

# An Analysis on the Effect of Selection on Exploration in Particle Swarm Optimization and Differential Evolution

Stephen Chen

*School of Information Technology*  
York University  
Toronto, Canada  
sychen@yorku.ca

Antonio Bolufé-Röhler

*School of Mathematical and Computational Sciences*  
University of Prince Edward Island  
Charlottetown, Canada  
aboluferohler@upei.ca

James Montgomery

*School of Technology, Environments and Design*  
University of Tasmania  
Hobart, Australia  
James.Montgomery@utas.edu.au

Tim Hendtlass

*Department of Computer Science and Software Engineering*  
Swinburne University of Technology  
Hawthorn, Australia  
thendtlass@swin.edu.au

**Abstract**—The goal of exploration to produce diverse search points throughout the search space can be countered by the goal of selection to focus search around the fittest current solution(s). In the limit, if all exploratory search points are rejected by selection, then the behaviour of the metaheuristic will be equivalent to one which performs no exploration at all (e.g. hill climbing). The effects of selection on exploration are clearly important, but our review of the literature indicates limited coverage. To address this deficit, we introduce new experiments which can specifically highlight the occurrence of “failed exploration” and its effects through selection that can trap a metaheuristic in a less promising part of the search space. We subsequently propose new lines of research to reduce the effects of selection and failed exploration which we believe are distinctly different from traditional lines of research to increase (pre-selection) exploration.

**Index Terms**—exploration, selection, metaheuristics, continuous domain search spaces

## I. INTRODUCTION

Exploration plays a critical role in the performance of metaheuristics. Two key exemplars of metaheuristics are Particle Swarm Optimization (PSO) [1] and Differential Evolution (DE) [2], and the role of exploration in these techniques has been studied extensively as can be seen in recent surveys [3], [4]. However, in their survey of PSO, Bonyadi and Michalewicz also raise the concern that much of this research has “only presented experimental results [on benchmark problem sets] and did not provide adequate discussion (neither from theoretical perspective, nor general discussions) on merits of the proposed approach.” [3] Extending from this perspective that final performance is too indirect a method to measure the effects of exploration during the search process, we believe that this entire line of research could benefit from a more direct experimental analysis of exploration.

The ability to measure exploration (and exploitation) is hindered by their lack of precise definitions. In particular, a broad survey of over 100 papers led Crepinšek, Liu, and Mernik to the unexpected conclusion that “The fact that until now exploration and exploitation have only been implicitly defined in EAs comes as a big surprise.” [5] Our definitions for continuous domains are based on dividing a multi-modal search space into attraction basins which each have a single local optimum. Each point in an attraction basin has a monotonic path of increasing (for maxima) or decreasing (for minima) fitness to its local optima. A search point (e.g. the current position of a particle) is then defined to be performing “exploration” if it is in a different attraction basin than its reference solution(s) (e.g. a particle’s *pbest* position), and it is defined to be performing “exploitation” if it is in the same attraction basin as (one of) its reference solution(s).

Recent research [6] has built upon the above definitions to highlight the crucial effects of selection on the exploratory activities of metaheuristics. The role of selection (as implemented in PSO, DE, and many other metaheuristics) is to accept or reject a candidate search solution. In the case of PSO and DE, this comparison is made with respect to a reference solution (i.e. *pbest* or *target*) which will be replaced if the search solution is accepted. For an exploratory search solution, the act of selection can lead to four possible outcomes: successful exploration, successful rejection, deceptive exploration, and failed exploration.

These four categories occur because selection is based on the fitness of the reference solution and the fitness of the search solution, whereas the purported goal of exploration is to find the most promising region of a search space. Our definitions do not define the “promise” of a region of a search space, or even what might be a “region of a search space”. However, we

can define an attraction basin and the fitness of an attraction basin (i.e. the fitness of the locally optimal solution within an attraction basin). For an exploratory search solution, these four possible outcomes (represented pictorially in Fig. 1) are described as follows for a minimization problem:

- Successful exploration: A fitter attraction basin represented by a fitter exploratory search solution is accepted to replace the less fit reference solution from a less fit attraction basin.
- Successful rejection: A less fit attraction basin represented by a less fit exploratory search solution is rejected to keep the fitter reference solution from a fitter attraction basin.
- Deceptive exploration: A less fit attraction basin represented by a fitter exploratory search solution is accepted to replace the less fit reference solution from a fitter attraction basin.
- Failed exploration: A fitter attraction basin represented by a less fit exploratory search solution is rejected to keep the fitter reference solution from a less fit attraction basin.

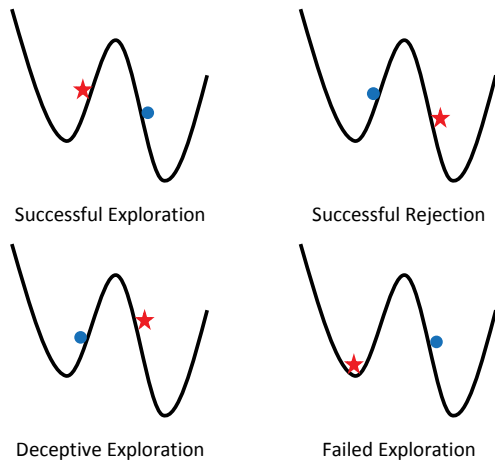


Fig. 1. Four categories of exploration after selection. In each sub-figure, the red star represents the reference solution and the blue dot represents the exploratory search solution.

