

Disburdening the Species Conservation Evolutionary Algorithm of Arguing with Radii

Catalin Stoean
University of Craiova
A. I. Cuza, No. 13
200585, Craiova, Romania
catalin.stoean@inf.ucv.ro

Mike Preuss
University of Dortmund
Otto-Hahn 14
44221, Dortmund, Germany
mike.preuss@cs.uni-dortmund.de

Ruxandra Stoean
University of Craiova
A. I. Cuza, No. 13
200585, Craiova, Romania
ruxandra.stoean@inf.ucv.ro

D. Dumitrescu
Babes-Bolyai University
M. Kogalniceanu, No. 1B
400084, Cluj, Romania
ddumitr@cs.ubbcluj.ro

ABSTRACT

The present paper investigates the hybridization of two well-known multimodal optimization methods, i.e. species conservation and multinational algorithms. The topological species conservation algorithm embraces the vision of the existence of subpopulations around seeds (the best local individuals) and the preservation of these dominating individuals from one generation to another, but detects multimodality by means of the hill-valley mechanism employed by multinational algorithms. The aim is to inherit the strengths of both parent techniques and at the same time overcome their flaws. The species conservation algorithm efficiently keeps track of several good search space regions at once, but is difficult to parametrize without prior problem knowledge. Conversely, the multinational algorithms use many function evaluations to establish subpopulations, but do not depend on provided radius parameter values. Experiments with all three algorithms are made on a wide range of test problems in order to investigate their advantages and shortcomings.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*heuristic methods*

General Terms

Algorithms, Experimentation

Keywords

Detect-multimodal mechanism, hybridization, multimodal evolutionary algorithms, species conservation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

1. INTRODUCTION

The species conservation technique [4] is a recent radii-based multimodal evolutionary algorithm that preserves the best individual in each subpopulation (or species) from one generation to another. In order to differentiate among the species, a radius parameter is introduced that is employed to calculate the similarity of every two individuals. If their distance is smaller than the given radius threshold, the two individuals are assigned to the same subpopulation. Although very efficient in conducted experiments, the method still falls under the unanimously accepted fact that there are difficulties stemming from the necessary distance measuring, especially in high dimensional search spaces. Besides adding a parameter to the algorithm that is hard to set accurately, it is well-known that contrast in distance values drops rapidly for more than 10 dimensions [2].

Shir and Bäck [7] tackle this problem by adapting the radii within their niching evolutionary algorithm. Conversely, we chose the alternative and aim to overcome this disadvantage by removing the radii and introducing an existing promising alternative method of subpopulation differentiation instead. The approach is topological instead of radii-based and has been proposed within the complex model of multinational evolutionary algorithms in [8], [9]. It detects whether two individuals track the same optimum by considering a set of additional candidate solutions between them and testing if one of these is weaker than both. If this is the case, they are assumed to track different optima and shall be distributed to different subpopulations.

Herein, we investigate a hybridization: Topological species conservation (TSC). It is hoped to combine the strengths of both techniques, while avoiding their weaknesses. On the one hand, the species conservation algorithm is simpler and more efficient than the multinational paradigm; on the other hand, the latter comes with a more natural decomposition of the population into species. By combining species conservation with an alternative mechanism for subpopulation determination we strive for simultaneously achieving two goals: Preservation of the best local individuals during runtime and natural differentiation between the multiple final solutions of the multimodal problem. In order to objectively validate

the new hybrid method, all considered test cases are also solved by the parent algorithms. Note that we deliberately abstain from using any stepsize adaptation mechanism as this would further complicate experimental analysis.

The paper is structured as follows. Sections two and three outline the concepts of species conservation and multinational algorithms. Section four describes TSC and section five presents the chosen test problems and experiments.

2. SPECIES CONSERVATION

The species conservation algorithm concentrates on two aspects [4]: The determination of species based on similarity criteria, where each subpopulation is dominated by a locally best individual, and the preservation of these individuals to the following generation of the evolutionary cycle.

2.1 Subpopulation Differentiation

A species is defined as a subset of individuals where the distance between every two members is less than the diameter given by a user-defined species radius. Each subpopulation is built around the best local individual, called species seed. A species is centred on its seed if for its every member the distance to it is less than the species radius. At the beginning of every generation, the mechanism that differentiates the species collects seeds by taking into account each individual of the population, in decreasing order of fitness, and testing whether it belongs to any species centred on the already found seeds (i.e. its distance to a certain known seed is less than the species radius) or is a new seed. The way of finding the seeds of each subpopulation is described by Algorithm 1. Population size is denoted by n .

