The Influence of Fitness Landscape Characteristics on the Search Behaviour of Particle Swarm Optimisers

by

Phlippie Bosman

Submitted in partial fulfillment of the requirements for the degree
Magister Scientae (Computer Science)
in the Faculty of Engineering, Built Environment and Information Technology
University of Pretoria, Pretoria

2019

Publication data:

Phlippie Bosman. The Influence of Fitness Landscape Characteristics on the Search Behaviour of Particle Swarm Optimisers. Master's dissertation, University of Pretoria, Department of Computer Science, Pretoria, South Africa, 2019.

Electronic, hyperlinked versions of this dissertation are available online, as Adobe PDF files, at:

http://cirg.cs.up.ac.za/

http://upetd.up.ac.za/UPeTD.htm

The Influence of Fitness Landscape Characteristics on the Search Behaviour of Particle Swarm Optimisers

by

Phlippie Bosman
E-mail: pbosman@cs.up.ac.za

Abstract

Fitness landscapes facilitate the analysis of optimisation problems in a detailed, yet intuitive manner. Such an analysis can be used to select an appropriate algorithm to solve the problem, based on the strengths and weaknesses of the algorithm. This requires an understanding of the effect of the various fitness landscape characteristics (FLCs) of the problem on the behaviour of the algorithm being considered. The effects of FLCs on the behaviour of particle swarm optimisers (PSOs) is still not well-understood. This dissertation uses a novel measure of PSO behaviour in terms of exploration and exploitation based on the rate of change of the swarm's diversity. The diversity rate-of-change (DRoC) measure is shown to be robust to various parameters, and consequently, to be an appropriate measure for quantifying PSO behaviour.

Using this DRoC measure to quantify PSO search behaviour, this dissertation then investigates correlations between individual FLCs and the search behaviour of PSOs. Some FLCs, such as searchability, deception, and funnels are found to correlate with PSO search behaviour in an intuitive fashion: when the FLCs indicate that a landscape is easier to solve for a PSO, the PSO tends to converge at a faster rate. The micro-ruggedness FLC correlates with PSO search behaviour counterintuitively: although landscapes with more ruggedness at the micro level could be considered harder to solve, they correspond with faster convergence in PSOs, rather than slower. Other FLCs, such as neutrality, macro-ruggedness and gradients, do not correlate significantly with PSO search behaviours.

Keywords: Optimisation, fitness landscapes, particle swarm optimisers, search behaviour, exploration, exploitation.

Supervisors : Prof. A. P. Engelbrecht
Dr. K. M. Malan
Department : Department of Computer Science
Degree : Master of Science

"I create phenomena in swarms, and paint with a full palette a gigantic and gaudy curtain before the abyss."

Nikos Kazantzakis

"Ain't no mountain high enough
Ain't no valley low enough"

Marvin Gaye, Tammi Terrell

Acknowledgements

I would like to acknowledge the following people, all of whom have been absolutely instrumental in the completion of this dissertation:

- My supervisors, Professor Andries Engelbrecht and Doctor Katherine Malan, not just for sharing their considerable expertise, thoughtful guidance, and seemingly inexhaustible patience; but also for their generosity, availability, compassion, and enthusiasm in times when I felt directionless and untethered;

- My friends from CIRG, and my friends who were pursuing their own postgraduate degrees in other fields, and my friends who have already completed their degrees, for relating, sympathising and celebrating in the various highs and lows of the journey, and for fixing the markers along the route, and for reminding me that none of us go it alone;

- My parents, Lourens and Ronelle, for everything.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Swarm-based metaheuristics have become an important area of focus within the field of optimisation [3, 28, 50]. These algorithms have been shown to succeed at finding global optima in complex optimisation problems where traditional optimisation techniques are susceptible to getting trapped in local optima [70], and have been used successfully in many optimisation applications [9, 14, 15, 20, 49]. The advantages of swarm-based metaheuristics are often attributed to their so-called emergent behaviour, which is the relatively sophisticated behaviour exhibited by a collection of individuals, arising from local interactions among the individuals, which exceeds the capabilities of any of the individuals acting alone [4, 18].

An important aspect of the behaviour of these algorithms lies in the balance of two complementary processes, namely exploration and exploitation [3, 70]. Through exploration (also called diversification), a swarm obtains a general view of the search space, enabling the algorithm to avoid getting trapped in local minima. Through exploitation (also called intensification), a swarm focuses its attention on a promising region of the search space, enabling the algorithm to improve its candidate solutions to the optimisation problem, and ideally, find the global optimum.

Due to these advantages, swarm-based metaheuristics generalise quite well to most optimisation problems. However, the so-called no free lunch theorems [67] suggest that any specific algorithm will only outperform random search on certain problems, and perform worse than random search on other problems. The assumptions for these theorems

have been shown to be invalid for many optimisation problem classes as well as many algorithms, including some in the swarm-based metaheuristics group [2, 65, 66], meaning that some algorithms may still be able to outperform others in general. Nevertheless, in practice, no single algorithm has yet been developed to outperform all others on all problems. As such, the algorithm selection problem [54] is still relevant with regards to swarm-based metaheuristics [34]; given a particular optimisation problem, a good algorithm is to be selected *for that problem* in order to find a good solution at a good rate.

This has lead to the ongoing development of many new algorithms. Unfortunately, in the search for algorithms that outperform others on specific problems, the influence of the characteristics of optimisation problems on the behaviour of algorithms is often ignored. In fact, this influence is still not well understood. This study intends to contribute to such an understanding.

1.1 Motivation

The behaviour of swarm-based metaheuristic algorithms, and especially the link between that behaviour and the algorithms' performance on specific problems, is still poorly understood, with few examples of such studies existing in the literature. Blum and Roli [3] provide an overview of how various metaheuristics implement exploration and exploitation; Clerc [7] provides an in-depth analysis of the exploitative behaviour of PSOs; and Yang and He [73] provide an analysis of an ideal balance between exploration and exploitation, concluding that there is still a significant gap between our understanding of metaheuristics and their actual behaviour.

Despite this poor understanding of algorithm behaviour, the field of meta-heuristics has recently seen a plethora of new nature-inspired swarm algorithms. For most of these new inventions, an understanding of the importance of a good balance between exploration and exploitation is stated, but the motivation for the new algorithm is not substantiated in terms of an understanding of what a good balance should be, nor by the shortcomings of existing algorithms. For example, new algorithms inspired by the behaviour of whales [43], bats [71], bird swarms [41], glowworms [31], and krill herds [16]

are all introduced without reference to any shortcomings in the behaviour of existing algorithms that the new algorithms aim to address. Other new algorithms, such as cuckoo search [72], as well as improvements suggested to existing algorithms [1, 6, 25, 39, 75], claim to achieve improved exploration or exploitation, without specifying the particular problems or circumstances under which the existing algorithms fail to explore or exploit sufficiently.

It is common within the field of optimisation to state that a particular algorithm, with a particular parameter configuration, has a particular behaviour, without accounting for the various optimisation problems that the algorithm might be applied to. While it is well known that an algorithm's parameters must be tuned to problem instances for optimal performance, this qualification is rarely applied when discussing the behaviour of algorithms.

1.2 Objectives

The aim of this study is to contribute to the undestanding of the behaviour of swarm-based metaheuristics, specifically with regards to the characteristics of optimisation problems. The two main contributions of this study are:

1. to propose a technique for numerically characterising swarm-based algorithm behaviour, and

2. to show how the characteristics of a problem can influence the behaviour of swarm-based metaheuristics.

The first contribution is a measure called the diversity rate-of-change (DRoC) of an algorithm, which indicates the rate at which an algorithm transitions from explorative to exploitative behaviour. The measurement allows for comparing the behaviour of different algorithm instances, or for comparing the behaviour of a single algorithm under different conditions. It is used in this study to investigate the influence of problem characteristics on the behaviour of algorithms. The measurement also has the potential to explain aspects of algorithm performance, and to improve parameter tuning.

The second contribution is given as a set of correlations between measurable characteristics of optimisation problems and algorithm behaviour. From these correlations, conclusions are drawn about the possible influence that certain problem characteristics have on the behaviour of algorithms.

In terms of swarm-based metaheuristics, this study is limited to particle swarm optimisers (PSOs) for the following reasons: these algorithms are simple to implement, requiring just a few parameters to be configured; many variants of the algorithms have been proposed that cover a wide range of behaviours, providing us with a sample of contrasting behaviours to study; and they are intuitive to understand, which facilitates discussions about expectations regarding their behaviours.

In terms of problem characteristics, this study focuses on the fitness landscape characteristics of optimisation problems. These characteristics have quantifiable and intuitive metrics that can be used to compare different optimisation problems. This allows us to study the correlations between the (numerically quantified) characteristics of the problems and the (numerically quantified) behaviours of algorithms, and to discuss those correlations.

Despite these limitations, the results should be applicable to other swarm-based metaheuristics and techniques for analysing optimisation problems. Drawing further links between problem characteristics and algorithms beyond the scope of this study is a promising area for future research.

1.3 Outline

Chapter 2 discusses the background concepts used in the study, namely PSO algorithms and some variants thereof, the notion of diversity in a PSO swarm, fitness landscapes, and metrics for quantifying specific FLCs. Chapter 3 presents a novel metric for quantifying PSO search behaviour with regards to exploration and exploitation, and shows the metric to be sufficiently robust with regards to some parameters that may be expected to negatively impact the accuracy of its measurements. Chapter 4 presents a modification to one of the FLC metrics introduced in Chapter 2; the chapter then addresses the primary research question by investigating correlations between FLC metrics and PSO

search behaviour. Conclusions are given in Chapter 5.

Chapter 2

Background

This chapter discusses the concepts used in this study. Section 2.1 discusses the particle swarm optimisation algorithm, and describes the variants of the algorithm that are applied to optimisation problems later in this study. Section 2.2 discusses diversity in swarms, as well as a method for quantifying diversity. Section 2.3 discusses the concept of fitness landscapes. Section 2.4 lists some characteristics that fitness landscapes can exhibit, as well as metrics that can be used to quantify those characteristics.

2.1 Particle Swarm Optimisers

Particle swarm optimisers (PSOs) are a class of stochastic population-based search algorithms inspired by the movement of flocks of birds in nature [11, 28]. PSOs work by randomly initialising a swarm of particles within the search space, with each particle representing a candidate solution to the optimisation problem. Each iteration of a PSO algorithm then updates the position of each particle to simulate movement through the search space. The direction and magnitude of each position update is determined based on stochastic elements, the particle's own momentum, and the particle's attraction towards positions in the search space that are known to be good from past personal experience, as well as from the experience of other particles in the neighbourhood of the particle.

Many variants of PSO have been proposed [17, 59], each exhibiting different behaviour

with regards to search. This study focuses on the classic PSOs with star, ring, and Von Neumann neighbourhood topologies [29] (global best, local best and Von Neumann PSO, respectively), the guaranteed convergence PSOs with the same neighbourhood topologies [62], the barebones and modified barebones PSO [27], and the social-only PSO [26]; these variants were selected because they exhibit a variety of search behaviours. The PSO variants are discussed below.

2.1.1 Global Best PSO

The global best (gbest) PSO is the basic version of the algorithm introduced by Kennedy and Eberhart [11, 28]. A swarm is initialised with n_s particles. Each particle's position at iteration $t = 0$, $\mathbf{x}_i(0)$, is initialised to a random position in the search space. The position of each particle at each subsequent iteration $t \geq 0$ is determined by adding a velocity vector to the particle's position at the previous iteration, given by

$$\mathbf{x}_i(t+1) = \mathbf{x}_i(t) + \mathbf{v}_i(t+1)\,, \tag{2.1}$$

where $\mathbf{v}_i(0) = \mathbf{0}$, and $\mathbf{v}_i(t)$ is the velocity vector of particle i at iteration t. The velocity vector is given by

$$\mathbf{v}_i(t+1) = \mathbf{v}_i^{inertia}(t) + \mathbf{v}_i^{cog}(t) + \mathbf{v}_i^{soc}(t)\,, \tag{2.2}$$

where $\mathbf{v}_i^{inertia}(t)$, $\mathbf{v}_i^{cog}(t)$, and $\mathbf{v}_i^{soc}(t)$ are the inertia, cognitive and social components of the velocity update equation, respectively.

The inertia component serves to smooth out the motion of particles [55], and is given by

$$\mathbf{v}_i^{inertia}(t) = w \cdot \mathbf{v}_i(t)\,. \tag{2.3}$$

The inertia component is the particle's velocity at the previous iteration, scaled by the inertia weight constant, $w \in [0, 1]$. For sufficiently large w values, the inertia component may serve to facilitate exploration.

The cognitive component serves to attract the particle towards the best position it has personally encountered so far, and is given by

$$\mathbf{v}_i^{cog}(t) = c_1 \cdot \mathbf{r}_1(t) \odot \big(\mathbf{y}_i(t) - \mathbf{x}_i(t)\big)\,, \tag{2.4}$$

where c_1 is a constant that controls the contribution of the cognitive component to the particle's velocity, $\mathbf{r}_1(t)$ is a random vector sampled from $U(0,1)^D$ (D being the dimensionality of the search space), $\mathbf{y}_i(t)$ is the particle's personal best position or *pbest*, which is the best position encountered by particle i by iteration t, and $\odot$ denotes element-wise vector multiplication.

Similarly, the social component serves to attract the particle towards the best position encountered by the entire swarm so far, and is given by

$$\mathbf{v}_i^{soc}(t) = c_2 \cdot \mathbf{r}_2(t) \odot \left(\hat{\mathbf{y}}(t) - \mathbf{x}_i(t)\right), \tag{2.5}$$

where c_2 is a constant controlling the contribution of the social component, $\mathbf{r}_2(t)$ is a random vector vector sampled from $U(0,1)^D$, and $\hat{\mathbf{y}}(t)$, the global best position or *gbest*, is the best position encountered by any particle in the swarm by iteration t, or equivalently, the fittest *pbest* over all particles in the swarm.

2.1.2 Local Best PSO

The local best (lbest) PSO is similar to the gbest PSO. The difference is that, for the lbest PSO, the social component of Equation (2.5) is defined as

$$\mathbf{v}_i^{soc}(t) = c_2 \cdot \mathbf{r}_2(t) \odot \left(\hat{\mathbf{y}}_i(t) - \mathbf{x}_i(t)\right), \tag{2.6}$$

where $\hat{\mathbf{y}}_i(t)$ is the neighbourhood best position or *nbest* of particle i, which is the fittest *pbest*, not of the entire swarm as with gbest PSO, but of a subset of the swarm that forms the neighbourhood of particle i. For the lbest PSO, the neighbourhood of each particle i consists of two other particles: particle $i-1$ and particle $i+1$. The neighbourhood topology of the lbest PSO is called a ring topology.

Note that the gbest PSO can be reformulated such that its social component is also given by Equation (2.6). The neighbourhood of each particle in the gbest PSO is then simply the entire swarm; such a neighbourhood topology is called a star topology.

2.1.3 Von Neumann PSO

Similar to the lbest PSO, the social component of the Von Neumann (VN) PSO [29] attracts each particle to the fittest *pbest* from the particle's neighbourhood, given by

Equation (2.6). With the VNPSO, the neighbourhood of each particle is determined as follows: The particles are arranged, in index order, into the smallest square 2-dimensional grid that can contain all the particles. For a swarm with n_s particles, the side length of such a grid is given by $\lceil \sqrt{n_s} \rceil$. Empty rows are deleted from the grid. The neighbourhood of any particle consists of the particles that are positioned directly above, below, to the left, and to the right of the particle in the grid. For example, a swarm with 10 particles is arranged into a grid as follows:

$$\begin{array}{cccc} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 10 & & \end{array}$$

where each cell indicates the index of a particle. The smallest square grid that can contain 10 particles is a 4×4 grid (a slightly smaller 3×3 grid can only contain 9 particles). The last row of the grid is empty, so the grid only contains 3 rows. The neighbours of particle 6 are particles 2 (above), 5 (to the left), 7 (to the right), and 10 (below). For particles on the edges of the grid, neighbours are selected by wrapping around; thus, in this example, the neighbours of particle 10 are particles 2 (below, wrapped around), 6 (above), and 9 (to the left, and also to the right, wrapped around).

2.1.4 Guaranteed Convergence PSO

Van den Bergh and Engelbrecht [62] have shown that the original PSO does not have guaranteed convergence to a local optimum. In the original gbest PSO as described in Section 2.1.1, a particle's cognitive component will have a value of 0 if the particle's position coincides with its *pbest* position. If the particle's position also coincides with the swarm's *gbest* position, its social component will also be 0, and its position update will only depend on its inertia; such a particle is susceptible to stagnate at a non-optimal position. Since the position of stagnation is the globally best position to which other particles in the swarm are attracted, the entire swarm is susceptible to stagnation at this position.

To overcome this problem, Van den Bergh and Engelbrecht [60] proposed the guaranteed convergence (GC) PSO. At each iteration, the index of the particle with the fittest

pbest, τ, is determined such that $\mathbf{y}_\tau = \hat{\mathbf{y}}$. The position of particle τ, $\mathbf{x}_\tau$, is updated according to an alternative velocity equation, given by

$$\mathbf{v}_\tau(t+1) = w \cdot \mathbf{v}_\tau(t) + \big(\hat{\mathbf{y}}(t) - \mathbf{x}_\tau(t)\big) + \rho(t)\big(1 - 2\mathbf{r}_2(t)\big)\,, \tag{2.7}$$

where the ρ term causes particle τ to perform a local search around its *pbest* within a bounding box, and the value of $\rho(t)$ scales the size of that bounding box. Other particles are updated according to the basic velocity equation for gbest PSO, given in Equation (2.2).

The value of ρ is determined as follows: An initial value of $\rho(0) = 1$ is used. Thereafter, at each iteration, ρ is updated according to

$$\rho(t+1) = \begin{cases} 2\rho(t) & \text{if \#successes} > s_c \\ \frac{1}{2}\rho(t) & \text{if \#failures} > f_c \\ \rho(t) & \text{otherwise,} \end{cases} \tag{2.8}$$

where #successes denotes the number of consecutive successes (position updates that resulted in a fitter *pbest*), #failures denotes the number of consecutive failures (position updates that did not result in a fitter *pbest*), and the constants s_c and f_c are threshold parameters for successes and failures, respectively. Whenever #successes is increased, #failures is reset to zero, and vice versa. To differentiate the algorithm from guaranteed convergence PSO variants for other network topologies, this algorithm is referred to as gbest guaranteed convergence PSO, or gbest GCPSO, in this study.

For the gbest PSO, each particle's neighbourhood consists of the entire swarm. Therefore, the entire swarm only has a single neighbourhood, and a single *gbest* particle. Conversely, for the lbest PSO and the VNPSO, each particle has a distinct neighbourhood, each with its own *nbest* particle, and each particle can be the *nbest* particle of multiple neighbourhoods.

Peer et al. [51] expanded GCPSO for algorithms with non-star neighbourhood topologies as follows. At each iteration, each particle is classified as an *nbest* particle if its *pbest* is the *nbest* of any of the neighbourhoods it belongs to; otherwise, it is classified as a non-*nbest* particle. Each *nbest* particle's position is then updated according to the alternative velocity update equation given by

$$\mathbf{v}_i(t+1) = w \cdot \mathbf{v}_i(t) + \big(\hat{\mathbf{y}}_i(t) - \mathbf{x}_i(t)\big) + \rho_i(t)\big(1 - 2\mathbf{r}_2(t)\big) \tag{2.9}$$

where $\rho_i(t)$ scales the bounding box of the local search performed by particle i. Non-*nbest* particles are updated normally according to Equation (2.6).