We believe that the case of “failed exploration” provides particular insight into the operation of metaheuristics. Specifically, the existence of failed exploration demonstrates that exploration is occurring in metaheuristics, so the standard focus of increasing exploration may be misdirected. We instead propose new lines of research to increase the rates of successful exploration/decrease the rates of failed exploration. The potential for these lines of research is highlighted by an experiment which eliminates the occurrence of failed exploration.

Before these experiments, we begin with a background into typical perspectives of exploration in PSO and DE. We also introduce the Rastrigin-based experiments that are used to observe and measure failed exploration. These experimental techniques are applied to PSO in Section III and to DE in

Section IV. We then conduct an experiment to show how metaheuristics could theoretically perform if failed exploration could be eliminated. Our discussion in Section VI contrasts traditional lines of research to increase and/or improve exploration with our suggested lines for future research. The paper then closes with a brief summary.

## II. BACKGROUND

This background section begins with some examples of how the term “exploration” is typically used in the literature for Particle Swarm Optimization and Differential Evolution. A key feature of existing discussions of exploration is that they are unconcerned with the effects of selection. In particular, the effects of selection are greatly influenced by exploitation in known attraction basins. The second subsection reviews an experiment designed to highlight the effects of exploitation and selection on future rates of failed exploration.

### A. Exploration and Exploitation in the Literature

The following are several examples of how the term exploration is typically used in the literature for Particle Swarm Optimization (PSO) [1]. “A larger inertia weight  $\omega$  facilitates global exploration (searching new areas) while a smaller inertia weight tends to facilitate local exploration to fine-tune the current search area.” [7] “Diversity is related to the notions of exploration and exploitation: the more diverse a swarm is, the more its particles are dispersed over the search space, and the more the swarm is exploring.” [8] “exploration: global search, exploring all over the search space to find promising regions and exploitation: local search, exploiting the identified promising regions to fine tune the search for the optimal solution” [9].

The literature for Differential Evolution (DE) [2] uses the term exploration with a similar connection to diversity and parametric variations as shown in the following examples. “PPCea ... increases the scaling factor  $F$  and control parameter for crossover  $Cr$  to further explore unvisited regions of the search space.” [10] “The desired equilibrium between exploration and exploitation in evolutionary algorithms can be achieved by controlling the population diversity.” [11] “At the beginning of the evolution process, the mutation operator of DE favors exploration. As evolution progresses, the mutation operator favors exploitation...” [12].

The search mechanisms used in PSO and DE lead to distinctly different search trajectories. To create a common context for all possible search trajectories, we focus only on an individual step which involves a new search solution. Specifically, we refer to search solutions and reference solutions, and the classification of a search solution as exploratory or exploitative occurs only in the context of the given reference solution(s). For example, a high-speed particle which is nominally performing exploration can still visit the attraction basin of its attractor/reference solution, and a diverse population of  $n$  hill-climbers conducts no exploration.

We believe that the provided examples of exploration from the literature are representative and that they demonstrate

two key weaknesses. The first weakness is that the presented concepts for exploration largely ignore the effects of selection, i.e. the process by which search solutions are accepted or rejected in their bid to become reference solutions. In DE, a population that is stuck at a local optimum might execute a series of jumps that are repeatedly rejected, so its search path might look like a star burst emanating from the original reference solution. In PSO, the search trajectories of particles may visit the periphery of many promising attraction basins, but all of this exploration is wasted effort if the reference solutions are never updated. [13]

The second weakness is that none of these common usages for exploration includes a verification that a solution intended to be exploratory is indeed in a different attraction basin. When most researchers use the term “exploration”, they really mean “potential for exploration”. Much like exploration pre-selection and post-selection, the potential for exploration is about the inputs for the search mechanism (e.g. random selection from a large range) as opposed to the output from that mechanism (e.g. a search solution that is in a distinctly different part of the search space from existing reference solutions). Without this verification, search solutions created with mechanisms that have the “potential for exploration” could instead lead to the creation of search solutions that our definitions would classify as “exploitation”.

### B. The Effect of Exploitation on Failed Exploration

The early exploitation of known attraction basins is actually an important aspect of some metaheuristics. For instance, one of the key features of DE is that it is self-scaling [2]: as population members cluster around a subset of local optima in the search space, the likelihood of producing shorter difference vectors increases, and short difference vectors are likely to lead to more exploitative solutions than exploratory solutions [12], [14]. This clustering is likely initially caused by a search solution (if the cluster represents multiple solutions in the same attraction basin) that was generated by a search action (e.g. the use of a large difference vector) that was intending to produce an exploratory search solution. Higher values of the difference vector scale  $F$  can prolong apparently exploratory behavior, but even  $F > 1$  does not prevent the population’s convergence [14]. In general, the ability to converge through self-scaling is viewed as a positive feature of DE, and this feature is often initiated by exploitation in known attraction basins.

