weicker, N. (1999). On evolution strategy optimization in dynamic environments. In *1999 Congress on Evolutionary Computation* (pp. 2039–2046). Piscataway, NJ: IEEE Service Center.

Weicker, K., & Weicker, N. (2000). Dynamic rotation and partial visibility. In *Proc. of the 2000 Congress on Evolutionary Computation* (pp. 1125–1131). Piscataway, NJ: IEEE Service Center.