The value of ρ for each particle is determined as follows. An initial value of $\rho_i(0) = 1$ is used. Thereafter, at each iteration, ρ_i is updated according to

$$\rho_i(t+1) = \begin{cases} 2\rho_i(t) & \text{if } \mathbf{x}_i \text{ is an } nbest \textbf{ and } \#successes_i > s_c \\ \frac{1}{2}\rho_i(t) & \text{if } \mathbf{x}_i \text{ is an } nbest \textbf{ and } \#failures_i > f_c \\ \rho(t) & \text{otherwise}, \end{cases} \tag{2.10}$$

where $\#successes_i$ and $\#failures_i$ are the number of consecutive successes (position updates resulting in improved fitnesses) and failures (position updates not resulting in improved fitnesses) achieved by particle i while it is an *nbest* particle. $\#successes_i$ is reset to 0 whenever $\#failures_i$ is incremented, and vice versa. $\#successes_i$ and $\#failures_i$ are also reset to 0 whenever particle i becomes a non-*nbest* particle, and are not updated while the particle remains a non-*nbest* particle. Note that, if $\mathbf{x}_i$ is not an *nbest* particle, ρ_i remains unchanged. This ensures that, whenever the particle becomes an *nbest*, it has a previous ρ value to update according to Equation (2.10).

2.1.5 Barebones PSO

Kennedy [27] observed that, in a particle swarm with the personal and neighbourhood best positions held constant and $c_1 = c_2$, multiple position updates yield a distribution of positions with a mean halfway between the personal and neighbourhood best positions. Moreover, particles in a PSO have been shown to converge on a theoretical attraction point, $\frac{1}{2}(\mathbf{y}_i + \hat{\mathbf{y}}_i)$, which is the position halfway between their personal and neighbourhood best positions, assuming $c_1 = c_2$ [57, 61].

Kennedy [27] proposed the barebones (BB) PSO. The BBPSO does not use velocities to update particle positions; instead, at each iteration, each particle's new position is sampled directly from a Gaussian distribution around the particle's theoretical attraction point, where the deviation of the distribution is based on the absolute value of the distance between the particle's *pbest* and *nbest* positions. For each particle i in a search space with dimensionality D, each component $j \in [1, D]$ of the particle's position vector

is updated according to

$$x_{ij}(t+1) \sim N\left(\frac{1}{2}\big(y_{ij}(t)+\hat{y}_{ij}(t)\big), \sigma_{ij}(t)\right), \tag{2.11}$$

where the deviation of the sample, $\sigma_{ij}(t)$, is given by $|y_{ij}(t)-\hat{y}_{ij}(t)|$.

2.1.6 Modified Barebones PSO

Kennedy [27] also proposed a modified version of the barebones PSO described in Section 2.1.5. In the modified barebones (MBB) PSO, the updated positions of particles are recombined with their personal best positions. Position component updates are given by

$$x_{ij}(t+1) = \begin{cases} y_{ij}(t) & \text{if } U(0,1) < 0.5 \\ N(\frac{1}{2}\big(y_{ij}(t)+\hat{y}_{ij}(t)\big), \sigma_{ij}(t)) & \text{otherwise.} \end{cases} \tag{2.12}$$

Particles in the MBBPSO are more cognitively-oriented than particles in the BBPSO, resulting in more explorative search behaviour.

2.1.7 Social-Only PSO

The social-only PSO (SPSO) [26] is a variant of the gbest PSO where the cognitive component has been removed. The velocity update equation for each particle only relies on its inertia component and social component, and is given by

$$\mathbf{v}_i(t+1) = \mathbf{v}_i^{inertia}(t) + \mathbf{v}_i^{soc}(t)\,, \tag{2.13}$$

where $\mathbf{v}_i^{inertia}(t)$ and $\mathbf{v}_i^{soc}(t)$ are the inertia and social components as defined in Equations (2.3) and (2.5). Particles in the SPSO are only attracted towards the swarm's *gbest* position. The SPSO exhibits relatively fast convergence and more exploitative behaviour compared to PSO algorithms where particles are also attracted towards their *pbest* positions.

2.1.8 Cognitive-Only PSO

The cognitive-only PSO (CPSO) is a variant of the gbest PSO where the social component has been removed. The velocity update equation for each particle only relies on its

inertia and cognitive component, and is given by

$$\mathbf{v}_i(t+1) = \mathbf{v}_i^{inertia}(t) + \mathbf{v}_i^{cog}(t)\,, \tag{2.14}$$

where $\mathbf{v}_i^{inertia}(t)$ and $\mathbf{v}_i^{cog}(t)$ are the inertia and cognitive components as defined in Equations (2.3) and (2.4).

Initially, particles in the CPSO have no velocity, and each particle's *pbest* position is its initial position. The particles will thus remain at their initial positions, unless additional methods are employed to let the particles start moving. One such method is to artificially generate a nearby *pbest* position for each particle.

Particles in the CPSO do not transfer any knowledge about the search space among one another. A CPSO swarm is effectively a collection of individual hill climbers. The CPSO is therefore not expected to exhibit convergent behaviour.

2.2 PSO Swarm Diversity

The diversity of a swarm refers to the dispersion of its particles' positions across the search space. In a more diverse swarm, particles are, on average, further apart. Olorunda and Engelbrecht [47] have found the average distance around the swarm centre to be a suitable measure of diversity. The measure is given by

$$\mathcal{D} = \frac{1}{n_s}\sum_{i=1}^{n_s}\sqrt{\sum_{j=1}^{D}(x_{ij} - \bar{x}_j)^2}\,, \tag{2.15}$$

where n_s is the swarm size, D is the dimensionality of the search space, x_{ij} is the j-th dimensional component of the i-th particle's position, and $\bar{x}_j$ is the average of the j-th dimensional component of all particles' positions, given by

$$\bar{x}_j = \frac{1}{n_s}\sum_{i=1}^{n_s} x_{ij}\,. \tag{2.16}$$

The diversity of a swarm is indicative of the swarm's search behaviour at a particular instant. A highly diverse swarm's particles are dispersed over a wide range of positions; such a swarm is exploring the search space. Conversely, for a swarm with low diversity, the particles are concentrated around a small region of the search space; such a swarm is exploiting a region of the search space.

2.3 Fitness Landscapes

Optimisation is the task of selecting the best element from a set, based on some criteria. An optimisation problem can be defined by a fitness function, $f : \mathcal{S} \to \mathbb{R}$, which maps candidate solutions from a search space to real values representing the fitness[1] of chosen solutions. For continuous optimisation problems, the search space $\mathcal{S} \subset \mathbb{R}^n$ is a subset of the Euclidean space, typically with a restricted domain. In order to solve the optimisation problem, the aim is to find a solution $\mathbf{x}^* \in \mathcal{S}$ such that $f(\mathbf{x}^*) \leq f(\mathbf{x}) \ \forall \mathbf{x} \in \mathcal{S}$ (assuming a minimisation problem; a maximisation problem can be transformed by simply inverting the sign of the result.) Such a solution is called the global minimum (assuming a minimisation problem; for a maximisation problem, it is called the global maximum).

In addition to a globally optimal solution, optimisation problems may have locally optimal solutions, called local optima. A local optimum is a solution $\mathbf{x}^*_{\mathcal{N}} \in S$ such that $f(\mathbf{x}^*_{\mathcal{N}}) \leq f(\mathbf{x}) \, \forall \, \mathbf{x} \in \mathcal{N}$, where $\mathcal{N} \subset \mathcal{S}$. A search technique is a technique or algorithm that is used to find an optimum of an optimisation problem. Some search techniques are simply used to find any optimum for a problem; however, the ideal is for search techniques to find the global optimum. Optimisation problems with local optima can mislead some search techniques to believe that their local optima are global optima. Therefore, local optima may impede the search for global optima. Optimisation problems with only a single, global optimum are called unimodal. Optimisation problems with additional, non-global (i.e. local) optima are called multimodal.

For fitness functions defined in one or two dimensions, one can visualise a *fitness landscape* (a concept attributable to Wright [68]): the vectors of candidate solutions correspond to coordinates on the terrain, and the fitnesses of those solutions are represented by the elevation of the terrain at those positions. For maximisation problems, the optima are situated at the peaks of mountains in the landscape. Similarly, for minimisation problems, the optima are situated at the depths of valleys in the landscape.

When regarding an optimisation problem as a fitness landscape, interesting tech-

[1] The term fitness here refers to the quality of a solution in the general sense, and is not limited to the context of evolution.

niques for analysing the problem become available. For example, consider a hill climbing optimisation technique. A hill climbing algorithm randomly selects an initial solution from the search space, and then iteratively replaces it with a similar solution that is fitter, until no similar solution can be found that is fitter. Hill climbing is considered to be a local optimisation technique, and as such, is susceptible to fail on multimodal optimisation problems. A multimodal maximisation problem can be visualised as a landscape with a global maximum at the peak of a single tall mountain, and local maxima at the peaks of smaller hills. It then becomes clear why the hill climbing technique is likely to fail: The hill climber simply rises to the peak of whichever hill is closest to its initial starting point; that hill is unlikely to be the tallest mountain in the landscape. In this example, the notion of a neighbourhood is implied: for a particular initial solution, the hill climber will only select a fitter solution from its immediate neighbourhood, which is the set of solutions that are nearby enough. In addition to the set of candidate solutions and the fitness function, the definition of the neighbourhood of a solution is critical to the shape of the fitness landscape. For example, suppose that a hill climber solving a fitness function defined in two dimensions is limited so that it can only move to new solutions in a single direction; despite the dimensionality of the fitness function, the hill climber perceives the landscape as a one-dimensional function.

2.4 Fitness Landscape Characteristics

Similar to landscapes in nature, fitness landscapes can exhibit various characteristics. For example, a landscape can have cliffs, or be steep, or be rough. Analysis of the characteristics of a fitness landscape can further our understanding of what makes it easier or harder for search techniques to solve optimisation problems. This study focuses on the ruggedness, neutrality, gradients, global landscape structure (underlying modality), deception, and searchability of fitness landscapes. These landscape characteristics, along with a scalar-valued metric for quantifying the presence and severity of each characteristic, are discussed below.

2.4.1 Ruggedness

Compare the two landscapes in Figure 2.1. They have similar underlying structures, but the landscape on the left is smooth, while the landscape on the right is rugged. A landscape is rugged if it has numerous "ups and downs", or sudden changes in fitness. Ruggedness thus refers to the level of variation in a fitness landscape. Rugged landscapes have many local optima [32, 48], which make them difficult to solve for certain search techniques.

Vassilev et al. [63] proposed a first entropic measure (FEM) of ruggedness with regards to neutrality in discrete search spaces. Malan and Engelbrecht [35] adapted the metric to quantify macro- and micro-ruggedness in continuous search spaces. The measurements are obtained as follows. The search space is first sampled using a progressive random walk [38] of n_t steps. The first entropic measure of macro-ruggedness, $\text{FEM}_{0.1}$, uses a step size of 10% of the range of the function domain during the walk. Similarly, the first entropic measure of micro-ruggedness, $\text{FEM}_{0.01}$, uses a step size of 1% of the domain range. The random walk produces a time series of fitness values, $f_1, ..., f_{n_t}$. From overlapping pairs of fitnesses in this time series, a string of symbols, $S(\varepsilon) = s_1 ... s_{n_t-1}$, $s_i \in \{\bar{1}, 0, 1\}$, is generated according to

$$s_i = \begin{cases} \bar{1} & \text{if } f_i - f_{i-1} < -\varepsilon \\ 0 & \text{if } |f_i - f_{i-1}| \leq \varepsilon \\ 1 & \text{otherwise}\,, \end{cases} \tag{2.17}$$

where ε controls the allowable error between fitnesses that will be considered equal. A group of three points from the walk, represented by two adjacent symbols, s_i and s_{i+1}, is considered a rugged group if $s_i \neq s_{i+1}$. From the set of all overlapping three-point groups sampled from the walk, the entropy of the sub-set of rugged groups is estimated using the information function,

$$H(\varepsilon) = \sum_{p \neq q} P_{[pq]} log_6 P_{[pq]}\,, \tag{2.18}$$

where pq is a substring of two elements in $S(\varepsilon)$, and $P_{[pq]}$ is the proportion of occurrences of pq in $S(\varepsilon)$, given by

$$P_{[pq]} = \frac{n_{[pq]}}{n_t - 1}\,, \tag{2.19}$$

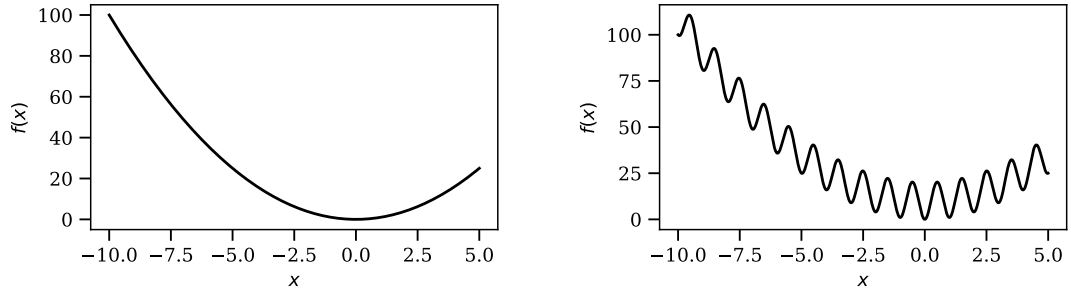

Figure 2.1: Example of a smooth (left) and a rugged (right) landscape

where $n_{[pq]}$ is the number of occurrences of the substring pq in $S(\varepsilon)$.

For larger values of ε, more fitnesses are considered equal, and the landscape becomes more neutral. The smallest value of ε for which the landscape becomes completely neutral is denoted as ε^* and is called the information stability. To obtain a scalar-valued metric of ruggedness, $H(\varepsilon)$ is calculated for various $\varepsilon \in [0, \varepsilon^*]$; the largest of these results is taken as the metric of ruggedness. The result is a value in $[0, 1]$, where 0 indicates a flat landscape, and 1 indicates maximal ruggedness.

2.4.2 Neutrality

Figure 2.2 shows two examples of landscapes with neutrality. The landscape on the left has a large connected neutral area on the left, and a smaller connected neutral area on the right. The landscape on the right has many small neutral areas. A fitness landscape exhibits neutrality if there are sections of the landscape where neighbouring points have nearly equal fitnesses. At such sections, the slope of the landscape is approximately zero. (Neutrality is distinct from smoothness in that a smooth landscape is not rugged, but may still have a non-zero slope to lead search algorithms towards fitter regions of the landscape.) Neutrality poses a challenge for many search algorithms, because zero-slope areas of a search space offer no information to guide search in a particular direction. Some search algorithms may also falsely interpret the lack of variation in fitness as an indication that the algorithm is converging on an optimum, when it might in fact just be moving through a neutral area.

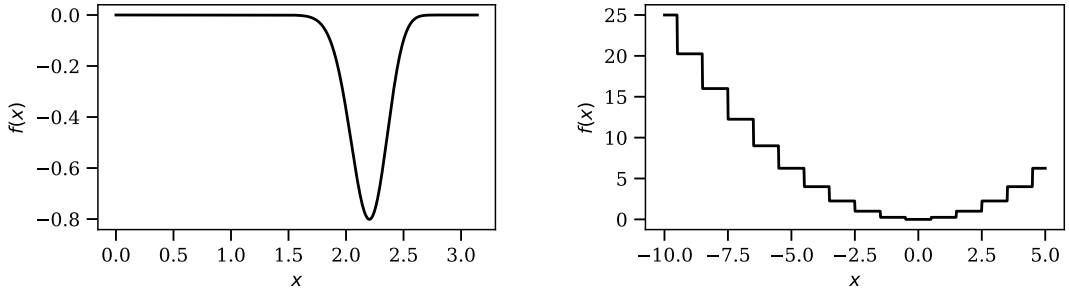

Figure 2.2: Examples of landscapes with neutrality

Van Aardt et al. [58] proposed two neutrality metrics, which are obtained as follows. The search space is sampled using progressive random walks of n_t steps, giving a time series of fitness values, $f_1, ..., f_{n_t}$. A string of symbols, $S(\varepsilon) = s_1, ..., s_{n_t-2}$, $s_i \in [0, \bar{0}]$, is generated from the fitness time series according to

$$s_i = \begin{cases} 0 & \text{if } \max(f_{i-1}, f_i, f_{i+1}) - \min(f_{i-1}, f_i, f_{i+1}) < \varepsilon \\ \bar{0} & \text{otherwise,} \end{cases} \tag{2.20}$$

where ε is the neutrality threshold, with higher values for ε resulting in more structures being considered neutral. Van Aardt et al. [58] used a neutrality threshold of 1×10^{-8}. The symbol 0 indicates a neutral 3-point structure in the walk, and $\bar{0}$ indicates a non-neutral structure.

From this string of symbols, two metrics are calculated. The first is the proportion of neutral structures in a walk, given by

$$\text{PN}(\varepsilon) = \frac{n_n}{|W|}\,, \tag{2.21}$$

where n_n is the number of neutral structures s in $S(\varepsilon)$ such that $s = 0$, and $|W|$ is the overall number of structures in $S(\varepsilon)$. The result is a value in $[0, 1]$, where 0 indicates that there is no neutrality in the landscape, and 1 indicates that the landscape is completely neutral. The metric gives an indication of overall neutrality, but does not provide an indication of the sizes of individual neutral regions. It would result in a similarly high value for a landscape with a single, large neutral plain as for a landscape with many small, disconnected neutral plains.

The second metric is the longest subsequence of neutral structures in a walk proportionate to the overall number of structures in the walk, given by

$$\text{LSN}(\varepsilon) = \frac{\max(|w_n|)}{|W|}\,, \tag{2.22}$$

where $\max(|w_n|)$ is the maximum number of steps in any neutral-only subsequence w_n of $S(\varepsilon)$ such that $s = 0 \;\forall\, s \in w_n$. The result is a value in $[0, 1]$, where 0 indicates a landscape with no neutrality, and 1 indicates a completely neutral landscape. The metric is based on the connectedness of neutral structures, and gives an indication of the size of neutral plains in the landscape.

2.4.3 Gradients

Compare the landscapes in Figure 2.3. The landscapes have similar shapes, but the basins of the local optima of the landscape on the right are steeper than those of the landscape on the left. An important difference between the two landscapes is that the landscape on the right has steeper gradients. The steepness of the gradients in a landscape can contribute to its difficulty. Some search algorithms may specifically struggle to optimise a landscape that has local optima in very deep basins.