The negative effects of exploitation in known attraction basins are presented through a detailed study on the Rastrigin function in [6]. The Rastrigin function shown in Equation 1 has a regular fitness landscape in which every point with integer values in all dimensions is a local optimum, and every other point belongs to the attraction basin of the local optimum that is determined by rounding each solution term to its nearest integer value. These features make it possible to quickly and easily determine the attraction basin of a search point and the fitness of the local optimum of this attraction basin. For

brevity, we will use “the fitness of an attraction basin” to mean the fitness of the (local) optimum of an attraction basin.

$$f(x) = 10n + \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i)) \quad (1)$$

Metaheuristics which accept or reject a search solution based on a comparison of its fitness with that of a reference solution expose themselves to the four categories of exploration described in Fig. 1. The study in [6] focuses on the category of failed exploration by generating a set of exploratory search solutions in attraction basins whose fitness is better than the reference solution’s attraction basin. It then counts how many of these exploratory search solutions would be accepted or rejected based on a comparison of the actual fitness of the search and reference solutions. The potential effects of exploitation and elitism are simulated by improving the fitness of the reference solution.

For a reference solution that has not experienced any local optimization (e.g. the first reference solution in a given attraction basin), the study in [6] shows that failed exploration occurs less than 50% of the time. However, if the reference solution is moved as little as 50% of the distance towards its local optimum (e.g. through exploitation), the rate of failed exploration can increase to over 99%. Fig. 2 helps illustrate this effect of how exploitation of reference solutions can lead to increased rates of failed exploration.




Fig. 2. Exploitation of a reference solution (red star on left) lowers the probability that future exploratory solutions from fitter attraction basins (blue circles in middle) will survive selection, and this rejection represents increasing rates of failed exploration.

Failed exploration is a problem in metaheuristics such as PSO and DE that do not maintain any information from rejected search solutions. For example, in an  $n$ -tuple bandit situation, Monte Carlo Tree Search can balance between known and unexplored states [15]. Further, an ant that leads to a rejected overall solution in Ant Colony Optimization (ACO) [16] can still update pheromone levels (if only by passive decay), so exploration and exploitation can occur in tandem without the negative effects presented above. However, the use of elitism in PSO and DE (among many other metaheuristics) leads to no benefits from failed exploration. Exploitation in

known attraction basins, which leads to increased rates of failed exploration, can thus reduce the effectiveness of these metaheuristics.

### III. SELECTION AND EXPLORATION IN PSO

The previously described experimental procedure using the Rastrigin function is now applied to Particle Swarm Optimization (PSO). Unlike the normal operation of PSO in which only the fitness and position of each solution is known, we can now observe and record the attraction basin and the fitness of that attraction basin for all solutions. This information allows us to explicitly state whether a current position represents exploration or exploitation with respect to its two attractors/reference solutions (i.e. its  $pbest$  and  $lbest$  positions), and to determine the effect of selection on an exploratory search solution (i.e. a classification into one of the four sub-categories of successful exploration, successful rejection, deceptive exploration, or failed exploration).

These experiments use a version of standard particle swarm optimization [17] with a ring topology. The key parameters specified from this standardization are  $\chi = 0.72984$ , and  $c_1 = c_2 = 2.05$  for the velocity updates given in Equation 2. Additional implementation details are the use of  $p = 50$  particles [17], zero initial velocities [18] and “Reflect-Z” for particles that exceed the boundaries of the search space (i.e. reflecting the position back into the search space and setting the velocity to zero) [19]. The source code for this implementation is available online [20].

$$v_{i+1,d} = \chi \{ v_{i,d} + c_1 \epsilon_1 (pbest_{i,d} - x_{i,d}) + c_2 \epsilon_2 (lbest_{i,d} - x_{i,d}) \} \quad (2)$$

Experiments involve 30 independent trials in  $n = 30$  dimensions using a fixed limit of  $10,000 \cdot n$  total function evaluations (FEs). Fig. 3 shows averages for the fitness of the best overall solution, the fitness of its attraction basin, and the fittest attraction basin represented by any  $pbest$  position. The key observation that we wish to highlight in this figure is the rapid convergence between the first and second plot lines. This convergence indicates that the best overall solution in PSO quickly approaches the local optimum in its attraction basin. Since this best overall solution is also a  $pbest$  reference solution, the study in [6] indicates that selection will cause future opportunities for successful exploration to become almost 100% cases of failed exploration instead (see Fig. 2). We will now show that failed exploration continues to occur in PSO even after the swarm has appeared to converge (i.e. all three plot lines have plateaued at the same level), but it is difficult to see this activity in aggregated data sets.