Algorithm 1 Seeds selection within the canonical species conservation algorithm

```

Seeds =  $\Phi$ ;
Sort population  $P$  decreasingly according to fitness;
for  $i = 1$  to  $n$  do
    found = FALSE;
    for every  $s$  in Seeds do
        if  $d(P_i, s) \leq \text{radius}$  then
            found = TRUE;
        end if
    end for
    if not found then
        Seeds = Seeds  $\cup$   $P_i$ ;
    end if
end for

```

2.2 Subpopulation Conservation

The method that is responsible for conserving the dominating individuals acts on the population after applying variation operators. For every seed s in turn, the worst fit individual w that belongs to the species centered on s (i.e. its distance to s is less than the species radius) is considered. If this individual exists and is less fit than s , then the former is replaced by the latter. Otherwise, if the species of s does not contain any individual at all (as they may have disappeared because of selection or variation operators), then the worst individual from the entire population w is replaced by s . The individual s that enters the population is prevented from further substitution during the current generation. Apart from conserving the fittest individuals from the entire population, this mechanism allows the preservation

of even less fit individuals, different enough from the global best ones, that could be at that moment positioned at the basis of an empty peak and thus become useful in future iterations. Subpopulation conservation is thus carried out through the conservation of the seeds of each species; the process is illustrated by Algorithm 2.

Algorithm 2 Seeds conservation within the canonical species conservation algorithm

```

Mark all individuals in  $P$  as unprocessed;
for every  $s$  in Seeds do
    Take worst unprocessed  $w$ , such that  $d(w, s) \leq \text{radius}$ ;
    if  $w$  exists then
        if  $f(w) < f(s)$  then
             $w = s$ ;
        end if
    else
        Take worst unprocessed  $w$  in  $P$ ;
         $w = s$ ;
    end if
    Mark  $w$  as processed;
end for

```

An important remark that concerns the evolutionary algorithm as a whole is that selection and recombination of individuals are both performed globally, irrespective of the different species they belong to. This may be counterproductive and resembles one of the critical parts of the algorithm. We can hope to improve its overall behavior by inserting an alternative approach.

At the end of the algorithm, the optima are selected from the seeds that were lastly chosen. This *Seeds* set will not contain only the desired optima, but also low fitness individuals that were stored because they were sufficiently different from all the other individuals. Therefore, another parameter, *solution acceptance threshold* r_f ($0 < r_f \leq 1$), is introduced so that all solutions that have high enough fitness to be selected: $x \in \text{Seeds}$ is selected if the inequality $f(x) \geq (f_{max} - f_{min}) \cdot r_f$ is satisfied, where $f(x)$ denotes the fitness of x , while f_{max} and f_{min} represent the maximum and minimum fitness in the final population.

3. MULTINATIONAL ALGORITHMS

The technique [8], [9] has a very interesting working metaphor: The whole population represents a world and each subpopulation is called a nation; each nation has a government which contains its fittest individuals, referred to as politicians. They define the policy of the nation, which is in fact the centroid of these individuals. Initially, the world consists a single nation. Then, the governments and, as a consequence, the policies are useful for the differentiation and the dynamics of subpopulations.

3.1 The Multinational Engine

The multinational idea is achieved through two mechanisms: Migration and merging. Within every generation, each individual is compared to the policy of its nation. If the two follow different optima, then the former migrates to a nation whose policy is tracking the same peak as itself. If there exist no such nation, then the individual forms a new nation, corresponding to a potentially new peak. Conversely, also at every generation, the policies of each nation are checked two by two to ensure that they do not track the

same optimum; if this is the case, then the two subpopulations are merged.

Selection can be performed on two levels: Either within each subpopulation or in a global but weighted fashion, i.e. the fitness of an individual is divided by the number of members in its corresponding nation, in a fashion that resembles sharing [3]. Finally, crossover is performed only between individuals of the same nation, as the combination of the genetic material of points that track different optima may lead to the appearance of offspring that are less fit than the parents. We assumed that mutation is also restricted to each nation, i.e. the offspring is accepted only if it remains within the premises.

3.2 The Detect-Multimodal Method

The verification of the relationship of two points in the search space, i.e. of the assumption that they track the same optimum, is performed through an approach originally called the *hill-valley* mechanism, but which is renamed here to *detect-multimodal* for reasons of simplicity. The function takes two individuals (points) as arguments and returns whether or not there is a valley between them in the fitness landscape, i.e. they track different optima. In order to reach that decision, a set of interior points between the two, based on user-defined gradations in the [0,1] interval, is generated. If the fitness of all interior points is higher than the minimal fitness of the two tested individuals, then it is concluded that they track the same optimum. Contrarily, if there exist such a point whose fitness is smaller than the minimal fitness of the two, then it is assessed that they follow different peaks. To conclude, *detect-multimodal* returns true if the two points follow different optima and false if they follow the same peak. The mechanism is described in Algorithm 3.

Algorithm 3 The detect-multimodal mechanism between two individuals x and y

```

i = 1;
found = FALSE;
while i < number of gradations and not found do
  for j = 1 to number of dimensions do
    interiorj = xj + (yj - xj) · gradationj;
  end for
  if f(interior) < min(f(x), f(y)) then
    found = TRUE;
  end if
end while
return found;

```

An important advantage of this manner of detecting multimodality is that a certain optimum is tracked by only one subpopulation, whereas the radii-based mechanism of species conservation allows the existence of several subpopulations that follow the same peak.