Malan and Engelbrecht [36] proposed two metrics for estimating gradients in a landscape. Both are based on a sample of the search space using a progressive Manhattan walk. In this type of walk, each step updates the position by a fixed step size u in a single randomly-selected dimension. The metrics are obtained as follows. A walk of n_t steps produces a sample of positions, $\mathbf{x}_1, \ldots, \mathbf{x}_{n_t}$. A sequence of the normalised gradient of each step in the walk, $g_1, \ldots, g_{n_t-1}$, is generated according to

$$g_t = \frac{(f(\mathbf{x}_{t+1}) - f(\mathbf{x}_t))/(f^{max} - f^{min})}{u/(\sum_{j=1}^{n_t} x_j^{max} - x_j^{min})}\,, \tag{2.23}$$

where u is the fixed step size used in the progressive Manhattan walk, f^{max} and f^{min} are the maximum and minimum fitness values encountered during the walk, and x_j^{max} and x_j^{min} are the search space bounds in dimension j.

Malan and Engelbrecht [36] derived two metrics from this sequence. The first metric

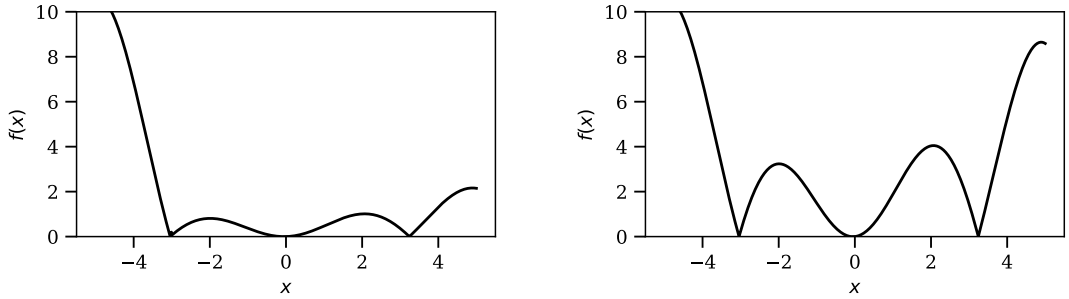

Figure 2.3: Example landscapes with different gradients

is the average gradient, given by

$$\mathrm{G}_{avg} = \frac{\sum_{t=1}^{n_t-1} |g_t|}{n_t - 1}\,. \tag{2.24}$$

Note that the absolute value of each gradient is used. This is so that the metric is not concerned with the direction of the gradient, and so that gradients across a landscape with opposite signs do not cancel one another out. The result is a positive real number, where higher values indicate the presence of steeper slopes. The metric does not, however, give any indication as to the distribution of the slopes across the search space.

The second metric is the standard deviation of the gradients, given by

$$\mathrm{G}_{dev} = \sqrt{\frac{\sum_{t=1}^{n_t-1} (\mathrm{G}_{avg} - |g_t|)^2}{n_t - 2}}\,. \tag{2.25}$$

The result is a positive real number. Lower values indicate that gradients are more evenly distributed across the landscape, and therefore that G_{dev} is a good indicator of the general gradient of the landscape. On the other hand, higher values indicate highly contrasting gradients at different areas of the landscape; at some areas, there may be gradients that are vertical or very steep, and at other areas, the landscape may be close to neutral.

2.4.4 Global Landscape Structure

Compare the landscapes in Figure 2.4. Both landscapes are similarly multimodal and rugged, but the landscape on the left has an underlying unimodal structure, whereas the

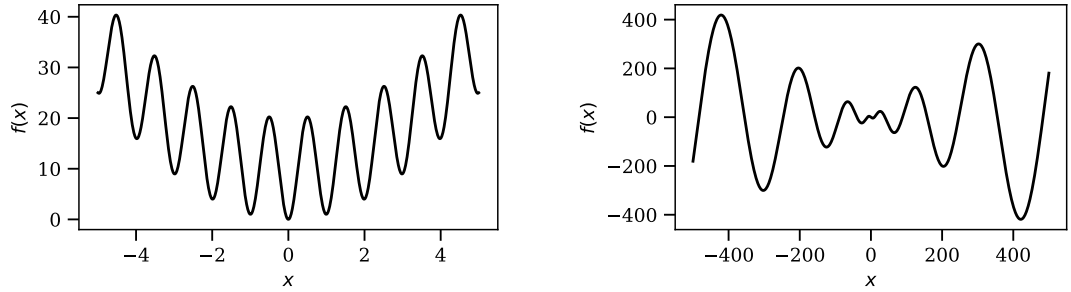

Figure 2.4: Example landscapes with underlying unimodal (left) and multimodal (right) global structures

landscape on the right has an underlying multimodal structure. An underlying unimodal structure can make a landscape easier to optimise for certain algorithms. Likewise, an underlying multimodal structure can make a landscape harder to optimise.

Lunacek and Whitley [33] introduced a dispersion metric for estimating global landscape structure. Malan and Engelbrecht [36] adapted the metric by normalising solution vectors to allow comparison among landscapes with different domains. The adapted metric is obtained as follows. First, n_a positions are sampled from a uniform random distribution of the search space. The sampled positions are normalised so that the domain of the search space is [0, 1] in all dimensions. From such a sample, the subset $S^*_{n_b}$ of the n_b fittest positions is determined, where $n_b < n_a$. Then, $\text{disp}(S^*_{n_b})$ is calculated as the average pair-wise Euclidean distance between the positions in $S^*_{n_b}$. The metric is then given by

$$\text{DM} = \text{disp}(S^*_{n_b}) - \text{disp}_D\,, \tag{2.26}$$

where D is the dimensionality of the search space, and disp_D is the average pair-wise Euclidean distance of a large sample taken from $U(0,1)^D$. The result is a value in $[-\text{disp}_D, \sqrt{D} - \text{disp}_D]$, where negative values indicate a landscape with an underlying unimodal structure, and positive values indicate a landscape with an underlying multimodal structure.

2.4.5 Deception

Gradients in a landscape that lead towards optima act as information to guide search towards those optima. A landscape's difficulty is influenced by the availability of such information, as well as the quality of the available information. Information is considered to be of high quality if it guides the search towards the global optimum. Information that leads away from the global optimum is considered deceptive. Compare the two landscapes in Figure 2.5. In the left landscape, the gradients lead towards the global optimum, while in the right landscape, the gradients lead away from global optimum. A landscape with readily available, high quality information is easier to solve than one where such information is scarce. The difficulty is further increased if a landscape contains deceptive information.

Jones and Forrest [24] proposed a fitness distance correlation metric for estimating deception in a landscape for a hill climbing algorithm. The metric indicates how well the fitness of points in the landscape correlate to their distance from the global optimum. In a landscape with a perfect correlation between the fitness of positions and their distance to the global optimum, moving closer to the global optimum always results in a better fitness; such a landscape perfectly guides search towards the global optimum. If moving closer to the global optimum results in worse fitness, a search technique could be deceived to direct its search away from the global optimum. The metric requires that the global optimum of the landscape be known, making it unsuitable for some practical scenarios. Malan [37] adapted the metric to lift this requirement by substituting the fittest position from a sample for the global optimum. This shifts the focus of the metric from deception in terms of the global optimum to deception in relation to a fitter solution. The adapted metric is obtained as follows. A sample of n_a points is taken from the landscape, along with their associated fitnesses, $F = \{f_1, \dots, f_{n_a}\}$. The fittest point from the sample, $\mathbf{x}^*$, is determined. The Euclidean distance from each sampled point to $\mathbf{x}^*$ is determined, giving $\text{Dist}^* = \{d_1^*, ..., d_{n_a}^*\}$. The metric is the correlation between F and Dist^*, given by

$$\text{FDC} = \frac{\sum_{i=1}^{n_a} (f_i - \bar{f})(d_i^* - \bar{d}^*)}{\sqrt{\sum_{i=1}^{n_a} (f_i - \bar{f})^2} \sqrt{\sum_{i=1}^{n_a} (d_i^* - \bar{d}^*)^2}}, \tag{2.27}$$

where $\bar{f}$ is the mean of F, and $\bar{d}^*$ is the mean of Dist^*. The result is a value in $[-1, 1]$.

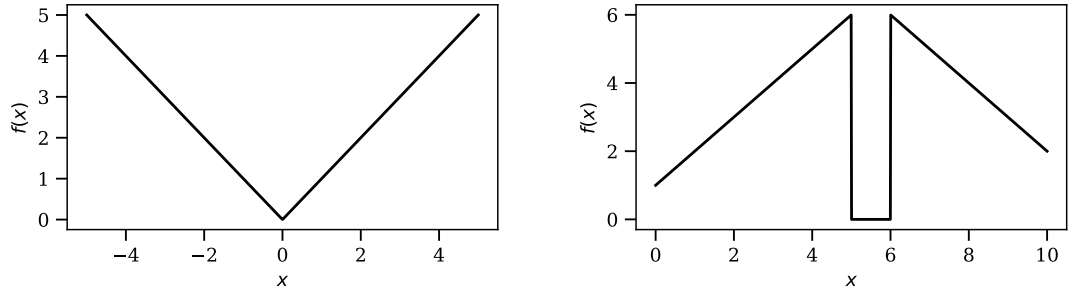

Figure 2.5: Example landscapes with gradients leading towards (left) and away from (right) global optima.

Assuming minimisation, higher values indicate landscapes with more high-quality information to guide the search towards the global optimum.

2.4.6 Searchability

In genetics, evolvability refers to the capacity to generate fit variants. Changes in the fitness landscape have been shown to influence evolvability [12, 40]. Therefore, evolvability not only characterises the individual or population undergoing evolution, but is also linked to the local structure of the search space. In optimisation, evolvability is limited to evolutionary algorithms. Malan [37] generalises evolvability by introducing the concept of searchability with regards to a particular search technique. Searchability refers to the capacity of a particular search technique to move to a fitter position in a landscape.

Verel et al. [64] introduced fitness clouds as a technique for analysing the evolvability of a landscape with regards to a particular genetic search operator. The technique was adapted by Malan and Engelbrecht [37] to analyse searchability with regards to PSO algorithms, and to derive scalar-valued results from the scatterplot generated by the original technique. The metric is calculated as follows. A particular variant of PSO is selected as the search operator (Malan and Engelbrecht selected social-only and cognitive-only PSO, giving two distinct metrics). Particles of the PSO are randomly initialised within the search space at positions $\mathbf{x}_1, ..., \mathbf{x}_{n_s}$. The particles' positions are

then updated using the chosen PSO variant, giving updated positions $\mathbf{x}'_1, ..., \mathbf{x}'_{n_s}$. For cognitive PSO, particles are initialised with artificially-generated personal best positions in order for the particles to build up momentum. The particles are initialised with zero velocity; therefore, each particle's position is updated twice in order for the inertia component of the velocity update equation to come into effect. Particles that have left the search space during the two position updates are discarded, leaving n_v valid initial positions, $\mathbf{x}_1, ..., \mathbf{x}_{n_v}$, and n_v updated positions, $\mathbf{x}'_1, ..., \mathbf{x}'_{n_v}$. The first metric, the fitness cloud index, is the proportion of valid particles for which the fitness was improved, given by

$$\text{FCI} = \frac{\sum_{i=1}^{n_v} g(i)}{n_v}\,, \tag{2.28}$$

where

$$g(i) = \begin{cases} 1 & \text{if } f(\mathbf{x}'_i) < f(\mathbf{x}_i) \\ 0 & \text{otherwise}\,. \end{cases} \tag{2.29}$$

The result is a value in $[0, 1]$, where higher values indicate more searchable landscapes. The metric has two variants corresponding to the variants of PSO that were selected for the position updates: FCI_{soc} is the fitness cloud index of a landscape obtained using updates of the social-only PSO algorithm, and FCI_{cog} is obtained using cognitive-only PSO.

Malan [34] found that, when taking multiple FCI measurements on a single landscape, the unpredictability of the metric is a good indicator of problem difficulty. The unpredictability of the two FCI metrics taken over multiple runs is determined as follows. Samples of a sufficiently large sample size n_a of measurements are taken on the landscape using FCI_{soc}, as well as FCI_{cog}, giving two samples of measurements, $\text{FCI}_{soc}(1), ..., \text{FCI}_{soc}(n_a)$ and $\text{FCI}_{cog}(1), ..., \text{FCI}_{cog}(n_a)$. (A sample size of $n_a = 30$ is typically deemed suitable.) The second metric, the fitness cloud index mean standard deviation, is then given by

$$\text{FCI}_{\bar{\sigma}} = \frac{\text{FCI}_{soc}(\sigma) + \text{FCI}_{cog}(\sigma)}{2}\,, \tag{2.30}$$

where $\text{FCI}_{soc}(\sigma)$ and $\text{FCI}_{cog}(\sigma)$ are the standard deviations of the FCI_{soc} and FCI_{cog} samples, respectively. The result is a value in approximately $[0, 0.509]$, where higher

values indicate that the FCI metrics are less predictable on the landscape, and thus that the landscape is less searchable.

2.5 Summary

This chapter provided an overview of the particle swarm optimisation (PSO) algorithms that are used in this dissertation, namely the global best (gbest) PSO, local best (lbest) PSO, Von Neumann (VN) PSO, guaranteed convergence (GC) variations with star, ring and Von Neumann topologies, barebones (BB) and modified barebones (MBB) PSOs, and social-only (S) PSO.

The chapter then discussed the notion of swarm diversity, and discussed a method to quantify a swarm's diversity at a particular instant by measuring the average distance of all particles around the swarm's center position.

Lastly, the chapter provided an overview of fitness landscapes derived from optimisation problems. Such fitness landscapes may exhibit certain characteristics which, if measured, can provide insight into the difficulties that the optimisation problems pose for the search techniques solving those problems. Existing metrics were discussed for a selection of these fitness landscape characteristics. The ruggedness of a landscape can be estimated on the micro- and macro-level using the first entropic measures for micro- and macro-ruggedness ($\text{FEM}_{0.01}$ and $\text{FEM}_{0.1}$, respectively). The proportion of a landscape's terrain that is neutral can be estimated by the proportion of neutral structures (PN) in a random walk, and the connectedness of the neutral terrain can be estimated by the length of the longest neutral subsequence (LSN) of such a random walk. The severity of gradients in a landscape can be estimated by the average and the standard deviation of the gradients between steps of a random walk (G_{avg} and G_{dev}, respectively). The underlying modality of a landscape can be estimated by the dispersion metric (DM). The deceptiveness of a landscape can be estimated by the fitness-distance correlation (FDC). A landscape's searchability with regards to particle swarm optimisers can be estimated by fitness cloud indices obtained with social and cognitive updates (FCI_{soc} and FCI_{cog}, respectively). Additionally, the average standard deviation of fitness cloud index measurements ($\text{FCI}_{\bar{\sigma}}$) can be used to estimate a landscape's searchability.

Chapter 3

Diversity Rate-of-Change

This chapter discusses a measure for quantifying the search behaviour of PSO algorithms. The measure, called the diversity rate-of-change (DRoC) measure, is based on instantaneous diversity measurements taken on a PSO swarm. Aspects of this chapter have been published in [5].

The chapter is structured as follows. Section 3.1 introduces the DRoC measure. Section 3.2 tests the DRoC measure's ability to show differences in the search behaviours of variants of the PSO algorithm. Section 3.3 tests the robustness of the DRoC measure with regards to the swarm size used with PSO algorithms. Similarly, Section 3.4 tests the DRoC measure's robustness with regards to the dimensionality of the search space, and Section 3.5 tests the measure's robustness with regards to the number of diversity measurements used to calculate the measure. Conclusions are summarised in Section 3.6.

3.1 Diversity Rate-of-Change Measure

Section 2.2 provided a method for measuring the diversity of a PSO swarm. Such a diversity measurement indicates a PSO's search behaviour in terms of exploration and exploitation; however, an individual diversity measurement only indicates an algorithm's behaviour at the particular instant at which the measurement was taken. In order to understand the search behaviour of an algorithm over its entire execution, many diversity

measurements must be taken at various iterations, and the trend over time must be studied. If an algorithm initially has high diversity, which then reduces rapidly, the algorithm can be understood to transition from explorative to exploitative behaviour rapidly. Conversely, if an algorithm reduces its diversity at a relatively slower rate, the algorithm can be understood to transition from explorative to exploitative behaviour at a slower rate. All of the PSO algorithms included in this study find a single solution in a static environment, with no processes implemented in order to manage diversity; therefore, as they transition from exploration to exploitation, all the PSO algorithms reduce their diversity in a roughly continuous manner.

In order to quantify the rate at which a PSO reduces its diversity, the diversity rate-of-change measure is proposed. Measurements are obtained as follows. A PSO is applied to an optimisation problem for n_t iterations. During the execution of the PSO, the diversity of the swarm at each iteration t is determined as

$$\mathcal{D}(t) = \frac{1}{n_s} \sum_{i=1}^{n_s} \sqrt{\sum_{j=1}^{D} (x_{ij}(t) - \bar{x}_j(t))^2}\,, \tag{3.1}$$

where $x_{ij}(t)$ is the j-th dimensional component of the i-th particle at iteration t, and $\bar{x}_j(t)$ is the average of the j-th dimensional component of all particles' positions at iteration t. All the diversity measurements taken during the execution of the PSO are approximated by the two-piecewise linear approximation,

$$y(t) \approx \mathcal{D}(t) \quad \text{for } 0 \leq t \ < n_t\,, \tag{3.2}$$

where the two-piecewise linear approximation takes the form

$$y(t) = \begin{cases} m_1 x + c & \text{for } 0 \ \leq t \ < t' \\ m_2(x + t') + m_1 t' + c & \text{for } t' \ \leq \ t \ \leq \ n_t\,, \end{cases} \tag{3.3}$$

where m_1 is the gradient of the first line segment, c is the y intersection of the first line segment, m_2 is the gradient of the second line segment, and t' is the t-value at which the two line segments cross. The values of m_1, m_2 and t' are chosen to minimise the least squares error (LSE) between $y(t)$ and $\mathcal{D}$, given by

$$\text{LSE} = \sum_{t=0}^{n_t - 1} \big(\mathcal{D}(t) - y(t)\big)^2\,. \tag{3.4}$$

In order to obtain an accurate approximation of the rate at which $\mathcal{D}(t)$ decreases, especially during the initial values of t, the initial y-value of the two-piecewise approximation is fixed to be the same as the initial y-value of the diversity measurements it approximates. This is achieved by simply setting c to $\mathcal{D}(0)$.

The rate at which a swarm decreases its diversity is indicated by the gradient of the first line of the linear approximation. If a swarm decreases its diversity at a relatively fast rate, the first line will have a relatively steeper slope. Conversely, if a swarm decreases its diversity at a relatively slow rate, the first line will have a relatively gradual slope. The diversity rate-of-change measure is therefore given by m_1, the gradient of the first line. The measurement is a negative number, where the magnitude of the measurement indicates the rate at which a swarm transitions from exploration to exploitation. A smaller DRoC measurement, or equivalently, a negative measurement with a greater magnitude, indicates that the swarm transitioned to explorative behaviour at a faster rate.

Consider the diversity measurements plotted in Figure 3.1a. During the first few iterations, the diversity measurements of the gbest PSO (red) are relatively smaller than those of the lbest PSO (blue). The gbest PSO transitions from explorative to exploitative behaviour at a faster rate compared to the lbest PSO for this specific problem. The two-piecewise linear approximations fitted to the diversity measurements of each PSO are plotted in Figure 3.1b. Additionally, the DRoC measurement for each algorithm, which is simply the slope of the first line segment of each two-piecewise approximation, is indicated in the top-right corner of Figure 3.1b. The DRoC for the gbest PSO is a smaller number than that of the lbest PSO, indicating that the gbest PSO reduced its diversity at a faster rate than the lbest PSO.