It is typical to present the results of experiments as the average of a statistically relevant set of trials (e.g. 30). However, many insights into the operation of metaheuristics can be lost through the process of averaging their behaviour. Two single trials have been selected from the 30 used to create Fig. 3 to highlight both good and bad search behaviours that might be observed in PSO under typical operating conditions. The next




Fig. 3. PSO, average of 30 trials. The difference between the second and third plot lines shows when the best solution is not from the fittest attraction basin that is represented by the  $pbest$  positions. The upward movement of the third plot line before convergence indicates that highly fit attraction basins routinely have their  $pbest$  positions moved to less fit attraction basins

two figures plot the actual fitness of the current best overall solution, the fitness of its attraction basin, and the fitness of the best attraction basin of any  $pbest$  position (e.g. an estimate of PSO’s potential performance with improved exploitation [13]).

In addition to the three plot lines (note: the second plot line is often obscured by the first plot line), individual dots display the fitness of attraction basins represented by exploratory search solutions (with respect to both  $pbest$  and  $lbest$  positions) that were rejected under the category of failed exploration. It is noted that the bands of the dots are a feature of the Rastrigin function’s search space which has some levels of fitness for attraction basins occur more frequently than others. Specifically, the existence of the bands in these figures does not imply anything about the reference solutions that led to the shown instances of failed exploration (e.g. having a fitness at the level of the gap).

Fig. 4 shows an example of PSO when it is operating well. First, clusters of failed exploration around the 100 error level stop around FE 40,000 which suggests that straggling particles have been able to move towards much fitter attraction basins. There is a clear step near FE 50,000 showing a new fittest overall particle being found in the best overall attraction basin. There is also very little failed exploration for attraction basins that are fitter than the best overall solution. Conversely, Fig. 5 shows a particularly bad example of stalling in PSO. There are two clear examples around FE 20,000 and FE 50,000 of  $pbest$  positions from highly fit attraction basins being eliminated from the swarm (i.e. deceptive exploration). There are also two large bands of dots representing a large amount of failed exploration occurring in attraction basins that are much fitter than the best overall solution. These dots indicate that PSO is not suffering from a lack of exploration, and it is instead suffering from an inability to achieve the category of successful exploration for the exploration that it is performing.




Fig. 4. PSO, single trial that highlights “good” operation. The fittest particle at least arrives at the local optimum of the fittest attraction basin among all of the  $pbest$  positions

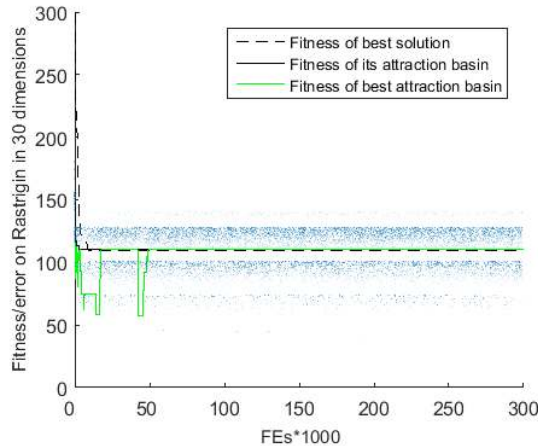


Fig. 5. PSO, single trial that highlights “poor” operation. Many fitter attraction basins are visited by the particles, and some of these attraction basins are even represented by  $pbest$  positions for a while. However, all of these highly promising exploratory search solutions do not affect the overall search trajectory, and the performance of the swarm does not benefit from this (pre-selection) exploration

#### IV. SELECTION AND EXPLORATION IN DE

We again use the Rastrigin function to allow us to conduct an in-depth analysis of exploration and the effects of selection on Differential Evolution (DE). There are obviously many limits to this study which are based on a single function. In particular, Rastrigin is globally convex, has attraction basins of similar size and shape, and it is highly regular and symmetric. Nonetheless, we believe the insights are relevant, unique, and likely to extend to a broad range of multi-modal fitness landscapes and problem domains. In particular, the four sub-categories for exploration that result from selection are not specifically correlated to the limitations of the Rastrigin function. Thus, even though the exact effects will of course vary for different functions, the ability to observe failed exploration

on the Rastrigin function provides useful insight into how DE might face similar challenges in other multi-modal search spaces.

$$y_i = r_1 + F(r_2 - r_3) \quad (3)$$

Our experiments with DE use an implementation of DE/rand/1/bin with typical parameters of population size  $p = 50$ , crossover  $Cr = 0.9$ , and scale factor  $F = 0.8$  [2], [21]. Each population member  $x_i$  is considered as a *target* for replacement by a candidate solution that is constructed in two steps: creation of an intermediate solution and crossover with  $x_i$ . During the creation of an intermediate solution  $y_i$  from three distinct random solutions  $r_1$ ,  $r_2$ , and  $r_3$  in Equation 3, the scale factor  $F$  affects the “step size” from  $r_1$  taken in the direction of the “difference vector” created with  $r_2$  and  $r_3$ .

$$x'_{i,d} = \begin{cases} y_{i,d} & u_d \leq Cr \\ x_{i,d} & u_d > Cr \end{cases} \quad (4)$$

Equation 4 defines how this intermediate solution is then crossed term-by-term in each dimension  $d$  of the search space with the target solution  $x_i$  to produce a new search/offspring solution  $x'_i$ .