4. THE TOPOLOGICAL SPECIES CONSERVATION ALGORITHM

The novel proposed algorithm inherits the notion of a seed (dominating individual) of a subpopulation and the idea of seeds conservation. Simultaneously, the new approach determines subpopulations through the *detect-multimodal* procedure and uses a weighted selection with the employment of global recombination, on the one hand, and local selection accompanied by local recombination, on the other hand.

4.1 Motivation

Hybridization of the two parent algorithms may have several advantages. Firstly, the efficiency of the species conservation method lies in the conservation of several locally fittest individuals (seeds): Thus, subpopulations are not lost completely, even if selection may skip all individuals within one population or they may disappear because of recombination and mutation. Conservation of the seeds of the found subpopulations prevents them from getting extinct. Secondly, replacing the radius-dependent method to differentiate subpopulations in favor of one that employs fitness discrepancies as in multinational algorithms, may have two advantages. On the one hand, one gets rid of a crucial parameter (for which it is very difficult to find a proper value, especially in higher dimensional problems) and, on the other hand, less fit individuals that are actually not promising, but merely different enough from the others, are rapidly detected. This is obvious especially for large plateaus contained in the fitness landscape: While the species conservation method would form a great number of subpopulations, multinational algorithms detect only one peak to follow.

Thirdly, in the TSC variant, when weighted selection is used and recombination is performed globally, seeds are conserved and subpopulations consequently cannot disappear. Although recombination between individuals from different subpopulations usually produces weaker individuals that do not belong to any of the subpopulations of their parents, more exploration of the search space is gained by its means. Finally, we managed to avoid the expensive behavior of the original multinational algorithm that, due to its subpopulations dynamics, that is achieved through migration and merging, uses a high number of fitness evaluations in a very small number of generations. By incorporating the preservation of multimodality through seed conservation and efficiently keeping track of each individual's subpopulation during evolution, we overcame this great disadvantage.

4.2 The Mechanics

An important improvement over the multinational algorithms stems from the reduced number of calls of the *detect-multimodal* procedure. This is achieved by a subpopulation dynamic similar to the one in species conservation and is amplified by a good tracking of each individual during the evolutionary cycle. Consequently, the seed selection procedure (Algorithm 4) fulfills a two-fold purpose: At the same time, seeds are detected and corresponding subpopulations built, and all individuals are distributed over these established subpopulations.

In the first generation of the TSC algorithm, as we do not have any a priori knowledge about the population, all individuals are considered to be in the same subpopulation. All individuals that belong to the same species have the same value for the ID variable, so, before evolution starts, they are all in the subpopulation labelled with $ID = -1$. Moreover, when the seed selection procedure is called for the first time, then the *detect-multimodal* mechanism has to be applied as a tool for distinguishing individuals from different subpopulations. In the next generations, they are differentiated only through their ID s.

The population is sorted decreasingly, according to the fitness evaluations, so that species are formed around the current fittest seeds. In each call of the procedure, the ID s are reordered according to the the fitter individuals in the

population. We denote by n the size of the population, P_i represents the i -th individual, while P_{iID} stands for its ID .

Algorithm 4 Seeds selection procedure within TSC

```

if first generation then
  for  $i = 1$  to  $n$  do
     $P_{iID} = -1$ ;
  end for
end if
 $Seeds = \Phi$ ;
 $count = 1$ ;
Sort population  $P$  decreasingly according to the fitness;
for  $i = 1$  to  $n$  do
   $found = \text{FALSE}$ ;
  for every  $s$  in  $Seeds$  do
    if first generation then
      if not(detect-multimodal( $P_i, s$ )) then
         $found = \text{TRUE}$ ;
         $P_{iID} = s_{ID}$ ;
      end if
    else
      if  $P_{iID} = s_{ID}$  then
         $found = \text{TRUE}$ ;
      end if
    end if
  end for
  if not  $found$  then
     $P_{iID} = count$ ;
     $Seeds = Seeds \cup P_i$ ;
     $count = count + 1$ ;
  end if
end for

```

After seeds selection and labelling of each individual with the corresponding ID , the usual steps of an evolutionary algorithm follow. For both selection and recombination, where there is the need for recognizing subpopulations, the ID s of the involved individuals are recalled. However, when an offspring is obtained after recombination or mutation, its ID is set to -1. The only exception was considered in case of recombination when both parents (or all parents in case of a multi-parent recombination is considered) belong to the same subpopulation, in which case the offspring inherits the ID from them.

After the selection and variation operators are applied, the conservation of the seeds takes place (Algorithm 5), in a similar manner as in the case of the species conservation algorithm. In comparison to the procedure in the original algorithm, besides the fact that we do not make use of any distance, here we have to treat the cases of the *free* individuals obtained after the application of the variation operators, i.e. individuals with $ID = -1$. f denotes the fitness function.

As mentioned before, we propose two versions for the novel technique, judging from the way selection and recombination are performed, i.e. either globally or locally. However, the general lines followed by both versions of the hybridized technique are the same and depicted in Algorithm 6.