3.2 Testing the DRoC Measure

This section tests the capability of the DRoC measure to quantify the search behaviour of PSOs. The section is structured as follows. Expected differences between some PSO variants with respect to their search behaviour in terms of exploration and exploitation are given in Section 3.2.1. The experimental procedure used in this section is discussed

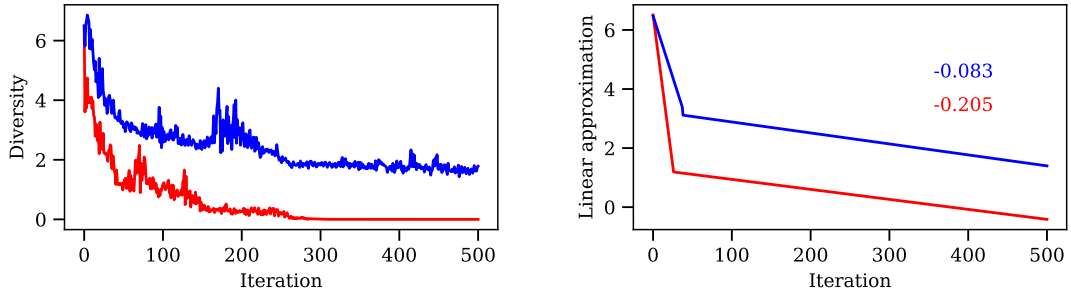

(a) Diversity measurements taken on the gbest PSO (red) and the lbest PSO (blue) over 500 iterations

(b) Two-piecewise linear approximations fitted to diversity measurements taken on the gbest PSO (red) and the lbest PSO (blue) over 500 iterations, with each algorithm's associated m_1 value indicated in the upper-right corner

Figure 3.1: Examples of DRoC measurements taken on two PSO algorithms

in Section 3.2.2. Results are discussed in Section 3.2.3. Conclusions are summarised in Section 3.2.4.

3.2.1 Expected PSO Behaviour

In order to evaluate the DRoC measure, the primary method used in this chapter is to compare DRoC measurements taken on each of a pair of PSOs, rank them, and determine whether they relate to one another in a sensible manner. For example, knowing that the SPSO has a relatively more connected neighbourhood topology than the lbest PSO does, it is sensible to expect the SPSO to reduce its diversity at a relatively faster rate compared to the lbest PSO, and therefore, to expect DRoC measurements taken on the SPSO to have a greater magnitude than those taken on the lbest PSO. If DRoC measurements taken on the SPSO and the lbest PSO contradict these expectations, that could indicate a failure of the DRoC measure to accurately indicate the respective algorithms' search behaviour. Such a method of investigation requires a set of pre-existing expectations about how different PSOs should compare to one another in terms of search behaviour. This section defines such a set of expectations.

If particles in a PSO algorithm base their position updates on their *nbest* positions, the neighbourhood topology of the algorithm is expected to influence the rate at which the swarm transitions from explorative to exploitative behaviour. Consider a swarm with neighbourhood topology $\mathcal{N}$. Suppose a particle in this swarm with index τ finds the best position, $\mathbf{x}^*$, encountered by the swarm so far. This particle becomes the neighbourhood best particle for all its neighbouring particles in $\mathcal{N}_\tau$. Other particles remain unaffected.

If $\mathcal{N}$ is a star topology, as with the gbest PSO, gbest GCPSO, and SPSO, then $\mathcal{N}_\tau$ is comprised of the entire swarm. The social component of each particle in the swarm attracts it towards the new *nbest* position immediately. Once all the particles get close to the position, the swarm can transition to an exploitative behaviour.

Conversely, if $\mathcal{N}$ is a grid topology, as with the VNPSO and VNGCPSO, then $\mathcal{N}_\tau$ is comprised of, at most, four particles. Only these particles are immediately attracted to the new *nbest*. Once they reach this position, they will adopt the *nbest* position as their *pbest* positions, and they will start attracting their neighbours. In turn, their neighbours also need to reach the *nbest* position before they start attracting *their* neighbours, and so on. The decreased connectedness of the grid topology slows down the attraction effect of the *nbest* position found by any particle, and is expected to slow down the swarm's transition from exploration to exploitation.

Similarly, if $\mathcal{N}$ is a ring topology, which is the least-connected topology included in this study, as with the lbest PSO and lbest GCPSO, then $\mathcal{N}_\tau$ is comprised of only two particles. Such a swarm is expected to transition from exploration to exploitation at the slowest rate compared to the other topologies.

Particles in the gbest PSO are relatively well-connected, but each particle's attraction to the swarm's *gbest* position is counteracted by its attraction to its own *pbest* position; the cognitive component of the gbest PSO facilitates exploration [13], and therefore slows down the swarm's transition to exploitation. In the SPSO, this cognitive component is absent. The SPSO is expected to transition to exploitation at the fastest rate compared to all other algorithms included in this study.

For the guaranteed convergence algorithms (gbest GCPSO, lbest GCPSO, and VN-GCPSO), the swarm performs a local search around each *nbest* position; this is done to prevent the swarm from stagnating prematurely by keeping the swarm in an explorative

state for longer. The guaranteed convergence algorithms are therefore expected to transition to exploitation slightly slower than their non-guaranteed convergence counterparts (gbest PSO, lbest PSO, and VNPSO), though this difference might not be significant.

The BBPSO differs from most of the algorithms in this study in that it does not update its particles' positions using a velocity update equation; instead, new positions are sampled around an attraction point. This attraction point is based on each particle's *pbest* and *nbest* positions, and a position update based on this attraction point is somewhat analogous to a velocity update with a cognitive and social component. However, particles' previous positions or velocities are not taken into account, so there is no analogue for an inertia component, which would facilitate exploration for large w values. Therefore, the BBPSO is expected to transition to exploration slightly faster than the gbest PSO and the gbest GCPSO.

The MBBPSO differs from the BBPSO in that particles' updated positions are recombined with their *pbest* positions in order to facilitate initial exploration. The MBBPSO is therefore expected to remain in an explorative state for longer than the BBPSO.

In the CPSO, particles are not expected to converge, so the algorithm's behaviour can not be described in terms of exploration and exploitation. The CPSO is therefore not included in this study.

Table 3.1 summarises the expectations regarding the search behaviour of PSOs relative to one another. The table shows, for the PSO in each row, whether that PSO is expected to reduce its diversity at a faster or slower rate than the PSO in each column.

For some comparisons, the two PSOs are expected to be only slightly different. In these cases, finding no significant difference in behaviour between the two PSOs will also be considered to be in line with expectations.

Some expectations are not defined. For instance, when comparing the gbest PSO to the MMBPSO, it is expected that the MBBPSO will exhibit slower convergence than the BBPSO, and that the BBPSO will exhibit faster convergence than the gbest PSO; however, the degrees of difference are not defined for either comparison, and so no direct expectation can be defined for the difference in behaviour between the gbest PSO and the MBBPSO. These comparisons are thus not used when evaluating the DRoC measure in this chapter.

Table 3.1: Expected comparative rate of diversity reduction between each pair of PSOs

	Gbest PSO	VNPSO	Lbest PSO
Gbest PSO	–	Faster	Faster
VNPSO	Slower	–	Faster
Lbest PSO	Slower	Slower	–
Gbest GCPSO	Slightly slower	Undefined	Undefined
VNGCPSO	Slower	Slightly slower	Undefined
Lbest GCPSO	Slower	Slower	Slightly slower
BBPSO	Slightly faster	Faster	Faster
MBBPSO	Undefined	Undefined	Undefined
SPSO	Faster	Faster	Faster

	Gbest GCPSO	VNGCPSO	Lbest GCPSO
Gbest PSO	Slightly faster	Faster	Faster
VNPSO	Undefined	Slightly faster	Faster
Lbest PSO	Undefined	Undefined	Slightly faster
Gbest GCPSO	–	Faster	Faster
VNGCPSO	Slower	–	Faster
Lbest GCPSO	Slower	Slower	–
BBPSO	Slightly faster	Faster	Faster
MBBPSO	Undefined	Undefined	Undefined
SPSO	Faster	Faster	Faster

	BBPSO	MBBPSO	SPSO
Gbest PSO	Slightly slower	Undefined	Slower
VNPSO	Slower	Undefined	Slower
Lbest PSO	Slower	Undefined	Slower
Gbest GCPSO	Slightly slower	Undefined	Slower
VNGCPSO	Slower	Undefined	Slower
Lbest GCPSO	Slower	Undefined	Slower
BBPSO	–	Faster	Slower
MBBPSO	Slower	–	Slower
SPSO	Faster	Faster	–

3.2.2 Experimental Procedure

The DRoC measure is tested as follows. DRoC measurements were obtained on various PSO algorithms running on various benchmark functions. All of the algorithms used, along with their respective parameters, are listed in Table 3.2. This study is not concerned with the performance of the algorithms, so no parameter tuning was performed; instead, well-known defaults were used. Each PSO was used with a swarm size of 25 particles. In order to prevent particles from roaming outside the bounds of the search space, the PSOs were only allowed to update their *pbest* and *nbest* positions if the new positions were within the search space bounds.

Table 3.3 lists all the benchmark functions used. For each benchmark function, each PSO was run 30 times from random starting positions, for 2000 iterations. A DRoC measurement was determined from these diversity measurements by using the SciPy library [23] to perform two-piecewise linear approximations. For each benchmark function, for each pair of PSO algorithms, the 30 DRoC measurements of each PSO algorithm were compared using a pair-wise Mann-Whitney U test with a 95% level of significance. The result of the rank test indicates whether the first PSO's DRoC measurements are significantly lower (-1) or higher (1) than those of the second PSO, or whether they are not significantly different (0). The rank result for each configuration is then compared against the intuitive expectations discussed in Section 3.2.1.

Table 3.2: Algorithms and control parameters used in this section

PSO Algorithm	Control parameters
Gbest PSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
Lbest PSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
VNPSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
Gbest GCPSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
Lbest GCPSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
VNGCPSO	$w = 0.729844$, $c_1 = 1.49618$, $c_2 = 1.49618$
BBPSO	(No control parameters to specify)
MBBPSO	Probability of recombination = 0.5
SPSO	$w = 0.729844$, $c_1 = 0$, $c_2 = 1.49618$

Table 3.3: Benchmark functions used in this section

Function name	Domain
Ackley [74]	$\mathbf{x} \in [-32, 32]^2$
Alpine [53]	$\mathbf{x} \in [-10, 10]^2$
Eggholder [45]	$\mathbf{x} \in [-512, 512]^2$
Goldstein-Price [74]	$\mathbf{x} \in [-2, 2]^2$
Griewank [74]	$\mathbf{x} \in [-600, 600]^2$
Levy 13 [45]	$\mathbf{x} \in [-10, 10]^2$
Michalewicz [44]	$\mathbf{x} \in [0, \pi]^2$
Quadric [74]	$\mathbf{x} \in [-100, 100]^2$
Quartic [74]	$\mathbf{x} \in [-1.28, 1.28]^2$
Rastrigin [74]	$\mathbf{x} \in [-5.12, 5.12]^2$
Rosenbrock [74]	$\mathbf{x} \in [-2.048, 2.048]^2$
Salomon [52]	$\mathbf{x} \in [-100, 100]^2$
Schwefel 2.22 [74]	$\mathbf{x} \in [-10, 10]^2$
Schwefel 2.26 [74]	$\mathbf{x} \in [-500, 500]^2$
Six-hump Camel Back [74]	$\mathbf{x} \in [-5, 5]^2$
Spherical [74]	$\mathbf{x} \in [-100, 100]^2$
Step [44]	$\mathbf{x} \in [-20, 20]^2$
Zakharov [44]	$\mathbf{x} \in [-5, 10]^2$

3.2.3 Results

Table 3.4 shows the rank between each pair of PSOs (columns) for each benchmark function (rows). The final three rows show the total number of each rank obtained across all the benchmark functions for each PSO pair. In the final three rows, the expected ranks for each PSO pair are indicated by shading the relevant rows; for pairs of PSOs where the difference is not expected to be significant, two rows are shaded, namely 0 (indicating an insignificant difference) and either -1 or 1, depending on the expected difference. The most-encountered rank is indicated by formatting the rank in bold.

When comparing measurements for the gbest PSO to those for the lesser-connected VNPSO, lbest PSO, VNGCPSO and lbest GCPSO, the gbest PSO converged at a faster

Table 3.4: Summarised results of pair-wise Mann-Whitney U tests for each pair of algorithms, for each benchmark function (continued on next page)

	Gbest PSO, VNPSO	Gbest PSO, Lbest PSO	Gbest PSO, Gbest GCPSO	Gbest PSO, VNGCPSO	Gbest PSO, Lbest GCPSO	Gbest PSO, BBPSO	Gbest PSO, SPSO	VNPSO, Lbest PSO	VNPSO, VNGCPSO	VNPSO, Lbest GCPSO	VNPSO, BBPSO	VNPSO, SPSO	Lbest PSO, Lbest GCPSO	Lbest PSO, BBPSO
Ackley	-1	-1	0	0	-1	1	0	-1	1	0	1	1	0	1
Alpine	0	-1	0	-1	-1	1	0	-1	0	-1	1	1	0	1
Eggholder	1	1	0	1	1	1	-1	1	0	1	0	-1	0	-1
Goldstein-Price	-1	-1	-1	-1	-1	1	0	0	-1	-1	1	1	-1	1
Griewank	0	-1	1	0	0	1	0	-1	1	0	1	1	1	1
Levy 13	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Michalewicz	-1	-1	0	-1	-1	1	0	-1	-1	-1	1	1	-1	1
Quadric	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Quartic	-1	-1	-1	-1	-1	1	0	-1	-1	-1	1	1	-1	1
Rastrigin	-1	-1	0	-1	-1	1	1	-1	0	-1	1	1	0	1
Rosenbrock	-1	-1	0	-1	-1	1	1	0	0	0	1	1	0	1
Salomon	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Schwefel 2.22	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Schwefel 2.22	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Six-hump	0	0	0	-1	-1	1	1	0	-1	-1	1	1	-1	1
Spherical	-1	-1	0	-1	-1	1	0	0	-1	-1	1	1	0	1
Step	-1	-1	0	-1	-1	1	0	0	1	0	1	1	0	1
Zakharov	-1	-1	0	-1	-1	1	0	0	0	0	1	1	0	1
Total -1's	**14**	**16**	2	**15**	**16**	0	1	6	5	7	0	1	4	1
Total 0's	3	1	**15**	2	1	0	**14**	**11**	**10**	**10**	1	0	**13**	0
Total 1's	1	1	1	1	1	**18**	3	1	3	1	**17**	**17**	1	**17**

Table 3.4: (Continued from previous page) Summarised results of pair-wise Mann-Whitney U tests for each pair of algorithms, for each benchmark function

	Lbest PSO, SPSO	Gbest GCPSO, VNGCPSO	Gbest GCPSO, Lbest GCPSO	Gbest GCPSO, BBPSO	Gbest GCPSO, SPSO	VNGCPSO, Lbest GCPSO	VNGCPSO, BBPSO	VNGCPSO, SPSO	Lbest GCPSO, BBPSO	Lbest GCPSO, SPSO	BBPSO, MBBPSO	BBPSO, SPSO	MBBPSO, SPSO
Ackley	1	0	-1	1	0	-1	1	0	1	1	-1	-1	0
Alpine	1	-1	-1	1	1	-1	1	1	1	1	-1	-1	1
Eggholder	-1	0	1	0	-1	1	0	-1	-1	-1	1	-1	-1
Goldstein-Price	1	-1	-1	1	1	-1	1	1	1	1	-1	-1	1
Griewank	1	0	-1	1	0	0	1	0	1	1	-1	-1	-1
Levy 13	1	0	-1	1	0	0	1	1	1	1	-1	-1	-1
Michalewicz	1	-1	-1	1	1	0	1	1	1	1	-1	-1	0
Quadric	1	-1	-1	1	0	0	1	1	1	1	-1	-1	1
Quartic	1	-1	-1	1	1	0	1	1	1	1	-1	-1	-1
Rastrigin	1	-1	-1	1	1	0	1	1	1	1	-1	0	1
Rosenbrock	1	-1	-1	0	1	0	1	1	1	1	-1	1	1
Salomon	1	0	-1	1	0	0	1	1	1	1	-1	-1	0
Schwefel 2.22	1	-1	-1	1	0	0	1	1	1	1	-1	-1	-1
Schwefel 2.22	1	-1	-1	1	0	0	1	1	1	1	-1	-1	-1
Six-hump	1	-1	-1	1	0	0	1	1	1	1	-1	-1	-1
Spherical	1	-1	-1	1	0	0	1	1	1	1	-1	-1	-1
Step	1	0	0	1	0	0	1	1	1	1	-1	-1	-1
Zakharov	1	-1	-1	1	0	0	1	1	1	1	-1	-1	1
Total -1's	1	**12**	**16**	0	1	3	0	1	1	1	**17**	**16**	**9**
Total 0's	0	6	1	2	**11**	**14**	1	2	0	0	0	1	3
Total 1's	**17**	0	1	**16**	6	1	**17**	**15**	**17**	**17**	1	1	6

rate for most benchmark functions, as expected. The measures also indicate that the gbest GCPSO reduced its diversity at a faster rate than the VNGCPSO and lbest GCPSO in most cases. The BBPSO, which has a star neighbourhood topology like the gbest PSO, was also found to reduce its diversity at a faster rate than the less-connected VNPSO, lbest PSO, VNGCPSO and lbest GCPSO in most cases.

When comparing measurements for the VNPSO to the lesser-connected lbest PSO, no significant difference was found between the two PSOs for most benchmark functions, contrary to expectations. However, for the cases where a significant difference was found between the two PSOs, the VNPSO was found to reduce its diversity at a faster rate in most cases. The significance of the difference in behaviour between the VNPSO and the lbest PSO may have been overestimated. Similarly, when comparing the VNPSO to the lesser-connected lbest GCPSO, and when comparing the VNGCPSO to the lbest GCPSO, no significant differences were found in most cases, with the few differences found aligning with expectations. These pairs of PSOs also appear to exhibit more similar behaviour than expected. Consider the difference in size between neighbourhoods in a VNPSO, whose neighbourhoods contain up to four particles, and an lbest PSO, whose neighbourhoods contain precisely two particles. The size of the neighbourhoods of these topologies are relatively similar compared to the neighbourhood of a gbest PSO, which contains all the particles in the swarm (25 in this study). This might explain why the VNPSO and the lbest PSO exhibit more similar behaviour than when comparing either of those algorithms with the gbest PSO.

When comparing the standard PSOs (gbest PSO, lbest PSO, VNPSO) to the GCPSO algorithms with the same neighbourhoods (gbest GCPSO, lbest GCPSO, VNGCPSO), no significant difference was measured for most of the benchmark functions, as expected. For the few differences found, the standard PSOs tended to converge at a faster rate than the guaranteed convergence variations, as expected.