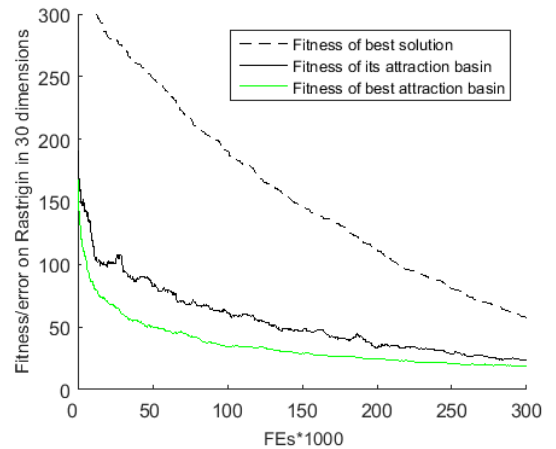


Fig. 6. DE, average of 30 trials. The difference between the second and third plot lines shows that the best solution is often not from the fittest attraction basin that is represented by any member of the population

The results for DE are shown in Fig. 6 which represents the average performance of 30 trials on Rastrigin in  $n = 30$  dimensions. It can be seen that the convergence of the fitness of the best solution (first plot line) with the fitness of its attraction basin (second plot line) occurs much more slowly in DE than in PSO. According to the study in [6], this lack of convergence can help reduce failed exploration, and it can also be seen that successful exploration (as evidenced by the downward trending of the second and third plot lines) ends later in DE than in PSO. The nature of convergence in DE benefits from further analysis of individual trials.

We again highlight two individual trials from the 30 used to produce the aggregate results shown in Fig. 6. The dots in

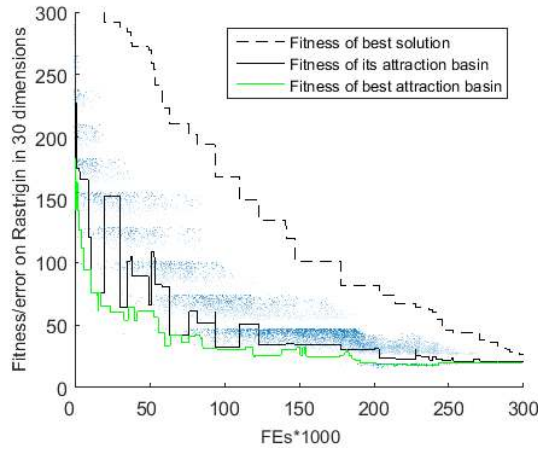


Fig. 7. DE, single trial that highlights “good” operation. The third plot line has relatively few upward motions which occur when the fittest attraction basin represented in the population is replaced by a new reference solution from a less fit attraction basin. There are also relatively few dots representing failed exploration

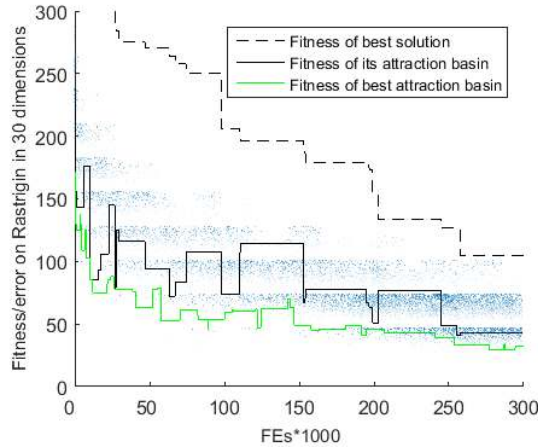


Fig. 8. DE, single trial that highlights “poor” operation. There are a larger number of upward movements in the third plot line, more dots representing failed exploration, and less convergence between the first and second plot lines. All of these factors reduce the performance of DE

these plots represent instances of failed exploration where the exploratory nature of a search solution is first determined with respect to the *target* reference solution. The first individual trial (Fig. 7) demonstrates good performance in DE and the second individual trial (Fig. 8) demonstrates poor performance in DE. In Fig. 7, the three plot lines converge which means that the best solution is approaching the local optimum in its attraction basin, and that this attraction basin is the best attraction basin among those represented by the entire population. There are also very few dots for failed exploration below the second and third plot lines. In contrast, Fig. 8 shows more dots of failed exploration below the third plot line, a persistent gap between the second and third plot lines, and a slower convergence of the first plot line to the second. These

observations suggest that multiple attraction basins are still represented in the final population and that DE is having a difficult time finding the local optimum in any of them.