In the case weighted selection is applied, the fitness of each individual is divided by the number of members in its subpopulation. This also encourages the exploration of the search space, as less fit individuals that are situated alone in weaker regions still have a good chance of being selected. In this case, the recombination takes place globally.

At the end of the algorithm, the *Seeds* set will contain only the found optima as there is not the case that low fitness individuals, that follow the same optimum as others from the *Seeds* set, are kept simply because they were sufficiently

Algorithm 5 Seeds conservation procedure within TSC

```

Mark all individuals in  $P$  as unprocessed;
for every  $s$  in  $Seeds$  do
  for  $i = 1$  to  $n$  do
    if  $P_{iID} = -1$  then
      if not(detect-multimodal( $P_i, s$ )) then
         $P_{iID} = s_{ID}$ ;
      end if
    end if
  end for
  Take worst unprocessed  $w$  from  $P$ , such that  $s_{ID} = w_{ID}$ ;
  if  $w$  exists then
    if  $f(w) < f(s)$  then
       $w = s$ ;
    end if
  else
    Take worst unprocessed  $w$  in  $P$ ;
     $w = s$ ;
  end if
  Mark  $w$  as processed;
end for

```

different from the other individuals. To conclude, this is another advantage over the species conservation algorithm as herein there is no need of another additional parameter i.e. *solution acceptance threshold*.

Algorithm 6 TSC algorithm

```

 $t = 0$ ;
Initialize  $P(t)$ ;
Evaluate  $P(t)$ ;
while not termination condition do
   $t = t + 1$ 
  Seeds selection;
  Selection of  $P(t)$  from  $P(t - 1)$ ;
  Recombination on  $P(t)$ ;
  Mutation on  $P(t)$ ;
  Evaluate  $P(t)$ ;
  Seeds conservation;
end while
Identify species seeds in  $Seeds$ 

```

An interesting event happens from time to time after the seeds conservation procedure, as some individuals created during recombination and/or mutation remain without assigned population. This happens if *detect-multimodal* applied between them and any of the existing seeds returns true every time. After the seeds conservation procedure, all these individuals remain with $ID = -1$, and consequently they enter the new generation as a subpopulation of its own. From the undertaken experiments, we noticed that the number of these individuals does not exceed 10 (with a population size of 200) and they only appear during the initial 10 up to 50 generations, depending on the considered problem. Nevertheless, these individuals are very important in the evolution as they provide and maintain diversity in the population. When they enter in the seed selection procedure, they suffer a change in this ID and they form a subpopulation like any other.

5. EXPERIMENTS

The test problems employed in the following experiments are F1 (Waves) from [8] and F2 (Six-hump Camel Back) from [4], [8], while we added F3 (De Jong), F4 (Rastrigin) and F5 (SKEWBI) to fill all cells of a table made of the different modality conditions, one, few, and many optima,

Table 1: Utilized benchmark functions

Function	Optima
$F1(x, y) = (0.3x)^3 - (y^2 - 4.5y^2)xy - 4.7\cos(3x - y^2(2 + x))\sin(2.5\pi x)$ $-0.9 \leq x \leq 1.2, -1.2 \leq y \leq 1.2$	10
$F2(x, y) = -((4 - 2.1x^2 + \frac{x^4}{3})x^2 + xy + (-4 + 4y^2)y^2)$ $-1.9 \leq x \leq 1.0, -1.1 \leq y \leq 1.1$	6
$F3(\vec{x}) = \sum_{i=1}^n (-x_i^2)$ $-5.12 \leq x_i \leq 5.12$	1
$F4(\vec{x}) = -20 - \sum_{i=1}^n (-x_i^2 + 10\cos(2\pi x_i))$ $-5.12 \leq x_i \leq 5.12$	1/many
$F5(\vec{x}) = -\sum_{i < n, i \text{ odd}} (x_i^4 + x_{i+1}^4 - x_i^2 + x_{i+1}^2 - 10x_i x_{i+1} + 0.25x_i + 20)$ $-5 \leq x_i \leq 5$	$2^{\frac{n}{2}}$

and of the different search space sizes, low and high dimensional. Recent simulation model based investigations [5] lead to the conjecture that complex multimodal optimization algorithms may perform better than simple multistart methods only if the number of optima is relatively low. F3 is tested to show that even in case of only one optimum, the considered methods still perform sufficiently well. The SKEWBI (skewed bimodal) problem is an extension of the biquadratic polynomial defined as f_1 in [6] into a higher (even sized) dimensional search space. In each pair of dimensions, an optimization algorithm has to locate the better one of two evenly sized basins of attraction. The problem is thus partly separable with a slowly growing number of optima of which one is the global one and $2^{\frac{n}{2}} - 1$ are local ones. The problems with the number of known peaks are depicted in Table 1.

Experiment 1: Direct performance comparison.

Pre-experimental planning: In our initial version of the algorithm, in each generation, when seeds were selected, we applied *detect-multimodal* procedure in order to establish the subpopulations. From our early experiments, we observed that this was too expensive with regard to the number of fitness evaluations, although the results were very accurate if we had a termination condition that measured only the number of generations. This is what convinced us to determine the subpopulations in the initial generation and then track each individual’s species as much as possible by using only the IDs. This substantially reduced the number of fitness evaluations.