When comparing measurements for the SPSO to those for PSOs with less-connected neighbourhoods (lbest PSO, VNSPO, lbest GCPSO, VNGCPSO), the SPSO was indicated to reduce its diversity at a faster rate in almost all cases, as expected. Contrary to expectations, all of the PSOs with less-connected neighbourhoods were measured as reducing their diversity at a faster rate than the SPSO for the Eggholder function. This

seems to indicate that the fitness landscape of a function may influence the DRoC; this possibility is further investigated in Chapter 4. The SPSO was expected to reduce its diversity at a faster rate than all other PSOs, but the DRoC measurements usually indicated no significant difference to the gbest PSO and the gbest GCPSO, although most of the remaining cases aligned with expectations, with only the Eggholder function producing a rank contrary to expectations. The difference in behaviour between the SPSO and the gbest PSO and the gbest GCPSO appears to have been overestimated. On the other hand, the BBPSO and MBBPSO were usually found to reduce their diversity at a faster rate than the SPSO, contrary to expectations.

Comparisons between measurements for the BBPSO and the MBBPSO align with the expectation that the BBPSO would reduce its diversity at a faster rate than the MBBPSO.

Both the gbest PSO and the gbest GCPSO were found to reduce their diversity at a slower rate than the BBPSO in most cases, as expected. Similarly, the gbest GCPSO was found to reduce its diversity at a faster rate than the MBBPSO in most cases, as expected.

3.2.4 Conclusions

This section tested the capability of the DRoC measure to indicate the search behaviour of PSO algorithms. The measure was evaluated by ranking DRoC measurements obtained on pairs of PSO algorithms, and testing whether those ranks aligned with the expected differences between the algorithms in terms of search behaviour. Most of the ranks did align with the expectations. In some cases where the ranks contradicted expectations, those contradictions were consistently observed for specific benchmark functions; Chapter 4 investigates whether the fitness landscape characteristics of these benchmark functions may have influenced the search behaviour of the PSOs in those cases. For the remaining cases where the ranks did not align with expectations, the deviation from expectations was not severe; in other words, no significant difference was usually found where one was expected, or vice versa, but few or no results outright contradicted expectations. These cases can simply be attributed to over- or underestimating the expected differences between the search behaviour of the relevant PSOs.

3.3 DRoC Robustness with Regards to Swarm Size

Section 3.2 showed that the DRoC measure is a good indicator of a PSO swarm's search behaviour under limited conditions. Specifically, for all tests, all PSOs were used with a swarm size of 25 particles. DRoC measurements are based on instantaneous diversity measurements taken on the swarm; therefore, the number of particles in a swarm may influence DRoC measurements taken on the swarm. Moreover, the number of particles in a swarm may influence the behaviour of the swarm in terms of exploration and exploitation, which may, in turn, influence DRoC measurements taken on the swarm.

This section investigates the robustness of the DRoC measure with regards to swarm size. The experimental procedure used is discussed in Section 3.3.1. The results are listed and discussed in Section 3.3.2. The findings of the section are summarised in Section 3.3.3.

3.3.1 Experimental Procedure

The procedure for testing the robustness of the DRoC measure with regards to swarm size was as follows. A number of experimental configurations were generated. Each configuration specifies two PSO algorithms to compare, a single swarm size which was used for both PSO algorithms, and a single benchmark function that both PSO algorithms were applied to. The algorithms that were used are listed in Table 3.2 along with their respective parameters. Each PSO was used with a swarm size of 5, 10, 25, 50, 75, 100 and 500 particles.

Table 3.5 lists all the benchmark functions used. The selected benchmark functions represent a variety of characteristics: the spherical function is smooth and unimodal,

Table 3.5: Benchmark functions used to test DRoC sensitivity to swarm size

Function name	Domain
Spherical [8]	$\mathbf{x} \in [-5.12, 5.12]^D$
Rastrigin [74]	$\mathbf{x} \in [-5.12, 5.12]^D$
Rosenbrock [74]	$\mathbf{x} \in [-2.048, 2.048]^D$
Weierstrass [44]	$\mathbf{x} \in [-0.5, 0.5,]^D$

the Rastrigin function is rugged on the micro scale and unimodal on the global scale, the Rosenbrock function is smooth and multimodal in four dimensions and higher, and Weierstrass function is rugged on the micro and macro scale. Each benchmark function was used in 5 dimensions. A configuration was generated for each combination of two unique PSOs, a swarm size, and a benchmark function.

Each configuration was then run as follows. The first PSO of the configuration was initialised with the specified swarm size, with the particles initialised at random positions within the bounds of the search space. The PSO was run on the benchmark function for 2000 iterations. The PSO was restricted from updating its *pbest* and *nbest* positions to positions outside the bounds of the search space. During its run, a diversity measurement was taken at each iteration. A DRoC measurement was determined for the first PSO from these diversity measurements. The process was then repeated for the second PSO of the configuration, resulting in a DRoC measurement for the second PSO.

Each configuration was run 30 times in order to obtain a statistically significant sample of DRoC measurements for each algorithm. For each configuration, the 30 DRoC measurements obtained for the first PSO were ranked against those obtained for the second PSO using a pair-wise Mann-Whitney U test with a 95% level of significance. The result of the rank test indicates whether the first PSO's DRoC measurements are significantly lower (-1) or higher (1) than those of the second PSO, or whether they are not significantly different (0). The rank result for each configuration was then compared against the expectations discussed in Section 3.2.1.

3.3.2 Results

Table 3.6 lists the rank results obtained by comparing the DRoC measurements of pairs of PSOs with a variety of swarm sizes. Only pairs of PSOs for which expectations have been hypothesised are listed and discussed here; the complete results are listed in Table B.1. Each sub-table lists the results for a particular pair of PSOs. Expected ranks are indicated in each sub-table for pairs of PSOs where significant differences were expected. Cells indicate the rank result for each benchmark function (rows), for each swarm size (columns). Rank results that contradict expectations are shaded.

Table 3.6: Ranks of DRoCs for PSO pairs with various swarm sizes

Swarm size:	5	10	25	50	75	100	500
(a) Gbest PSO, VNPSO					Expectation: -1		
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	0	-1	-1	-1	-1	-1	-1
Weierstrass	0	0	-1	-1	-1	-1	-1
(b) Gbest PSO, Lbest PSO					Expectation: -1		
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(c) Gbest PSO, Gbest GCPSO					Expectation: 0 or -1		
Spherical	0	0	0	0	0	0	0
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	0	-1	0	0	0	-1	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(d) Gbest PSO, VNGCPSO					Expectation: -1		
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(e) Gbest PSO, Lbest GCPSO					Expectation: -1		
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(f) Gbest PSO, BBPSO					Expectation: 0 or 1		
Spherical	1	1	1	1	1	1	1

Continued on next page ...

Table 3.6 (continued)

Swarm size:	5	10	25	50	75	100	500
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	1	0	1	0	0	0	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(g) Gbest PSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	0
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	0	0	-1	-1	-1
Weierstrass	1	1	1	1	1	0	0
(h) VNPSO, Lbest PSO						Expectation:	-1
Spherical	0	0	0	-1	-1	-1	-1
Rastrigin	0	0	0	1	1	1	1
Rosenbrock	0	0	-1	-1	-1	-1	-1
Weierstrass	0	0	-1	-1	-1	-1	-1
(i) VNPSO, VNGCPSO						Expectation:	0 or -1
Spherical	0	0	-1	-1	-1	-1	-1
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(j) VNPSO, Lbest GCPSO						Expectation:	-1
Spherical	-1	0	-1	-1	-1	-1	-1
Rastrigin	0	0	0	1	1	1	1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(k) VNPSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1

Continued on next page ...

Table 3.6 (continued)

Swarm size:	5	10	25	50	75	100	500
(l) VNPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(m) Lbest PSO, Lbest GCPSO					Expectation:	0 or -1	
Spherical	-1	0	0	0	-1	-1	-1
Rastrigin	0	0	0	0	1	0	1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(n) Lbest PSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(o) Lbest PSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(p) Gbest GCPSO, VNGCPSO						Expectation:	-1
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(q) Gbest GCPSO, Lbest GCPSO						Expectation:	-1
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1

Continued on next page ...

Table 3.6 (continued)

Swarm size:	5	10	25	50	75	100	500
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(r) Gbest GCPSO, BBPSO						Expectation:	0 or 1
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	1	1	1	0	0	1	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(s) Gbest GCPSO, SPSO						Expectation:	1
Spherical	1	1	0	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	0	0	-1	0	-1
Weierstrass	1	1	1	1	1	1	0
(t) VNGCPSO, Lbest GCPSO						Expectation:	-1
Spherical	0	0	0	0	0	0	-1
Rastrigin	0	-1	0	1	1	1	1
Rosenbrock	0	0	-1	-1	-1	-1	-1
Weierstrass	0	0	0	0	0	0	-1
(u) VNGCPSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	-1	-1	-1	-1	0	0
(v) VNGCPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(w) Lbest GCPSO, BBPSO						Expectation:	1

Continued on next page ...

Table 3.6 (continued)

Swarm size:	5	10	25	50	75	100	500
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	-1	-1	-1	-1	0	0
(x) Lbest GCPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(y) BBPSO, MBBPSO						Expectation:	-1
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	0	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	0	-1	0	0	0	0
(z) BBPSO, SPSO						Expectation:	1
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	0	0	-1	0	-1	-1	-1
Weierstrass	1	1	1	1	1	1	1
(aa) MBBPSO, SPSO						Expectation:	1
Spherical	1	0	-1	0	-1	-1	-1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1

Of the 756 results, 619 are in line with expectations. The remaining results are discussed below.

Unexpected Results for Small Swarms

A number of configurations in Table 3.6 contradicted expectations for PSOs used with small swarm sizes, namely:

- The gbest PSO and the VNPSO on the Spherical function
- The gbest PSO and the lbest PSO on the Spherical function
- The gbest PSO and the VNGCPSO on the Spherical function
- The VNPSO and the lbest PSO on the Spherical, Rosenbrock and Weierstrass functions
- The VNPSO and the lbest GCPSO on the Spherical function
- The VNPSO and the BBPSO on the Rastrigin function
- The lbest PSO and the BBPSO on the Rastrigin function
- The VNGCPSO and the lbest GCPSO on the Rosenbrock function
- The VNGCPSO and the BBPSO on the Rastrigin function
- The lbest GCPSO and the BBPSO on the Rastrigin function
- The BBPSO and the MBBPSO on the Rastrigin function

For each case identified above, Figure 3.2 plots the mean diversity at each iteration of the two PSOs being compared. Many of the plots demonstrate that, although their corresponding rank results are unexpected, those rank results are indeed accurate. For example, Figure 3.2a shows that, on average, the gbest and VNPSOs transitioned from exploration to exploitation at a highly similar rate. This is in line with the corresponding rank of 0 in Table 3.6, which indicated that there was no significant difference between the search behaviour of the two algorithms. Similarly, in each of Figures 3.2b–e, h–j, m, and n, the diversities of the two PSOs being compared are very similar to one another. All of these comparisons were thus correctly ranked as 0. In Figures 3.2f, l, q, and s, the diversity measurements for one of the PSOs decrease in a relatively less linear fashion,

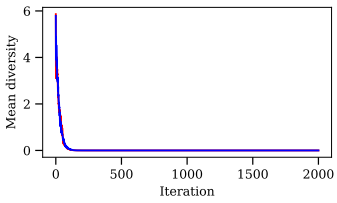

(a) Gbest PSO (red) and VNPSO (blue) with 5 particles, Spherical

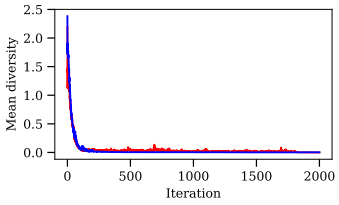

(b) Gbest PSO (red) and VNPSO (blue) with 5 particles, Rosenbrock

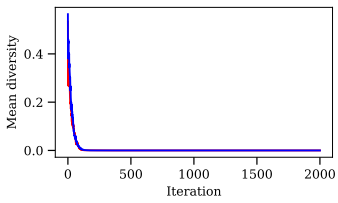

(c) Gbest PSO (red) and VNPSO (blue) with 5 particles, Weierstrass

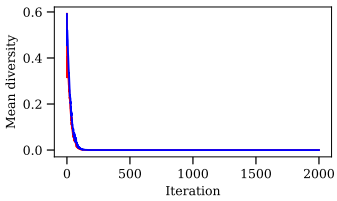

(d) Gbest PSO (red) and VNPSO (blue) with 10 particles, Weierstrass

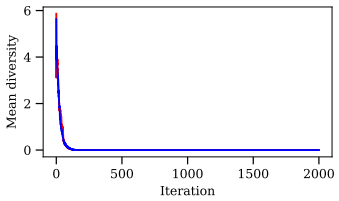

(e) Gbest PSO (red) and Lbest PSO (blue) with 5 particles, Spherical

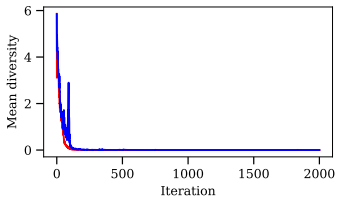

(f) Gbest PSO (red) and VNGCPSO (blue) with 5 particles, Spherical

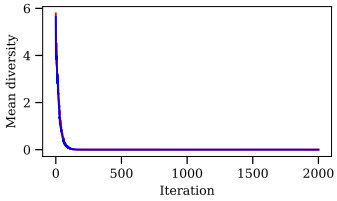

(g) VNPSO (red) and Lbest PSO (blue) with 5 particles, Spherical

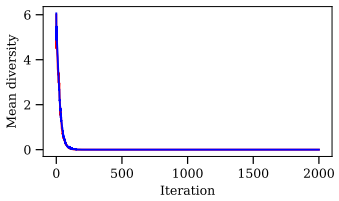

(h) VNPSO (red) and Lbest PSO (blue) with 10 particles, Spherical

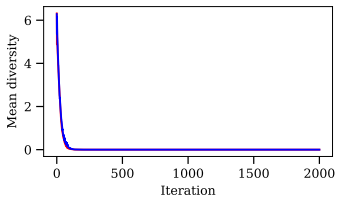

(i) VNPSO (red) and Lbest PSO (blue) with 25 particles, Spherical

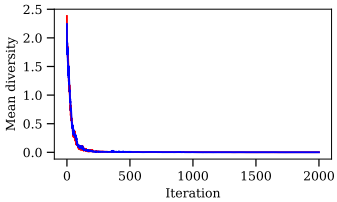

(j) VNPSO (red) and Lbest PSO (blue) with 5 particles, Rosenbrock

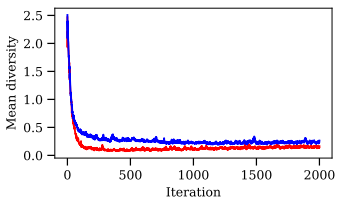

(k) VNPSO (red) and Lbest PSO (blue) with 10 particles, Rosenbrock

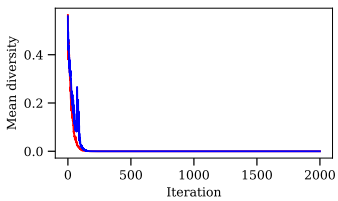

(l) VNPSO (red) and Lbest PSO (blue) with 5 particles, Weierstrass

Figure 3.2: Mean diversities at each iteration for each pair of PSOs where the DRoC measurements ranked as 0 with small swarms, contrary to expectations (Continued on next page)

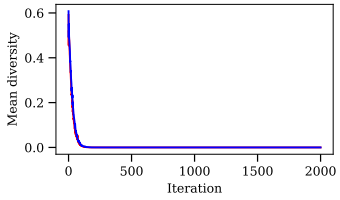

(m) VNPSO (red) and Lbest PSO (blue) with 10 particles, Weierstrass

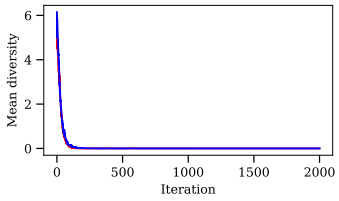

(n) VNPSO (red) and Lbest GCPSO (blue) with 10 particles, Spherical

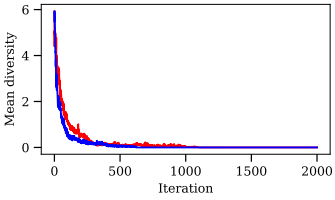

(o) VNPSO (red) and BBPSO (blue) with 5 particles, Rastrigin

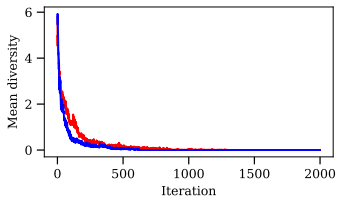

(p) Lbest PSO (red) and BBPSO (blue) with 5 particles, Rastrigin

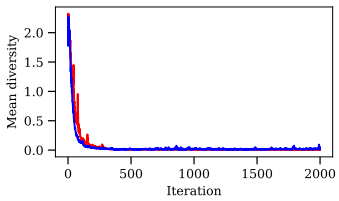

(q) VNGCPSO (red) and Lbest GCPSO (blue) with 5 particles, Rosenbrock

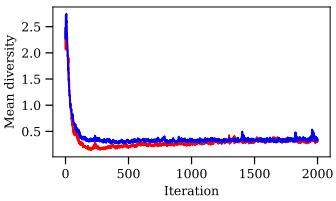

(r) VNGCPSO (red) and Lbest GCPSO (blue) with 10 particles, Rosenbrock

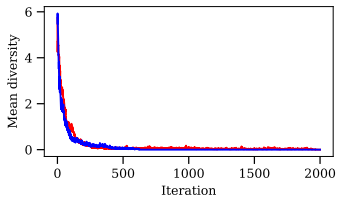

(s) VNGCPSO (red) and BBPSO (blue) with 5 particles, Rastrigin

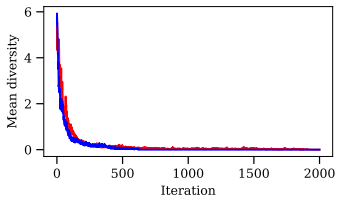

(t) Lbest GCPSO (red) and BBPSO (blue) with 5 particles, Rastrigin

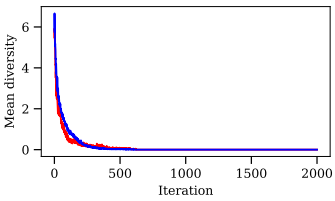

(u) BBPSO (red) and MBBPSO (blue) with 5 particles, Rastrigin

Figure 3.2: (Continued from previous page) Mean diversities at each iteration for each pair of PSOs where the DRoC measurements ranked as 0 with small swarms, contrary to expectations

exhibiting erratic spikes in diversity. However, for both cases, the spikes in diversity measurements occur while the PSOs are in an explorative state, and in both cases, the two PSOs being compared complete transitioning to an exploitative state after roughly the same number of iterations. This is in line with the corresponding ranks of 0 as seen in Table 3.6.

Some of the plots indicate that the ranked DRoC measurements failed to capture the difference in search behaviour between the PSO algorithms being compared. In Figures 3.2k, o, p, r, t, and u, it is clear that one PSO generally started exploiting the search space at a slightly earlier iteration than the other, which contradicts their corresponding rank results of 0. Notably, in some of these figures, the diversity measurements of both PSOs being compared are very similar to one another up until the point when one PSO transitions from exploration to exploitation before the other.