It should also be noted that between DE (see Fig. 7 and Fig. 8) and PSO (see Fig. 4 and Fig. 5), the second and third plot lines in DE show much more oscillation. Upward movements in these plot lines which represent the fitness of the attraction basins indicate that deceptive exploration has occurred. Reference solutions which approach local optima greatly reduce both successful exploration and deceptive exploration. However, even though the best overall attraction basin is often lost in DE, population diversity can support ongoing improvement in the third plot line. We therefore believe that the effects of (increased) deceptive exploration are sufficiently small in popular population-based metaheuristics (e.g. PSO and DE) that our primary focus can now be to observe and reduce the effects of failed exploration.

## V. ELIMINATING FAILED EXPLORATION IN PSO AND DE

The key contribution of the preceding experiments has been to allow measurements and observations on the fitness of attraction basins. If the goal of exploration is to find “promising areas of the search space”, the fitness of the exact search solution which is sampled from this promising area should be irrelevant. By being able to measure both the fitness of a solution and the fitness of its attraction basin, it is now possible to specifically observe when the performance of a metaheuristic is due to exploration (e.g. finding new attraction basins with better fitness) or exploitation (e.g. finding a fitter solution within an existing attraction basin).

The experimental design also allows the amount of failed exploration to be observed. The plotted dots in the individual trials (e.g. Fig. 4 and Fig. 5 for PSO) represent the fitness of attraction basins for rejected exploratory search solutions that are fitter than the attraction basins of their reference solutions. When these dots are below the third plot line, they represent a missed opportunity to move the overall search trajectory of the metaheuristic closer to the global optimum. These dots also indicate that metaheuristics which appear to have converged (e.g. see Fig. 5) might still be performing large amounts of exploration. Since a lack of exploration is not the limiting factor in the performance of these metaheuristics, attempts to increase (pre-selection) exploration (e.g. [3], [4]) might not be the most effective means by which to improve the performance of metaheuristics. We believe directly addressing and reducing the levels of failed exploration might be a more productive line of research.

The potential of this line of research is highlighted by the following experiment which eliminates the occurrence of failed exploration (and also deceptive exploration) in Particle Swarm Optimization and Differential Evolution. This elimination is achieved by a simple modification to the objective function. Taking advantage of the Rastrigin function which allows us to easily determine the attraction basin for every location in the search space, we make the fitness of a solution to be the fitness of its attraction basin (with ties broken

by their actual fitness). In this experiment, the previously described standard implementations of PSO and DE are still used, and this allows us to isolate the effects of selection. The search/exploration mechanisms employed by PSO and DE still generate search solutions as before, but the modified fitness function makes failed exploration (and deceptive exploration) impossible.

The modified function evaluation is implemented for Rastrigin in  $n = 30$  dimensions. Thirty independent trials are performed using the same termination condition of  $10,000 \cdot n$  total function evaluations. Standard PSO converges to the global optimum in 28 of the 30 trials (and to the second best local optimum in the other two trials). Fig. 9 shows that this convergence is rapid and monotonic when there is no failed exploration, and that the fitness of the attraction basin of the best solution/particle (e.g.  $g_{best}$ ) matches the fittest attraction basin of any particle (i.e. all of the attraction basins represented by the  $p_{best}$  positions).

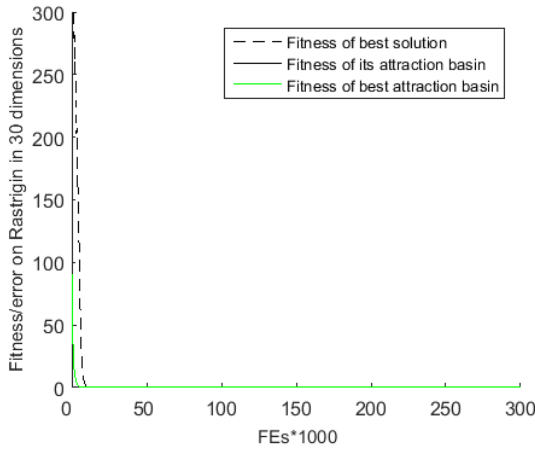


Fig. 9. PSO, average of 30 trials with failed exploration eliminated. The fittest solution is always from the fittest known attraction basin, so the second and third plot lines are identical

The implemented canonical version of DE is able to converge to the global optimum for all 30 trials with the modified fitness function. Fig. 10 shows that this convergence is rapid and monotonic, and that the fitness of the attraction basin of the best overall solution matches the fittest attraction basin for all of the solutions in the population. The key purpose of these experiments is to demonstrate that canonical versions of PSO and DE can find the global optimum in a multi-modal search space when failed exploration (and deceptive exploration) is eliminated. Compared to the vast amounts of research that attempt to increase (pre-selection) exploration [3], [4], it appears that reducing failed exploration could be a more promising line of attack.