Task: The first task of the experiment was twofold. First of all, we wanted to measure the ability of all three techniques to find all the peaks, in low and high dimensions; for this reason, we chose $F1$, $F2$ and $F5$ (for higher dimensions). Second of all, we desired to measure the global optimization capabilities with $F3$, $F4$ and $F5$ in low and high dimensions. The second task envisaged the setting of an objective termination condition, i.e. a fixed number of fitness evaluations, and the performing of manual parameter tuning not only with the aim of giving an unbiased comparison between the final results, but also of showing the simplicity (or conversely, the intricacy) of choosing the parameters.

Setup: We established a fixed number of fitness evaluations of 10^5 for all considered techniques and kept the po-

pulation size up to 200 for all conducted tests within this experiment. For the current experiment, all parameter tunings were manually made; for each parameter setting, we performed 30 runs and outlined the average results. When comparing these parameter sets to the ones achieved by a semi-automated tuning method (SPO, see [1]), it turned out that the manually achieved values generally range near the median of the configurations tested by the tuning method, with a considerable distance to the top performance. That is, they are neither particularly bad nor good. This indicates that finding the best parameter settings for all three algorithms is at least not straightforward. The tuning results are omitted here due to lack of space.

For the first class of functions, where we verify the ability of the techniques to locate all the peaks, we use two measures: One refers the percentage of runs out of the maximum number when all peaks are detected (i.e. *all peaks found*) and the other one gives an average of the number of peaks found in all 30 runs (i.e. *average number of peaks*). A value of 0% for the former is not an indication of the fact that the algorithm completely failed, but merely that it was not able to detect all optima in any run as several found peaks out of the known number does not count as successful. We consider that a peak is found when there is at least one individual that has the fitness closer to it than a certain ϵ . We chose $\epsilon = 0.1$ for all considered test problems. Regarding the second class of functions, where the aim is to locate the global optimum, besides the above measurements, we also computed the average of the best fitness found in each of the 30 runs (i.e. *average best fitness*).

For all considered approaches, we employed tournament selection of size 2, intermediary recombination (two-parent (for global), multi-parent (for local) and one offspring) and mutation with normal perturbation. The offspring replaces the weaker parent only if it is fitter. The values for the parameters used for each version of the three techniques are depicted in Table 2. Vector size refers to the number of interior points for the *detect-multimodal* procedure.

Results/Visualization: The obtained results of each approach are outlined in Table 3. For $F1$ and $F2$ results seem to be slightly better for the two variants of the TSC algorithm. For $F5$ with 4 variables the best results were obtained by the canonical species conservation algorithm followed by the multinational algorithm using local selection, with some competitive values, closely tracked by TSC with local selection. Nevertheless, for 10 and 20 dimensions, the best results are achieved by TSC with local selection.

All used algorithms proved that they still have the capability to properly solve unimodal functions like $F3$, despite the fact that we used up to 10 variables for the function. Regarding the *average best fitness* values, one can observe that the TSC and canonical species conservation algorithm proved to be slightly more precise than the multinational.

As concerns $F4$, results are balanced between both variants of TSC and the canonical species conservation algorithm. However, for 10 dimensions, the average best fitness values for both TSCs overpass the value for the same parameter obtained for species conservation. As the question of tracking the global optimum for the $F5$ function arises, obtained values for the average best fitness in Table 3 show that in higher dimensions (10 or 20) TSC goes side by side with the species conservation algorithm and performs much better than the multinational algorithm.

Table 2: Parameter values for the TSC, the original species conservation and the multinational algorithms