Figure 3.3 shows two-piecewise linear approximations of the diversity measurements of the configurations listed above, for which the ranked DRoC results failed to indicate the difference in search behaviour between the two PSOs being compared. In Figure 3.3a and d, the slopes of the first lines of the approximations are very similar among the two PSOs being compared. The DRoC measure is only based on the slope of the first line, which would cause the DRoC measurements taken on the pairs of algorithms listed above to be very similar, and to be ranked as not statistically significantly different. This demonstrates a limitation of the DRoC measurement: it only indicates the rate at which a swarm decreases its diversity during the exploration phase, without indicating the duration of this explorative phase. Keeping this limitation in mind, the DRoC measurements did correlate accurately with the diversity measurements in terms of the rate at which the PSOs decreased their diversities.

Figures 3.3b, c, e and f each demonstrates a notable difference between the first approximation lines of the PSOs being compared. For each of these figures, the difference is actually in line with expectations. For example, in Figure 3.3b, the slope of the first line corresponding to the lbest PSO is slightly more gradual than that of the BBPSO, indicating that the BBPSO reduced its diversity at a faster rate than the lbest PSO. As stated in Section 3.2.1, this is as expected; the BBPSO uses a more connected star topology than the lbest PSO's ring topology. The diversity measurements in Figure 3.2n also indicate that the BBPSO did, indeed, reduce its diversity at a slightly faster rate during the explorative phase. However, as indicated by the rank result of 0 for each of these cases, this difference was not statistically significant.

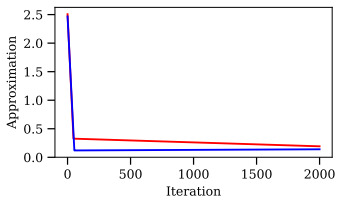

(a) Lbest PSO (red) and VNPSO (blue) with 10 particles, Rosenbrock

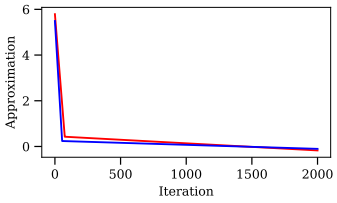

(b) Lbest PSO (red) and BBPSO (blue) with 5 particles, Rastrigin

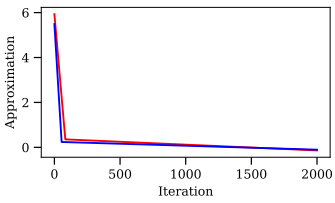

(c) VNPSO (red) and BBPSO (blue) with 5 particles, Rastrigin

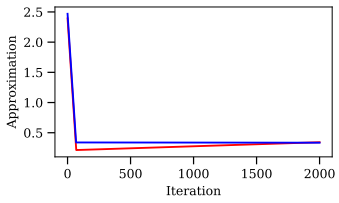

(d) VNGCPSO (red) and Lbest GCPSO (blue) with 10 particles, Rosenbrock

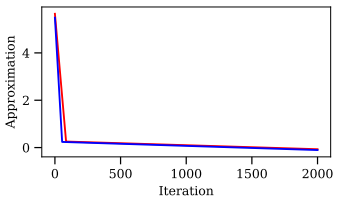

(e) Lbest GCPSO (red) and BBPSO (blue) with 5 particles, Rastrigin

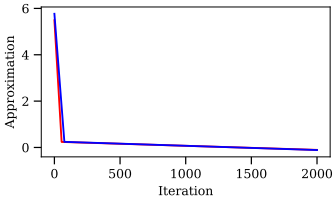

(f) BBPSO (red) and MBBPSO (blue) with 5 particles, Rastrigin

Figure 3.3: Mean two-piecewise linear approximations of diversity measurements for small swarms

Unexpected Results for Large Swarms

Some results in Table 3.6 for PSOs used with high swarm sizes contradicted expectations, namely:

- The gbest PSO and the SPSO on the Spherical and Weierstrass functions
- The gbest GCPSO and the SPSO on the Weierstrass function
- The lbest PSO and the lbest GCPSO on the Rastrigin function

For each case identified above, Figure 3.4 shows the mean of the diversity measurements of the PSOs being compared. Figures 3.4a–c demonstrate that the gbest PSO and the SPSO did indeed exhibit highly similar search behaviour in the respective cases, as indicated by the ranks of zero. The search behaviour of the gbest GCPSO was also

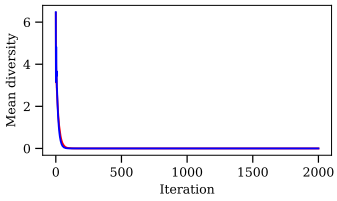

(a) Gbest PSO (red) and SPSO (blue), 500 particles, Spherical, ranked 0

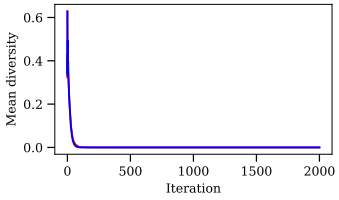

(b) Gbest PSO (red) and SPSO (blue), 100 particles, Weierstrass, ranked 0

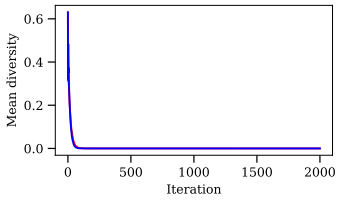

(c) Gbest PSO (red) and SPSO (blue), 500 particles, Weierstrass, ranked 0

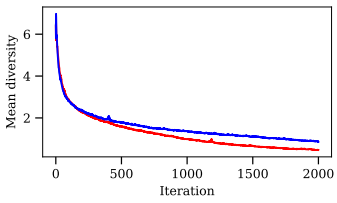

(d) Lbest PSO (red) and Lbest GCPSO (blue), 75 particles, Rastrigin, ranked 1

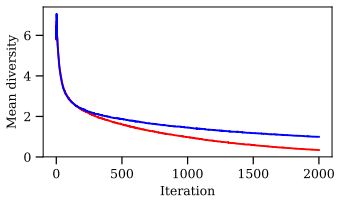

(e) Lbest PSO (red) and Lbest GCPSO (blue), 500 particles, Rastrigin, ranked 1

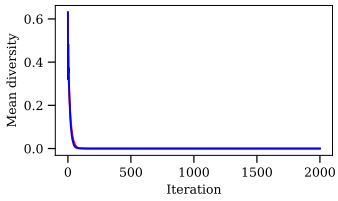

(f) Gbest GCPSO (red) and SPSO (blue), 500 particles, Weierstrass, ranked 0

Figure 3.4: Mean diversities at each iteration for each pair of PSOs where the DRoC measurements did not compare as expected with large swarms

similar to that of the SPSO on the Weierstrass function with a large swarm, as shown in Figure 3.4f, which corresponds with the rank of 0 for this configuration.

In Figures 3.4d and e, it appears that the lbest GCPSO reduced its diversity at a slower rate than the lbest PSO on the Rastrigin function with 75 and 500 particles, as expected; however, in both cases, the ranks indicated that the lbest GCPSO reduced its diversity at a faster rate. The average two-piecewise linear approximations for each of these cases are shown in Figures 3.5a and b. In these figures, it is clear that the slope for the first line of the lbest GCPSO is slightly steeper than that of the lbest PSO. The DRoC measure appears to fail to capture the behaviour of the swarms for these cases, possibly due to the relatively high diversity that the lbest GCPSO maintains for most of its execution.

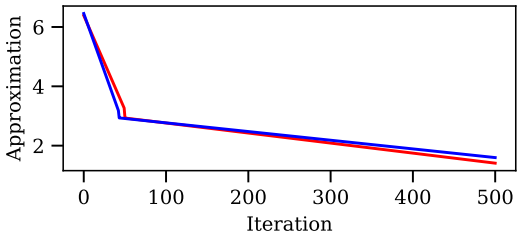

(a) Lbest PSO (red) and Lbest GCPSO (blue), 75 particles, Rastrigin, ranked 1

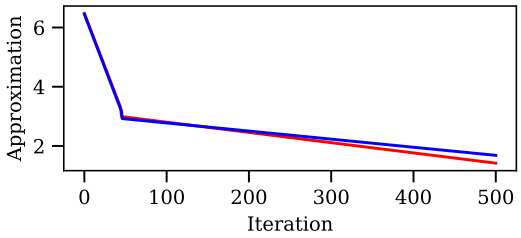

(b) Lbest PSO (red) and Lbest GCPSO (blue), 500 particles, Rastrigin, ranked 1

Figure 3.5: Mean two-piecewise linear approximations of diversity measurements for large swarms

Unexpected Results Across Swarm Sizes

For some PSO pairs and benchmark functions, some of the results in Table 3.6 contradicted expectations for all or most swarm sizes. This was the case for:

- The gbest PSO and the BBPSO on the Weierstrass function
- The gbest PSO and the SPSO on the Rosenbrock function
- The VNPSO and the lbest PSO on the Rastrigin function
- The VNPSO and the lbest GCPSO on the Rastrigin function
- The VNPSO and the BBPSO on the Weierstrass function
- The lbest PSO and the BBPSO on the Weierstrass function
- The gbest GCPSO and the BBPSO on the Weierstrass function
- The gbest GCPSO and the SPSO on the Rosenbrock function
- The VNGCPSO and the lbest GCPSO on the Spherical, Rastrigin and Weierstrass functions
- The VNGCPSO and the BBPSO on the Weierstrass function

- The BBPSO and the MBBPSO on the Weierstrass function
- The BBPSO and the SPSO on the Spherical and Rosenbrock functions
- The MBBPSO and the SPSO on the Spherical function

For each case identified above, Figure 3.6 plots the average diversity for all swarm sizes at each iteration. Figures 3.6a, e–g, l and m demonstrate that, for cases involving the BBPSO or the MBBPSO applied to the Weierstrass, the BBPSO and the MBBPSO failed to transition to an exploitative state altogether. Despite the fact that these PSOs were restricted to only update their *pbest* and *gbest* positions to positions within the search space bounds, the particles left the search space and remained out of bounds for the duration of the algorithms' runs. For cases such as these, the DRoC measurements were zeros, indicating that the swarm increased its diversity instead of decreasing it[1]. The DRoC measure is based on the assumption that the PSO being measured will, loosely speaking, decrease its diversity over time. Therefore, a limitation of the measure is that it can only be used to measure the rate at which a PSO transitions from exploration to exploitation for cases where the PSO makes the transition. Nonetheless, given that the diversity-increasing behaviour of the BBPSO and MBBPSO algorithms in these configurations, the ranks correctly indicated the relative rate at which those PSOs reduced their diversities.

In figures 3.6i and 3.6k, it is clear that the two PSOs being compared exhibited highly similar search behaviour. The corresponding ranks of 0 are thus correct.

In figures 3.6n and 3.6p, the BBPSO and the MBBPSO reduced their diversities at slightly faster rates than the SPSOs, as correctly indicated by the ranks of -1.

Figures 3.6b, h and o, in which diversities are shown for a smaller selection of iterations, each shows that one of the PSOs being compared increased its diversity somewhat after the initial reduction in diversity. For each of these cases, Figures 3.7a–c shows the average diversities alongside the resulting two-piecewise linear approximations of each PSO being compared. The figures demonstrates the effect that the increasing diversity

[1] The reason why the DRoC values were zeros instead of positive numbers is because of an implementation detail of the DRoC measure: in order to aid the two-piecewise linear approximation process, the slope of the first line was limited to a maximum value of zero.

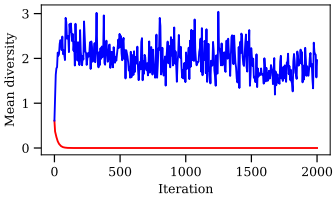

(a) Gbest PSO (red) and BBPSO (blue), Weierstrass

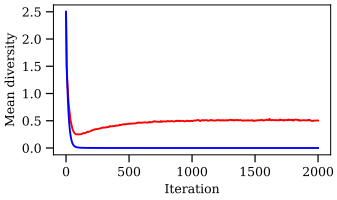

(b) Gbest PSO (red) and SPSO (blue), Rosenbrock

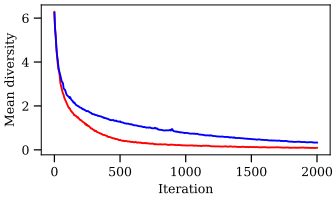

(c) VNPSO (red) and Lbest PSO (blue), Rastrigin

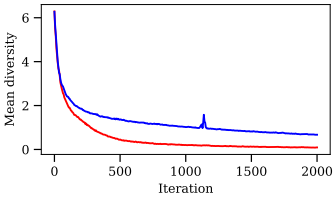

(d) VNPSO (red) and Lbest GCPSO (blue), Rastrigin

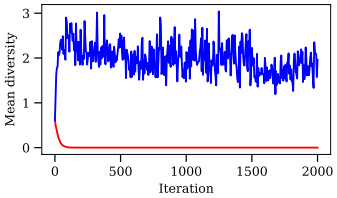

(e) VNPSO (red) and BBPSO (blue), Weierstrass

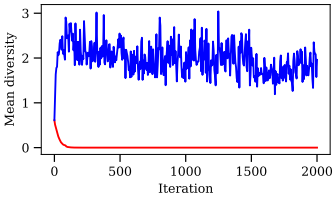

(f) Lbest PSO (red) and BBPSO (blue), Weierstrass

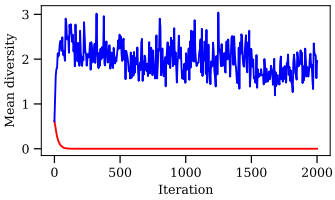

(g) Gbest GCPSO (red) and BBPSO (blue), Weierstrass

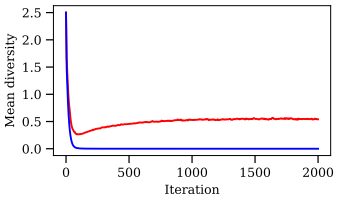

(h) Gbest GCPSO (red) and SPSO (blue), Rosenbrock

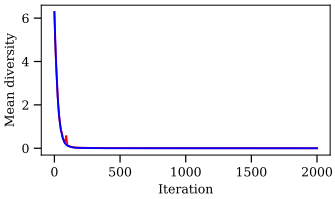

(i) VNGCPSO (red) and Lbest GCPSO (blue), Spherical

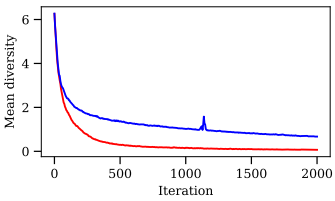

(j) VNGCPSO (red) and Lbest GCPSO (blue), Rastrigin

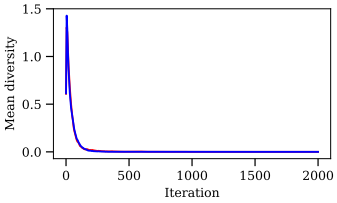

(k) VNGCPSO (red) and Lbest GCPSO (blue), Weierstrass

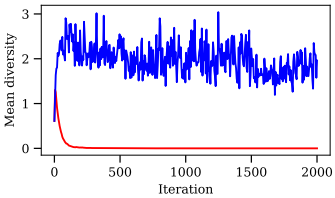

(l) VNGCPSO (red) and BBPSO (blue), Weierstrass

Figure 3.6: Mean diversities at each iteration for each pair of PSOs where the DRoC measurements did not compare as expected across swarm sizes (Continued on next page)

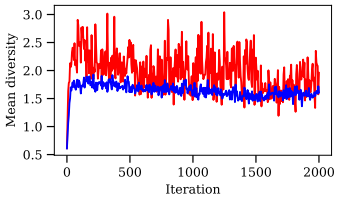

(m) BBPSO (red) and MBBPSO (blue), Weierstrass

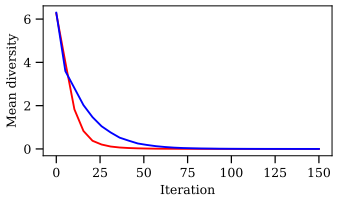

(n) BBPSO (red) and SPSO (blue), Spherical

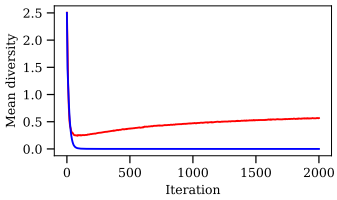

(o) BBPSO (red) and SPSO (blue), Rosenbrock

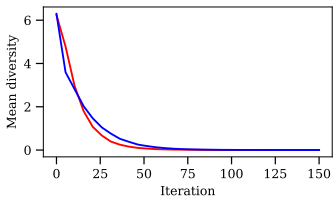

(p) MBBPSO (red) and SPSO (blue), Spherical

Figure 3.6: (Continued from previous page) Mean diversities at each iteration for each pair of PSOs where the DRoC measurements did not compare as expected across swarm sizes

had on the fitted two-piecewise linear approximations. For each of these cases the break point between the two lines was fitted relatively early for the PSO that was expected to reduce its diversity at a slower rate compared to the other PSO. This resulted in similar or even steeper first lines for the PSOs where those lines were expected to be less steep.

In each of Figures 3.6c, d and j, the Von Neumann topology-based PSO is shown to reduce its diversity at a faster rate than the ring topology-based PSO, as expected. However, in each of these figures, the diversities of the two PSOs are very similar during the initial iterations. This explains why the pairs of PSOs are ranked 0 for smaller swarm sizes, but not why they are ranked as 1 for higher swarm sizes. For each of these cases, Figures 3.7d–f shows the average diversities alongside the resulting two-piecewise linear approximations of each PSO being compared. The figures demonstrate that, while the lbest PSO and the lbest GCPSO had a relatively smooth transition from explorative to exploitative behaviour compared to the VNPSO and the VNGCPSO, the two-piecewise linear approximations were fit to indicate that the transitions occurred relatively early

compared to the VNPSO and VNGCPSO, and that the rate at which the lbest PSO decreased its diversity up until that point was relatively fast.

3.3.3 Summary

This section investigated the DRoC measure when used for PSO algorithms with various swarm sizes. Based on expectations regarding the behaviour of various PSOs, the measure provided accurate results in most cases, with a few exceptions.

For many of the exceptions, it was found that the PSOs exhibited unexpected behaviour, which was then successfully reflected in the DRoC measurements for these cases.

For some cases where the PSOs were used with small swarms, the measure accurately indicated the rate at which the PSOs decreased their diversities, but not the duration of their explorative behaviour; this is a limitation of the measure.

For some cases where the PSOs were used with large swarms, one of the two PSOs being compared exhibited a very smooth transition from explorative to exploitative behaviour. In these cases, it would be sensible to interpret the transition as relatively slow; however, the two-piecewise linear approximation on which the DRoC measure is based was fitted with the break point between the two lines positioned at a relatively early iteration. From these cases, it appears that, when a PSO makes a slow and smooth transition from exploration to exploitation, and continues to decrease its diversity during the exploitative phase, the DRoC measure may fail to provide an accurate measurement.

For some pairs of PSOs on some benchmark functions, for most or all swarm sizes, one of the PSOs being compared exhibited a trend of increasing diversity after the initial transition from exploration to exploitation. In these cases, the two-piecewise linear approximation was also fitted with the break point between the two lines positioned at a relatively early iteration. It appears that the DRoC measure is not robust in cases where the diversity of one PSO increases after transitioning from exploration to exploitation.