## VI. DISCUSSION

The conducted experiments leverage the simplicity of the Rastrigin function, which allows the attraction basin and the fitness of that attraction basin to be easily determined for every

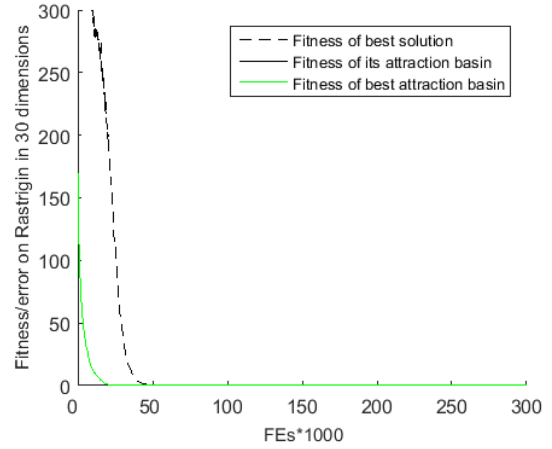


Fig. 10. DE, average of 30 trials with failed exploration eliminated. The fittest solution is always from the fittest known attraction basin, so the second and third plot lines are identical

location in the search space. For more difficult problems (e.g. composition functions [22]), the convoluted topologies can make similar experiments impractical. However, increasing problem difficulty tends to expose new limitations of metaheuristics (e.g. the curse of dimensionality [23]) as opposed to reducing the effects of limitations exposed on simpler problems. We are thus confident that the limitations identified by our current experiments on the Rastrigin function will be relevant to any metaheuristic which can experience failed exploration.

Our experiments focus on the performance of PSO and DE, which are two distinctly different metaheuristics. In the experiments of Section V, the fitness evaluation has been changed to eliminate the possibility of “failed exploration” (and deceptive exploration). When it is no longer possible to reject an exploratory search solution from a fitter attraction basin because its fitness is worse than the fitness of a reference solution from a less fit attraction basin, both PSO and DE rapidly converge towards the global optimum. This result suggests that the standard versions of PSO and DE do not suffer from a lack of exploration.

The experiments on PSO in Section III and on DE in Section IV show how well these metaheuristics perform in practice. In particular, it should be noted that it is possible for both PSO and DE to continuously perform exploration even though their nominal performance (e.g. as represented by the fitness of the best overall solution) has stalled. These observations demonstrate the effects of failed exploration. However, a large amount of the existing literature that attempts to improve the performance of PSO and DE focuses on how solutions are created and/or the diversity of the population. The effects of selection and of failed exploration receive much less attention in the literature.

In contrast to the extensive body of research which largely ignores the effects of selection [3], [4], we propose two alternate lines of research. Failed exploration occurs when an

exploratory search solution is rejected based on its fitness. However, as much of the above work aims to increase diversity, it is noted that a solution's fitness measures neither its diversity (e.g. minimum distance away from a current member of the population) nor novelty (e.g. being in a new area/attraction basin regardless of distance from other solutions). Therefore, attempts to address diversity or novelty which do not also address the unrelated and often counter productive goals of fitness-based selection criteria may not be able to overcome this critical handicap. Fitness is a poor measure of the novelty required for exploration. One line of suggested research is to develop diversity or novelty based criteria for the selection/survival of exploratory search solutions.

Another line of suggested research is to design new selection mechanisms for metaheuristics. For example, Leaders and Followers [6] attempts to reduce failed exploration by restricting comparisons to search solutions and reference solutions with similar relative fitness with respect to their local optima. The assumption is that the relative fitness of followers will be similar to the relative fitness of leaders when these two populations have similar median fitness. Although this assumption has not been analyzed in detail, LaF does perform better than base versions of PSO and DE in multi-modal search spaces [6]. LaF is also simpler than these base versions (e.g. has fewer parameters), let alone the modified versions of these metaheuristics (e.g. with adaptive parameters [4]). To avoid the proliferation of new metaphors and metaheuristics [24], this line of research should attempt to focus on existing categories for search techniques [25].

## VII. SUMMARY

Large amounts of research has been dedicated towards increasing "exploration" in metaheuristics. However, most of this research focuses on exploration pre-selection, and on the evaluation of exploratory search solutions based on their fitness. We believe the effects of selection and its creation of the four sub-categories of successful exploration, deceptive exploration, successful rejection, and failed exploration are also critical for the study of exploration in metaheuristics. Our experimental analysis shows how the elimination of failed exploration in an ideal case can greatly improve the performance of PSO and DE in a multi-modal search space. We suggest that the design of new selection mechanisms, especially those based on new metrics which can more directly measure the potential of exploratory search solutions, should be pursued as a direction for future research.