	range	F1	F2	F3			F4			F5			
	problem dimension	2d	2d	2d	5d	10d	2d	5d	10d	2d	4d	10d	20d
TSC													
<i>Global selection</i>													
Population size	10-200	200	200	50	50	50	150	150	190	200	200	200	200
Crossover probability	0.0-1.0	0.50	0.6	0.40	0.40	0.40	0.70	0.70	0.90	0.90	0.90	0.90	0.90
Mutation probability	0.0-1.0	0.80	0.9	0.60	0.60	0.60	0.80	0.30	0.20	0.90	0.90	0.20	0.20
Mutation strength	0.01-20.0	0.60	0.8	0.40	0.40	0.40	0.70	1.6	0.60	0.50	0.50	0.20	1.20
Vector size	1-10	3	8	4	4	4	1	1	2	4	4	1	1
<i>Local selection</i>													
Population size	10-200	200	200	50	50	50	150	140	200	200	200	200	200
Crossover probability	0.0-1.0	0.90	0.60	0.40	0.40	0.40	0.50	0.95	0.90	0.90	0.90	0.90	0.80
Mutation probability	0.0-1.0	0.90	0.90	0.60	0.60	0.60	0.60	0.3	0.20	0.90	0.30	0.30	0.20
Mutation strength	0.01-20.0	0.50	0.80	0.40	0.40	0.40	0.70	0.7	0.70	0.50	2	2	1
Vector size	1-10	4	8	4	4	4	1	1	2	4	2	2	2
Species conservation													
Population size	10-200	200	200	100	100	100	200	200	200	200	200	200	200
Crossover probability	0.0-1.0	0.40	0.70	0.40	0.40	0.40	0.40	0.70	0.70	0.70	0.70	0.70	0.70
Mutation probability	0.0-1.0	0.60	0.20	0.60	0.60	0.60	0.80	0.50	0.40	0.40	0.40	0.40	0.20
Mutation strength	0.01-20.0	0.10	0.20	0.40	0.40	0.40	0.1	0.50	1.2	0.50	0.50	0.70	5
Radius	0.01-25.0	0.10	0.20	1	4	6	0.1	3	8	0.50	1.5	3	12
Multinational													
<i>Global selection</i>													
Population size	10-200	200	200	100	100	100	200	100	100	200	200	200	200
Crossover probability	0.0-1.0	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
Mutation probability	0.0-1.0	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80
Mutation strength	0.01-20.0	0.60	0.20	0.10	0.10	0.30	4	20	30	5	1	0.20	0.50
Vector size	1-10	3	10	3	3	3	1	1	1	3	3	3	3
Government size	1-10	2	5	2	2	2	10	8	8	2	2	2	2
<i>Local selection</i>													
Population size	10-200	200	200	100	100	100	100	100	100	200	200	200	200
Crossover probability	0.0-1.0	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
Mutation probability	0.0-1.0	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80	0.80
Mutation strength	0.01-20.0	0.20	0.20	0.10	0.10	0.10	2	20	30	1	1	0.20	0.50
Vector size	1-10	3	10	3	3	3	1	1	1	3	3	3	3
Government size	1-10	2	5	2	2	2	10	10	8	2	2	2	2

Observations: Despite the fact that $F1$ and $F2$ also appear in [8], we did not use the provided results as we wanted to make a fair comparison. Therefore, the fixed number of fitness evaluations brings changes in the final results as it makes both versions of the multinational algorithm run for about 23 generations for $F1$ and 58 for $F2$. Results obtained in our paper overpass the original results for $F1$, but conducts to poorer results for $F2$ than the ones presented in [8]. It has to be also outlined that the canonical species conservation algorithm reaches about 370 generations for $F1$ and 415 for $F2$, while TSC with global selection reaches around 180 generations for each of the two functions while the one with local selection leads to about 430. To conclude, due to the limit of fitness evaluations, the multinational undergoes a much smaller number of generations, which is most likely not enough to properly locate the desired optima.

Discussion: Although by tracking the subpopulation ID for each individual we gain a large amount of fitness evaluations, this also brings along a disadvantage. The algorithm starts with the subpopulations that are found in the initial generation and then continues with their preservation. The drawback lies in the fact that the *detect-multimodal* mechanism does not establish subpopulations in a completely accurate manner: It may happen that the two individuals for which the mechanism is applied follow two distant peaks and the interior points may situate on other hills between the ones of the initial individuals; if the interior points all have the fitness evaluation higher than the minimum evaluation

of the two initial points, then these are said to be in the same subpopulation. Thus, there may be the case that this happens to some individuals in the first generation and consequently they will bring noise to their subpopulations. In trying to minimize the erroneously placed individuals, one could apply *detect-multimodal* method each time the seeds are selected, although at the cost of a raise in the number of fitness evaluations.

To conclude, the TSC algorithm, in the way it is now, clearly gives better results than the multinational algorithm, as it has the advantage of much more generations for the same number of evaluation calls; at the same time, results obtained by the TSC algorithm are very competitive with the results given by the competitive species conservation algorithm, at the advantages of the elimination of the radius parameter that is hard to set up and of the expensive necessary runtime that is spent in distance measuring.

Experiment 2: Radius/vector size dependency.

Pre-experimental planning: In preliminary experiments, we established test cases on which species conservation and TSC were able to detect a larger number of optima. To enhance measurable differences, the required accuracy had to be lowered so that placing an individual inside the corresponding basin in the final population suffices to indicate a found optimum. As experiment 1 shows that the multinational performance is not competitive to the two others, we consider only species conservation and TSC. As first expe-

Table 3: Results for the TSC, original species conservation and the multinational algorithms

	F1	F2	F3	F4	F5
problem dimension	2d	2d	2d	5d	10d
TSC					
<i>Global selection</i>					
All peaks found (%)	100	96.67	100	100	100
Average number of peaks	10	5.97	1	1	1
Average best fitness	-	-	-10^{-12}	-10^{-10}	-10^{-5}
<i>Local selection</i>					
All peaks found (%)	96.67	100	100	100	100
Average number of peaks	9.97	6	1	1	1
Average best fitness	-	-	-10^{-13}	-10^{-9}	-10^{-5}
Species conservation					
All peaks found (%)	86.67	100	100	100	100
Average number of peaks	9.87	6	1	1	1
Average best fitness	-	-	-10^{-13}	-10^{-9}	-10^{-5}
Multinational					
<i>Global selection</i>					
All peaks found (%)	6.66	0	100	100	100
Average number of peaks	7.57	3.3	1	1	1
Average best fitness	-	-	-10^{-10}	-10^{-6}	-0.002
<i>Local selection</i>					
All peaks found (%)	30	96.66	100	100	100
Average number of peaks	9	5.97	1	1	1
Average best fitness	-	-	-10^{-10}	-10^{-6}	-10^{-4}