Aside from the limitations discussed above, the DRoC measure is concluded to be quite robust with respect to the swarm size of the PSO.

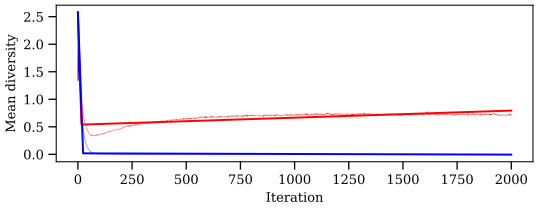

(a) Diversities and linear approximation for Gbest PSO (red) and SPSO (blue), Rosenbrock

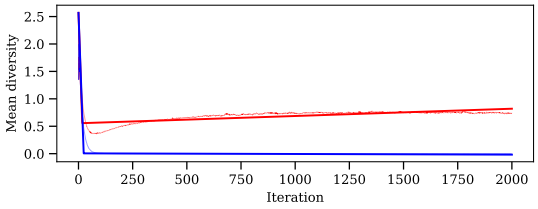

(b) Diversities and linear approximation for Gbest GCPSO (red) and SPSO (blue), Rosenbrock

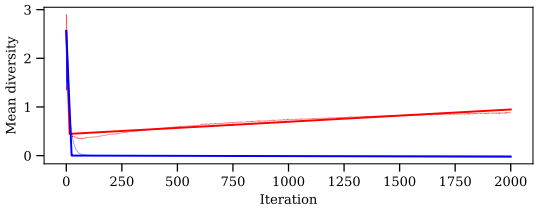

(c) Diversities and linear approximation for BBPSO (red) and SPSO (blue), Rosenbrock

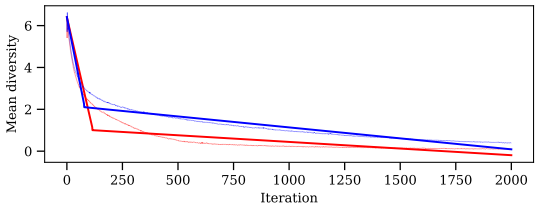

(d) Diversities and linear approximation for VNPSO (red) and Lbest PSO (blue), Rastrigin

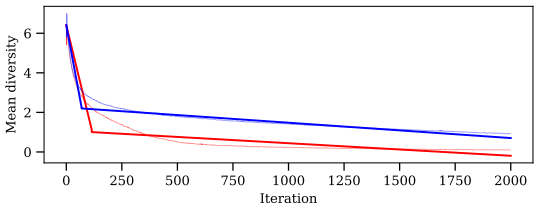

(e) Diversities and linear approximation for VNPSO (red) and Lbest GCPSO (blue), Rastrigin

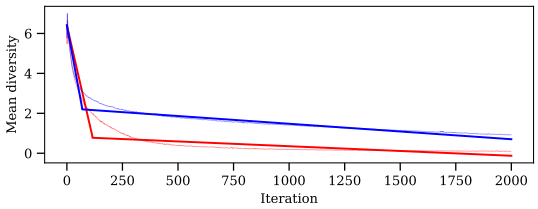

(f) Diversities and linear approximation for VNGCPSO (red) and Lbest GCPSO (blue), Rastrigin

Figure 3.7: Mean diversity measurements at each iteration for swarm sizes ranging from 50 to 500 particles, with mean two-piecewise linear approximations

3.4 DRoC Robustness with Regards to Dimensionality

The dimensionality of a search space influences the diversity of a swarm. Therefore, the dimensionality of a search space may impact the accuracy of the DRoC measure.

Similar to Section 3.3, this section investigates the robustness of the DRoC measure with regards to the dimensionality of the search space. The experimental procedure is provided in Section 3.4.1. Results are listed and discussed in Section 3.4.2. The findings of the section are summarised in Section 3.4.3.

3.4.1 Experimental Procedure

The procedure for testing the DRoC measure's robustness with regards to search space dimensionality is similar to the procedure used in Section 3.3.1, but instead of testing for various swarm sizes, various dimensionalities were used. A configuration was generated for each pair of PSOs listed in Table 3.2, each benchmark function listed in Table 3.5, and each of the following dimensionalities: 5, 25, 50, 100, 500 and 1000. Each PSO was used with 25 particles initialised at random positions in the search space, and run for 2000 iterations. Each configuration was run 30 times, giving 30 DRoC measurements for each PSO in the configuration. The DRoC measurements for the two PSOs in each configuration were then ranked against one another to determine which, if any, were significantly greater than the other.

3.4.2 Results

Table 3.7 lists the rank results obtained by comparing the DRoC measurements of pairs of PSOs in search spaces with various dimensionalities. As with Section 3.3.2, this section only discusses results for which expectations have been previously discussed; the complete results are given in Table B.2. Each sub-table lists the results for a particular pair of PSOs, with expected ranks indicated in each sub-table, and unexpected results shaded.

Table 3.7: Ranks of DRoCs for PSO pairs in various dimensionalities

Dimensionality:	2	5	25	50	100	500	1000
(a) Gbest PSO, VNPSO					Expectation:		-1
Spherical	-1	-1	-1	-1	0	0	0
Rastrigin	-1	-1	-1	0	0	0	0

Continued on next page ...

Table 3.7 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Rosenbrock	-1	-1	-1	-1	0	0	0
Weierstrass	-1	-1	-1	-1	0	0	0
(b) Gbest PSO, Lbest PSO						Expectation:	-1
Spherical	-1	-1	-1	-1	-1	0	-1
Rastrigin	-1	-1	-1	0	-1	0	-1
Rosenbrock	-1	-1	-1	-1	0	-1	0
Weierstrass	-1	-1	-1	-1	0	0	-1
(c) Gbest PSO, Gbest GCPSO					Expectation:	0 or	-1
Spherical	0	0	0	0	-1	0	0
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	0	0	-1	0	0	0	0
Weierstrass	-1	-1	0	0	0	0	0
(d) Gbest PSO, VNGCPSO						Expectation:	-1
Spherical	-1	-1	-1	1	0	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	-1	1	-1	-1	0
Weierstrass	-1	-1	-1	0	0	0	-1
(e) Gbest PSO, Lbest GCPSO						Expectation:	-1
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	0	-1	0	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
(f) Gbest PSO, BBPSO						Expectation:	0 or 1
Spherical	1	1	0	1	0	0	0
Rastrigin	1	1	0	1	0	0	0
Rosenbrock	1	1	-1	-1	0	-1	0
Weierstrass	1	-1	-1	-1	1	1	1
(g) Gbest PSO, SPSO						Expectation:	1

Continued on next page ...

Table 3.7 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Spherical	0	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	0	1	1	1	1	1
Weierstrass	0	1	1	1	1	1	1
(h) VNPSO, Lbest PSO						Expectation:	-1
Spherical	0	0	-1	-1	0	0	-1
Rastrigin	-1	0	-1	0	-1	0	-1
Rosenbrock	0	-1	-1	-1	0	-1	0
Weierstrass	0	-1	-1	0	0	0	-1
(i) VNPSO, VNGCPSO					Expectation:	0 or	-1
Spherical	-1	-1	1	1	0	-1	-1
Rastrigin	0	0	1	1	-1	-1	-1
Rosenbrock	0	-1	1	1	0	-1	0
Weierstrass	-1	-1	-1	1	0	0	-1
(j) VNPSO, Lbest GCPSO						Expectation:	-1
Spherical	-1	-1	0	0	0	-1	-1
Rastrigin	-1	0	0	0	-1	-1	-1
Rosenbrock	0	-1	-1	0	-1	-1	-1
Weierstrass	-1	-1	-1	0	-1	-1	-1
(k) VNPSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	0	0	0
Rastrigin	1	1	1	1	0	0	0
Rosenbrock	1	1	0	0	0	0	0
Weierstrass	1	-1	-1	0	1	1	1
(l) VNPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1

Continued on next page ...

Table 3.7 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Weierstrass	1	1	1	1	1	1	1
(m) Lbest PSO, Lbest GCPSO					Expectation: 0 or -1		
Spherical	0	0	1	1	0	-1	-1
Rastrigin	0	0	1	1	0	-1	-1
Rosenbrock	0	-1	0	1	-1	-1	-1
Weierstrass	-1	-1	-1	0	0	-1	0
(n) Lbest PSO, BBPSO						Expectation: 1	
Spherical	1	1	1	1	0	0	1
Rastrigin	1	1	1	1	0	0	1
Rosenbrock	1	1	1	1	0	0	0
Weierstrass	1	-1	-1	1	1	1	1
(o) Lbest PSO, SPSO						Expectation: 1	
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(p) Gbest GCPSO, VNGCPSO						Expectation: -1	
Spherical	-1	-1	-1	0	0	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	-1	0	0	-1	0
Weierstrass	-1	-1	-1	0	0	0	-1
(q) Gbest GCPSO, Lbest GCPSO						Expectation: -1	
Spherical	-1	-1	-1	-1	0	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	0	-1	-1
(r) Gbest GCPSO, BBPSO						Expectation: 0 or 1	
Spherical	1	1	0	1	0	0	0

Continued on next page ...

Table 3.7 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Rastrigin	1	1	0	1	-1	0	0
Rosenbrock	0	1	-1	-1	0	0	0
Weierstrass	1	-1	-1	0	1	1	1
(s) Gbest GCPSO, SPSO						Expectation:	1
Spherical	0	0	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	0	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(t) VNGCPSO, Lbest GCPSO						Expectation:	-1
Spherical	0	0	-1	-1	0	0	1
Rastrigin	0	0	-1	-1	0	1	0
Rosenbrock	0	-1	-1	-1	0	0	0
Weierstrass	0	0	-1	-1	0	0	0
(u) VNGCPSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	0	1	1
Rastrigin	1	1	1	1	0	1	1
Rosenbrock	1	1	0	-1	0	1	0
Weierstrass	1	-1	-1	-1	1	1	1
(v) VNGCPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(w) Lbest GCPSO, BBPSO						Expectation:	1
Spherical	1	1	1	1	0	1	1
Rastrigin	1	1	1	1	0	1	1
Rosenbrock	1	1	1	0	1	1	1
Weierstrass	1	-1	-1	1	1	1	1

Continued on next page ...

Table 3.7 (continued)

Dimensionality:	2	5	25	50	100	500	1000
(x) Lbest GCPSO, SPSO						Expectation:	1
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
(y) BBPSO, MBBPSO						Expectation:	-1
Spherical	-1	-1	0	1	1	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	1	1	-1	0	-1
Weierstrass	-1	-1	1	1	0	0	0
(z) BBPSO, SPSO						Expectation:	1
Spherical	-1	-1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	-1	1	1	1	1	1
Weierstrass	-1	1	1	1	1	1	1
(aa) MBBPSO, SPSO						Expectation:	1
Spherical	-1	-1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	1	1	1	1	0	1

Of the 756 results, 586 are as expected. The remaining results are discussed below.

Figure 3.8 shows three examples of test cases where one or both of the PSOs failed to transition from exploration to exploitation. As discussed in Section 3.3.2, the DRoC measure assumes that the PSO being measured will decrease its diversity. Cases where the PSO increases its diversity result in non-negative DRoC measurements.

Table 3.8 lists the average DRoC measurements obtained for each PSO, benchmark function, and dimensionality studied in this section, with cells highlighted for non-negative DRoC values.

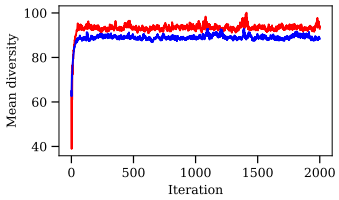

(a) Gbest PSO (red) and Lbest PSO (blue), Spherical in 500 dimensions, ranked 0

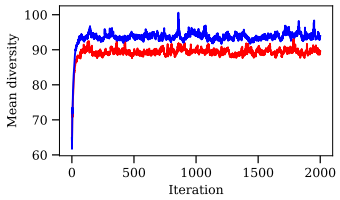

(b) Lbest PSO (red) and VNPSO (blue), Rastrigin in 500 dimensions, ranked 0

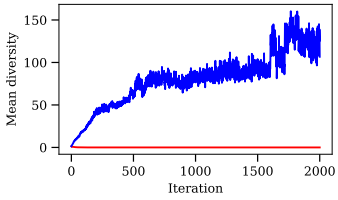

(c) SPSO (red) and MBBPSO (blue), Weierstrass in 25 dimensions, ranked -1

Figure 3.8: Example mean diversities for PSO pairs where one or both PSOs did not reduce diversity

Table 3.8: Average DRoC measurements for various dimensionalities

Dimensionality:	2	5	25	50	100	500	1000
(a) Gbest PSO							
Spherical	-0.206	-0.219	-0.146	-0.038	0	0	0
Rastrigin	-0.129	-0.121	-0.081	-0.004	0	0	0
Rosenbrock	-0.037	-0.102	-0.054	-0.01	0	0	0
Weierstrass	-0.019	-0.02	-0.012	-0.002	0	0	0
(b) VNPSO							
Spherical	-0.166	-0.165	-0.087	-0.019	0	0	0
Rastrigin	-0.076	-0.058	-0.035	0	0	0	0
Rosenbrock	-0.024	-0.054	-0.031	-0.003	0	0	0
Weierstrass	-0.017	-0.013	-0.007	0	0	0	0
(c) Lbest PSO							
Spherical	-0.158	-0.154	-0.07	-0.001	0	0	0
Rastrigin	-0.057	-0.058	-0.023	0	0	0	0
Rosenbrock	-0.024	-0.048	-0.022	0	0	0	0
Weierstrass	-0.014	-0.011	-0.004	0	0	0	0

Continued on next page . . .

Table 3.8 (continued)

Dimensionality:	2	5	25	50	100	500	1000
(d) Gbest GCPSO							
Spherical	-0.212	-0.247	-0.155	-0.043	0	0	0
Rastrigin	-0.131	-0.115	-0.082	-0.007	0	0	0
Rosenbrock	-0.04	-0.088	-0.047	-0.012	0	0	0
Weierstrass	-0.015	-0.015	-0.011	-0.002	0	0	0
(e) VNGCPSO							
Spherical	-0.14	-0.141	-0.104	-0.044	-0.006	-0.003	-0.009
Rastrigin	-0.072	-0.059	-0.061	-0.016	-0.001	-0.001	-0.005
Rosenbrock	-0.021	-0.039	-0.034	-0.013	-0.001	0	-0.002
Weierstrass	-0.004	-0.005	-0.005	-0.002	0	0	0
(f) Lbest GCPSO							
Spherical	-0.146	-0.141	-0.084	-0.023	-0.002	0	0
Rastrigin	-0.064	-0.064	-0.034	-0.001	0	0	0
Rosenbrock	-0.021	-0.035	-0.023	-0.004	0	0	0
Weierstrass	-0.004	-0.005	-0.003	0	0	0	0
(g) BBPSO							
Spherical	-0.483	-0.428	-0.146	-0.054	0	0	0
Rastrigin	-0.171	-0.147	-0.076	-0.028	0	0	0
Rosenbrock	-0.061	-0.124	-0.034	-0.005	0	0	0
Weierstrass	-0.041	-0.006	0	-0.068	-5.359	-0.131	-4.953
(h) MBBPSO							
Spherical	-0.269	-0.286	-0.136	-0.068	-0.016	0	0
Rastrigin	-0.111	-0.087	-0.054	-0.036	-0.002	0	0
Rosenbrock	-0.04	-0.066	-0.039	-0.017	0	0	0
Weierstrass	-0.022	-0.001	-0.056	-0.269	-1.029	-5.071	-4.454
(i) SPSO							
Spherical	-0.217	-0.257	-0.37	-0.514	-0.933	-2.271	-3.128
Rastrigin	-0.187	-0.239	-0.409	-0.604	-0.96	-2.223	-3.098

Continued on next page ...

Table 3.8 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Rosenbrock	-0.072	-0.1	-0.156	-0.215	-0.366	-0.899	-1.249
Weierstrass	-0.02	-0.024	-0.039	-0.062	-0.098	-0.218	-0.302

Non-negative DRoC measurements were observed for benchmark functions with dimensionalities of 100 and higher for nearly all PSOs, with the exception of the BBPSO and the MBBPSO on the Weierstrass function, and the SPSO on all functions. This is because the particles of a PSO are more likely to roam outside the bounds of the search space for larger-dimensional problems [46]. Many of the unexpected results in Table 3.7 can thus be attributed to the fact that the PSOs did not transition to exploitative behaviour.

Excluding all cases where one or both PSOs yielded non-negative DRoC measurements, some unexpected results still remain. Figure 3.9 shows the average diversities for some of these unexpected cases. As the figure demonstrates, there were additional cases where the DRoC measurements were negative numbers, but the PSOs failed to transition to exploitation. Notably, for all of these examples, the PSO that did not reduce its diversity was either the BBPSO or the MBBPSO, and the benchmark function was the Weierstrass function.

For most of the remaining cases, the rank results were unexpected but correct, given the actual behaviour of the PSOs in question. Figure 3.10 shows the average diversities of each pair of PSOs that were correctly ranked as 1, despite contrary expectations. Similarly, Figures 3.11 and 3.12 show the average diversities of PSO pairs that were correctly ranked as 0 and -1, respectively.

The remaining cases are shown in Figure 3.13. Notably, each of Figures 3.13a–j show that both PSOs spent all or most of their 2000-iteration runs reducing their diversities. As discussed in Section 3.3.2, a limitation of the DRoC measure is that it may fail to provide a good estimate of the iteration at which a PSO transitions from exploration to exploitation if the PSO continually decreases its diversity during the iterations taken into account. Figure 3.13k shows that the BBPSO increased its diversity after it has transitioned from exploration to exploitation.