## REFERENCES

- [1] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Neural Networks, 1995. Proceedings., IEEE International Conference on*, vol. 4, IEEE, 1995, pp. 1942–1948.
- [2] R. Storn and K. Price, "Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces," *Journal of global optimization*, vol. 11, no. 4, pp. 341–359, 1997.
- [3] M. R. Bonyadi and Z. Michalewicz, "Particle swarm optimization for single objective continuous space problems: a review," 2017.
- [4] S. Das, S. S. Mullick, and P. N. Suganthan, "Recent advances in differential evolution—an updated survey," *Swarm and Evolutionary Computation*, vol. 27, pp. 1–30, 2016.
- [5] M. Črepinšek, S.-H. Liu, and M. Mernik, "Exploration and exploitation in evolutionary algorithms: a survey," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, p. 35, 2013.
- [6] Y. Gonzalez-Fernandez and S. Chen, "Leaders and followers — a new metaheuristic to avoid the bias of accumulated information," in *Evolutionary Computation (CEC), 2015 IEEE Congress on*. IEEE, 2015, pp. 776–783.
- [7] Y. Shi and R. Eberhart, "Parameter selection in particle swarm optimization," in *Evolutionary programming VII*. Springer, 1998, pp. 591–600.
- [8] P. Bosman and A. P. Engelbrecht, "Diversity rate of change measurement for particle swarm optimisers," in *International Conference on Swarm Intelligence*. Springer, 2014, pp. 86–97.
- [9] N. Lynn and P. N. Suganthan, "Heterogeneous comprehensive learning particle swarm optimization with enhanced exploration and exploitation," *Swarm and Evolutionary Computation*, vol. 24, pp. 11–24, 2015.
- [10] S.-H. Liu, M. Mernik, D. Hrnčič, and M. Črepinšek, "A parameter control method of evolutionary algorithms using exploration and exploitation measures with a practical application for fitting sovova's mass transfer model," *Applied Soft Computing*, vol. 13, no. 9, pp. 3792–3805, 2013.
- [11] D. Zaharie, "Control of population diversity and adaptation in differential evolution algorithms," in *Proc. of MENDEL*, vol. 9, 2003, pp. 41–46.
- [12] A. Salman, A. P. Engelbrecht, and M. G. Omran, "Empirical analysis of self-adaptive differential evolution," *European Journal of operational research*, vol. 183, no. 2, pp. 785–804, 2007.
- [13] D. Tamayo-Vera, S. Chen, A. Bolufé-Röhler, J. Montgomery, and T. Hendtlass, "Improved exploration and exploitation in particle swarm optimization," in *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*. Springer, 2018, pp. 421–433.
- [14] J. Montgomery, "Differential evolution: Difference vectors and movement in solution space," in *Evolutionary Computation, 2009. CEC'09. IEEE Congress on*. IEEE, 2009, pp. 2833–2840.
- [15] K. Kuanusont, R. D. Gaina, J. Liu, D. Perez-Liebana, and S. M. Lucas, "The n-tuple bandit evolutionary algorithm for automatic game improvement," in *Evolutionary Computation (CEC), 2017 IEEE Congress on*. IEEE, 2017, pp. 2201–2208.
- [16] M. Dorigo, V. Maniezzo, and A. Coloni, "Ant system: optimization by a colony of cooperating agents," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 26, no. 1, pp. 29–41, 1996.
- [17] D. Bratton and J. Kennedy, "Defining a standard for particle swarm optimization," in *Swarm Intelligence Symposium, 2007. SIS 2007. IEEE*. IEEE, 2007, pp. 120–127.
- [18] A. Engelbrecht, "Particle swarm optimization: Velocity initialization," in *Evolutionary Computation (CEC), 2012 IEEE Congress on*. IEEE, 2012, pp. 1–8.
- [19] S. Helwig, J. Branke, and S. Mostaghim, "Experimental analysis of bound handling techniques in particle swarm optimization," *IEEE Transactions on Evolutionary computation*, vol. 17, no. 2, pp. 259–271, 2013.
- [20] "Pso source code," June 2017. [Online]. Available: [https://www.researchgate.net/publication/259643342\\_Source\\_code\\_for\\_an\\_implementation\\_of\\_Standard\\_Particle\\_Swarm\\_Optimization\\_-\\_revised?ev=prf\\_pub](https://www.researchgate.net/publication/259643342_Source_code_for_an_implementation_of_Standard_Particle_Swarm_Optimization_-_revised?ev=prf_pub)
- [21] J. Montgomery and S. Chen, "An analysis of the operation of differential evolution at high and low crossover rates," in *Evolutionary Computation (CEC), 2010 IEEE Congress on*. IEEE, 2010, pp. 1–8.
- [22] J. Liang, B. Qu, P. Suganthan, and A. G. Hernández-Díaz, "Problem definitions and evaluation criteria for the cec 2013 special session on real-parameter optimization," *Computational Intelligence Laboratory, Zhengzhou University, Zhengzhou, China and Nanyang Technological University, Singapore, Technical Report*, vol. 201212, pp. 3–18, 2013.
- [23] S. Chen, J. Montgomery, and A. Bolufé-Röhler, "Measuring the curse of dimensionality and its effects on particle swarm optimization and differential evolution," *Applied Intelligence*, vol. 42, no. 3, pp. 514–526, 2015.
- [24] K. Sörensen, "Metaheuristics—the metaphor exposed," *International Transactions in Operational Research*, vol. 22, no. 1, pp. 3–18, 2015.
- [25] M. A. Lones, "Metaheuristics in nature-inspired algorithms," in *Proceedings of the Companion Publication of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 2014, pp. 1419–1422.