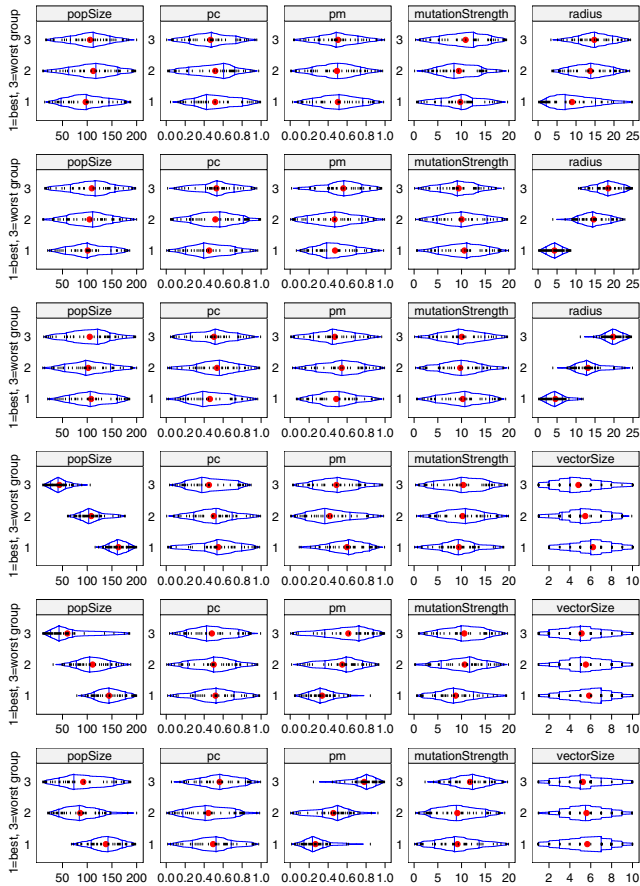


Figure 1: Parameter effect split plots, from top to bottom: Species conservation on F1, 2 dimensions, and F5, 10 and 20 dimensions, TSC on the same three problems. Pc and pm stand for crossover and mutation probability, respectively

Table 4: Peak detection (relaxed accuracy) in species conservation and TSC with local selection, values taken over 100 LHD samples/30 replicates

	non-uni	mean	max	min	std
F1 (2 d/10 peaks)					
species conserv.	0.12	1.32	10.0	1.0	1.37
TSC	1.0	7.87	10.0	3.17	1.67
F5 (10 d/32 p)					
species conserv.	0.42	7.20	32.0	1.0	11.33
TSC	1.0	2.72	5.67	1.07	1.06
F5 (20 d/1024 p)					
species conserv.	0.6	23.97	186	1.0	45.99
TSC	0.91	2.68	5.8	1.0	1.08

periments had not proved great differences between the global and local TSC variants, we concentrate on the local one.

Task: Detect how sensitive species conservation reacts to radius parameter values, and if the vector size parameter of TSC is similarly critical. We start with the hypothesis that performance is highly dependent on an accurately chosen radius value but abstain from the use of statistical tests due to the exploratory character of this experiment.

Setup: The investigated test problems are F1 in 2 dimensions, and F5 in 10 and 20 dimensions. For each of the two algorithms, we perform a LHD (Latin Hypersquare Design) in the allowed parameter space given in table 2. Each of the 100 design points is replicated 30 times and the mean number of detected peaks is recorded. In addition to standard measures, we also determine the non-unimodal end population fraction, that is the ratio of design points for which at least once out of the 30 replicates more than one peak is detected. All runs stop at a maximum of 10^5 evaluations.

Results/Visualization: Table 4 presents the obtained data for all functions and algorithms, including the non-unimodal fraction (non-uni column). In fig. 1, the 6 tested cases are depicted as effect split plots. The set of 100 design points is split into 3 (nearly) equal groups according to the chosen

quality criterion¹, and the single parameter values for these groups are drawn separately. This enables recognizing a concentration of certain parameter values that corresponds to the quality value, even if the effect is not linear, e.g. if best and worst performing design points share an extreme (i.e. small) parameter range whether medium configurations have this parameter set at opposite (i.e. large) values.

Observations: Interestingly, location of most peaks on the waves test problem (F1) seems to be much easier with the TSC than with the species conservation algorithm, as seen in tab. 4. For most (88%) of the design points of the latter, only one peak could be located. For the regularly spaced F5 test cases with larger search spaces, TSC still mostly attains more than one peak, but cannot be configured in a way that a large number of peaks is found, as is obviously possible for the species conservation algorithm. Concerning the parameter effects documented in fig. 1, we can state that the radius parameter becomes very important for high dimensions, whether the other species conservation parameters do not exhibit a clear structure. For the TSC, the parameter importance seems to depend on the problem at stake. For F1 and F5 in 10 dimensions, the population size is most relevant, while for F5 in 20 dimensions, the mutation probability must be set to a small value for achieving good performance. The vector size appears to be non-critical.

Discussion: At a first glance, it surprises to see the species conservation algorithm failing so often on F1 with only 2 search space dimensions. However, it may be necessary to consider the geometrical form of the basins of attraction in addition to number of peaks and the search space size. F1 possesses highly elliptical basins, so that obtaining an optimal radius for a 'search spheroid' is difficult if one cannot chose different values in the different dimensions. Interestingly, one may use this property to backwards deduce how spheroidal the basins are: If no radius can be obtained that delivers good performance, the basins may be formed too irregularly. However, species conservation performs clearly superior for the F5 test cases, if properly parameterized.

TSC has difficulties with keeping more than 5-10 separate subpopulations, while large population sizes are generally advisable. However, our parameter investigations indicate that the vector size parameter may be set to a fixed value (e.g. 1), so that only 4 parameters are left to consider instead of 5. Although not tested in this experiment, it can be expected that for non-spheroid and non-regularly spaced basin sets, the topological approach is better suited than the spheroidal distance-based approach. The results obtained with TSC on F1 point in this direction.

6. CONCLUSIONS AND FUTURE WORK

This paper is not primarily about a new algorithm. It is about comparing existing algorithms and the search for a method that provides a useful combination of their advantages. We therefore reimplemented these, namely the species conservation and the multinational (evolutionary) algorithm. To our knowledge, such a direct comparison has not been done before. From the obtained results, we can conclude that the suggested hybrid algorithm performs better than the original multinational approach in all test cases and approximately on the same level as the species conservation approach. For the irregularly formed basin struc-

ture of F1 (waves), the new approach is better while being simpler to configure, as it has no radius parameter. On problems with more regularly formed basin structures, the species conservation algorithm is slightly better because it is able to cover more basins if well parametrized.

A first step to extend the current work would be to consider more irregularly formed and/or spaced test problems in order to see where the turning point between the TSC and the original species conservation algorithm is. It may also be useful to drive the hybridization even further and insert techniques based on relative (instead of absolute) distances into the TSC. Furthermore, the ability to keep more subpopulations in different basins alive may be improved by reconsidering the handling of non-assigned (the 'free') individuals. Currently, only main parameter effects have been investigated. Once the parameter interactions are understood better, adding mutation stepsize adaptation mechanisms will surely open a path leading to substantially increased performance.

7. REFERENCES

- [1] BARTZ-BEIELSTEIN, T. *Experimental Research in Evolutionary Computation – The New Experimentalism*. Natural Computing Series. Springer, Berlin, 2006.
- [2] BEYER, K. S., GOLDSTEIN, J., RAMAKRISHNAN, R., AND SHAFT, U. When is "nearest neighbor" meaningful? In *International Conference on Database Theory (ICDT 1999)* (London, UK, 1999), Springer-Verlag, pp. 217–235.
- [3] EIBEN, A. E., AND SMITH, J. E. *Introduction to Evolutionary Computing*. Springer-Verlag, Berlin, Heidelberg, NY, 2003.
- [4] LI, J. P., BALAZS, M. E., PARKS, G. T., AND CLARKSON, P. J. A species conserving genetic algorithm for multimodal function optimization. *Evolutionary Computation* 10, 3 (2002), 207–234.
- [5] PREUSS, M. Niching Prospects. In *Bioinspired Optimization Methods and their Applications (BIOMA 2006)* (2006), Jozef Stefan Institute, Ljubljana, Slovenia, pp. 25–34.
- [6] PREUSS, M., NAUJOKS, B., AND RUDOLPH, G. Pareto set and EMOA behavior for simple multimodal objective functions. In *Parallel Problem Solving from Nature (PPSN 2006), Lecture Notes in Computer Science* (Berlin, Germany, 2006), vol. 4193, Springer-Verlag, pp. 513–522.
- [7] SHIR, O. M., AND BAECK, T. Niche radius adaptation in the cma-es niching algorithm. In *Parallel Problem Solving from Nature (PPSN 2006), Lecture Notes in Computer Science* (Berlin, Germany, 2006), vol. 4193, Springer-Verlag, pp. 142–151.
- [8] URSEM, R. K. Multinational evolutionary algorithms. In *Congress of Evolutionary Computation (CEC 1999)* (Piscataway, NJ, 1999), vol. 3, IEEE Press, pp. 1633–1640.
- [9] URSEM, R. K. Multinational GAs: Multimodal optimization techniques in dynamic environments. In *Genetic and Evolutionary Computation Conference (GECCO 2000)* (2000), vol. 1, Morgan Kaufmann, pp. 19–26.

¹measured mean number of peaks in this case