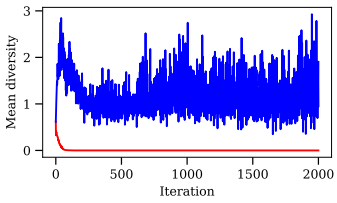

(a) Gbest PSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

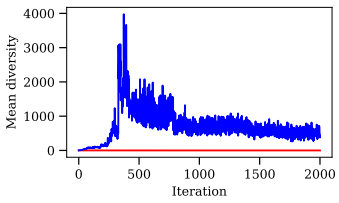

(b) Gbest PSO (red) and BBPSO (blue), Weierstrass in 25 dimensions

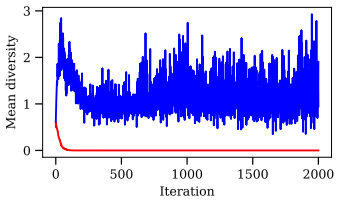

(c) VNPSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

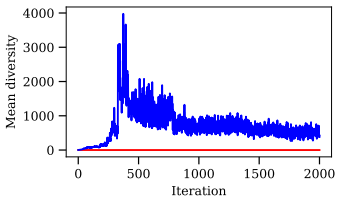

(d) VNPSO (red) and BBPSO (blue), Weierstrass in 25 dimensions

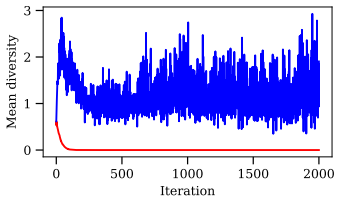

(e) Lbest PSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

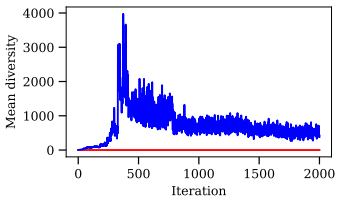

(f) Lbest PSO (red) and BBPSO (blue), Weierstrass in 25 dimensions

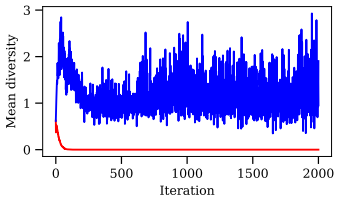

(g) Gbest GCPSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

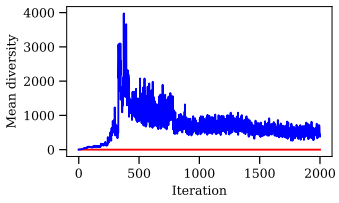

(h) Gbest GCPSO (red) and BBPSO (blue), Weierstrass in 25 dimensions

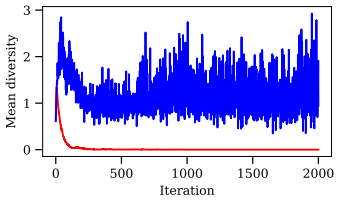

(i) VNGCPSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

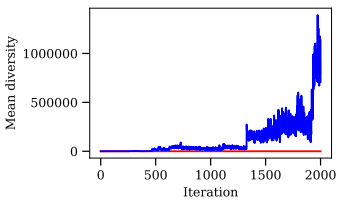

(j) VNGCPSO (red) and BBPSO (blue), Weierstrass in 50 dimensions

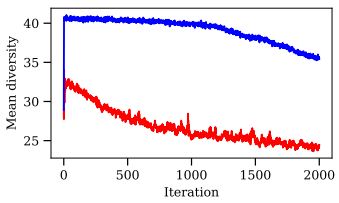

(k) Lbest GCPSO (red) and BBPSO (blue), Spherical in 100 dimensions

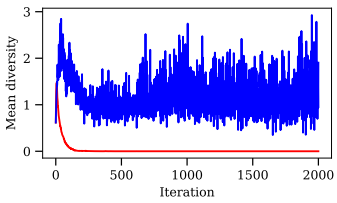

(l) Lbest GCPSO (red) and BBPSO (blue), Weierstrass in 5 dimensions

Figure 3.9: Mean diversities for PSO pairs where one or both PSOs did not reduce diversity
(Continued on next page)

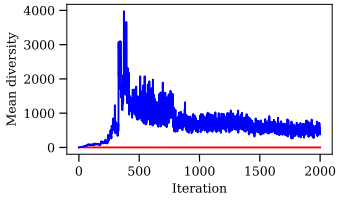

(m) Lbest GCPSO (red) and BBPSO (blue), Weierstrass in 25 dimensions

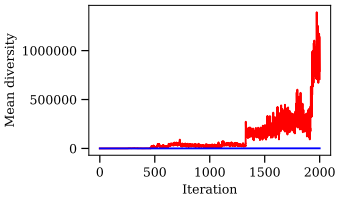

(n) BBPSO (red) and MBBPSO (blue), Weierstrass in 50 dimensions

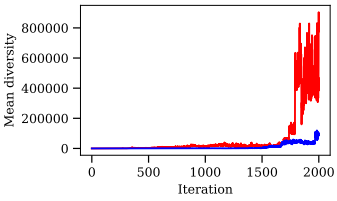

(o) BBPSO (red) and MBBPSO (blue), Weierstrass in 100 dimensions

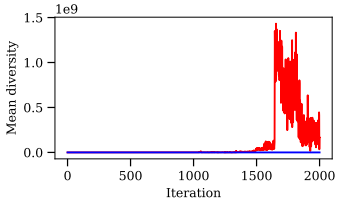

(p) BBPSO (red) and MBBPSO (blue), Weierstrass in 500 dimensions

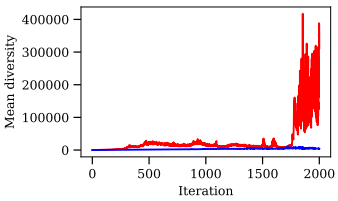

(q) BBPSO (red) and MBBPSO (blue), Weierstrass in 1000 dimensions

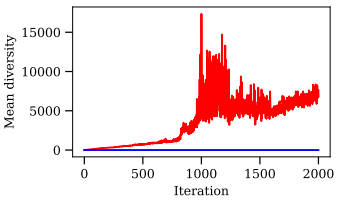

(r) MBBPSO (red) and SPSO (blue), Weierstrass in 500 dimensions

Figure 3.9: (Continued from previous page) Mean diversities for PSO pairs where one or both PSOs did not reduce diversity

3.4.3 Summary

This section investigated the robustness of the DRoC measure with regards to the dimensionality of the search space. It was found that, when used in search spaces with relatively high dimensionalities, the PSOs often did not exhibit the expected search behaviour, where the swarm's diversity would initially be relatively high, followed by a period where the diversity is reduced as the swarm transitions from exploration to exploitation, and finally settling in an exploitative state with a relatively low diversity. Instead, in high-dimensional search spaces, the PSOs usually increased their diversities for the duration of the experiment (i.e. 2000 iterations), and failed to begin reducing their diversities. In these cases, the assumption about the general search behaviour that the DRoC measure is based on did not hold; this was found to be the case with dimensionalities of 100 and higher in most cases. Consequently, when comparing pairs

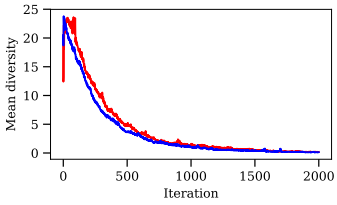

(a) Gbest PSO (red) and VNGCPSO (blue), Spherical, 50 dimensions

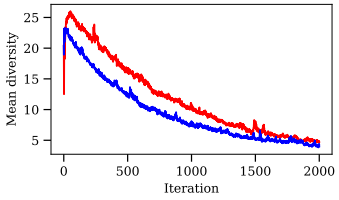

(b) Gbest PSO (red) and VNGCPSO (blue), Rastrigin, 50 dimensions

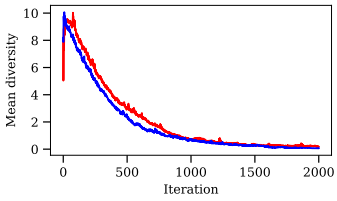

(c) Gbest PSO (red) and VNGCPSO (blue), Rosenbrock, 50 dimensions

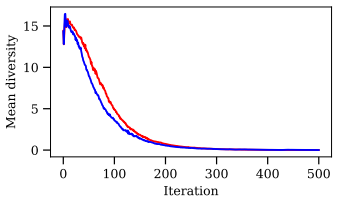

(d) VNPSO (red) and VNGCPSO (blue), Spherical, 25 dimensions

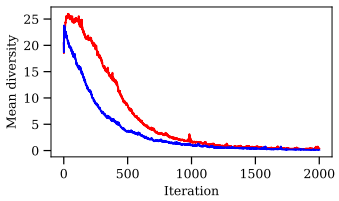

(e) VNPSO (red) and VNGCPSO (blue), Spherical, 50 dimensions

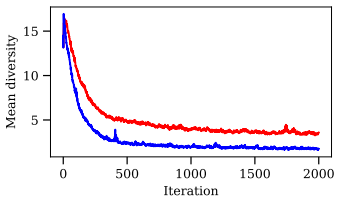

(f) VNPSO (red) and VNGCPSO (blue), Rastrigin, 25 dimensions

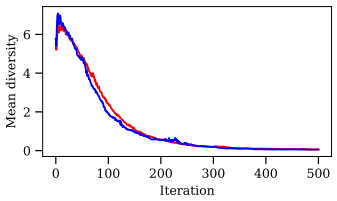

(g) VNPSO (red) and VNGCPSO (blue), Rosenbrock, 25 dimensions

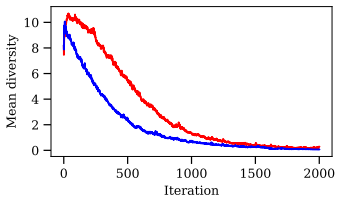

(h) VNPSO (red) and VNGCPSO (blue), Rosenbrock, 50 dimensions

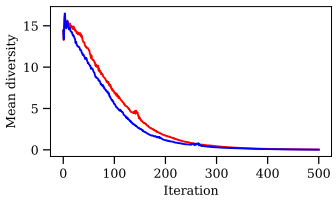

(i) Lbest PSO (red) and Lbest GCPSO (blue), Spherical, 25 dimensions

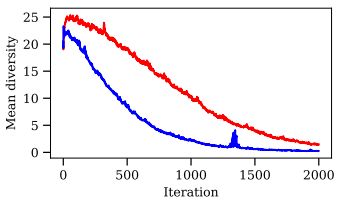

(j) Lbest PSO (red) and Lbest GCPSO (blue), Spherical, 50 dimensions

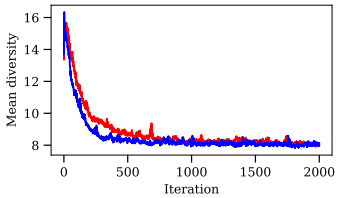

(k) Lbest PSO (red) and Lbest GCPSO (blue), Rastrigin, 25 dimensions

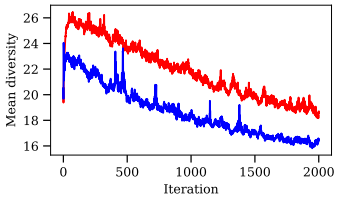

(l) Lbest PSO (red) and Lbest GCPSO (blue), Rastrigin, 50 dimensions

Figure 3.10: Mean diversities for PSO pairs that were correctly ranked as 1, despite contrary expectations (Continued on next page)

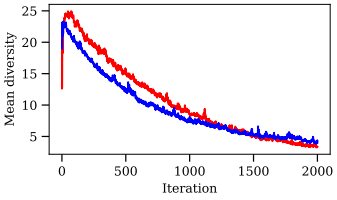

(m) Gbest GCPSO (red) and VNGCPSO (blue), Rastrigin, 50 dimensions

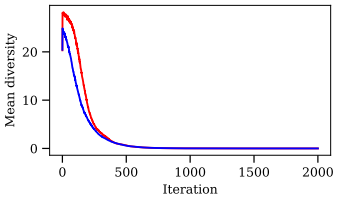

(n) BBPSO (red) and MBBPSO (blue), Spherical, 50 dimensions

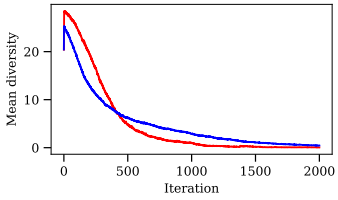

(o) BBPSO (red) and MBBPSO (blue), Rastrigin, 50 dimensions

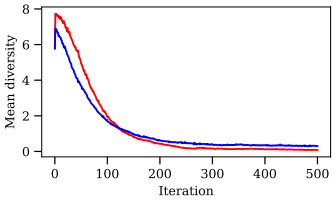

(p) BBPSO (red) and MBBPSO (blue), Rosenbrock, 25 dimensions

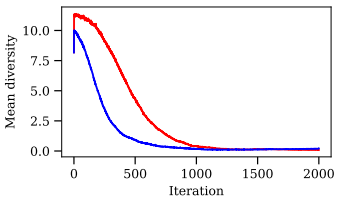

(q) BBPSO (red) and MBBPSO (blue), Rosenbrock, 50 dimensions

Figure 3.10: (Continued from previous page) Mean diversities for PSO pairs that were correctly ranked as 1, despite contrary expectations

of PSOs against one another in such cases, the rates at which the PSOs reduced their diversities often contradicted the expectations as set out in Section 3.2.1. Regardless, non-negative DRoC measurements were found to indicate that the PSOs did not reduce their diversities; thus the DRoC measure still proved to be a reliable indicator of PSO search behaviour.

For cases where the PSOs did reduce their diversities in the expected manner, the DRoC measure was found to be a good indicator of the search behaviour of the PSOs, regardless of the dimensionality of the search space.

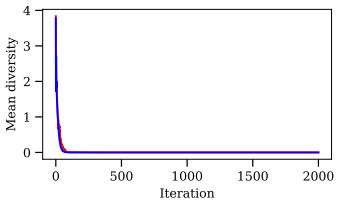

(a) Gbest PSO (red) and SPSO (blue), Spherical, 2 dimensions

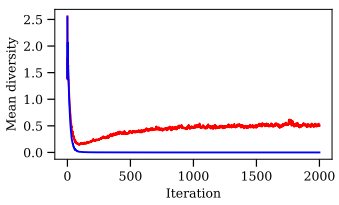

(b) Gbest PSO (red) and SPSO (blue), Rosenbrock, 5 dimensions

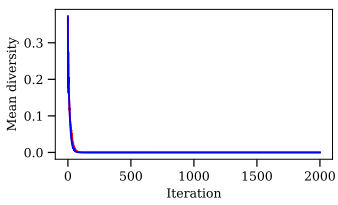

(c) Gbest PSO (red) and SPSO (blue), Weierstrass, 2 dimensions

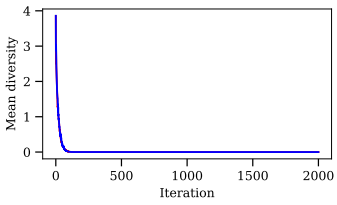

(d) VNPSO (red) and Lbest PSO (blue), Spherical, 2 dimensions

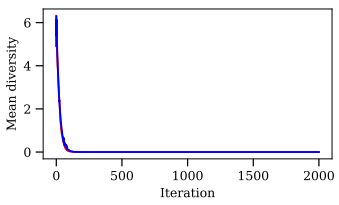

(e) VNPSO (red) and Lbest PSO (blue), Spherical, 5 dimensions

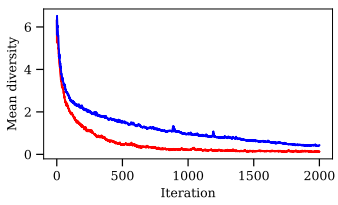

(f) VNPSO (red) and Lbest PSO (blue), Rastrigin, 5 dimensions

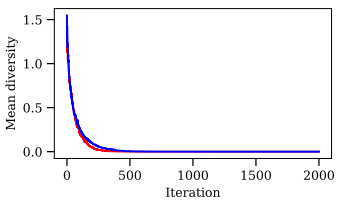

(g) VNPSO (red) and Lbest PSO (blue), Rosenbrock, 2 dimensions

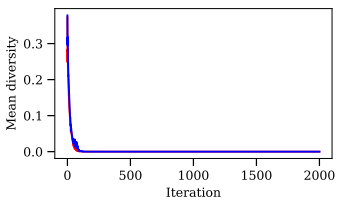

(h) VNPSO (red) and Lbest PSO (blue), Weierstrass, 2 dimensions

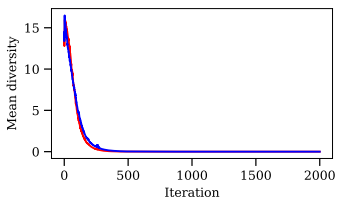

(i) VNPSO (red) and Lbest GCPSO (blue), Spherical, 25 dimensions

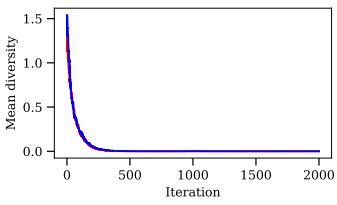

(j) VNPSO (red) and Lbest GCPSO (blue), Rosenbrock, 2 dimensions

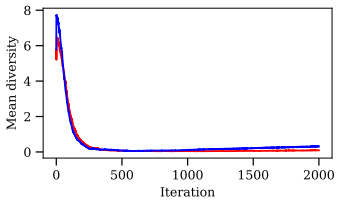

(k) VNPSO (red) and BBPSO (blue), Rosenbrock, 25 dimensions

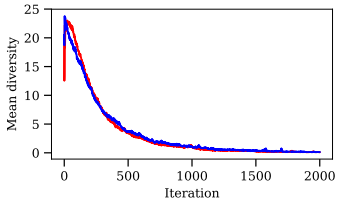

(l) Gbest GCPSO (red) and VNGCPSO (blue), Spherical, 50 dimensions

Figure 3.11: Mean diversities for PSO pairs that were correctly ranked as 0, despite contrary expectations (Continued on next page)

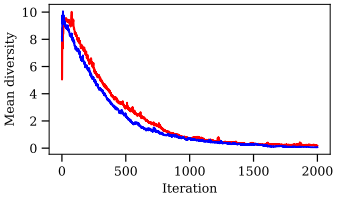

(m) Gbest PSO (red) and VNGCPSO (blue), Rosenbrock, 50 dimensions

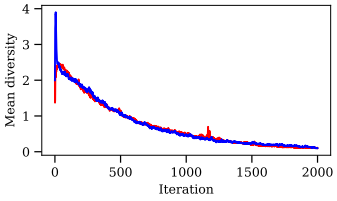

(n) Gbest GCPSO (red) and VNGCPSO (blue), Weierstrass, 50 dimensions

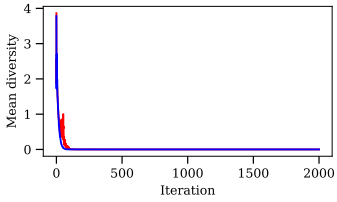

(o) Gbest GCPSO (red) and SPSO (blue), Spherical, 2 dimensions

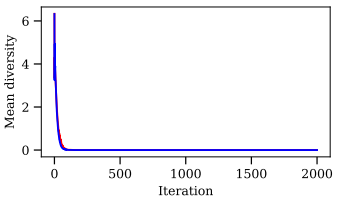

(p) Gbest GCPSO (red) and SPSO (blue), Spherical, 5 dimensions

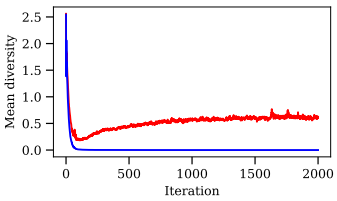

(q) Gbest GCPSO (red) and SPSO (blue), Rosenbrock, 5 dimensions

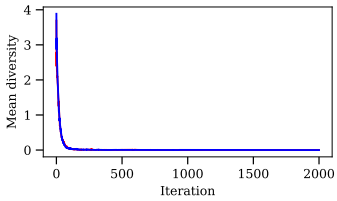

(r) VNGCPSO (red) and Lbest GCPSO (blue), Spherical, 2 dimensions

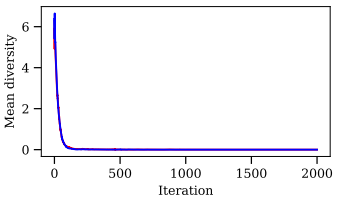

(s) VNGCPSO (red) and Lbest GCPSO (blue), Spherical, 5 dimensions

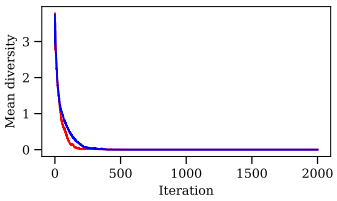

(t) VNGCPSO (red) and Lbest GCPSO (blue), Rastrigin, 2 dimensions

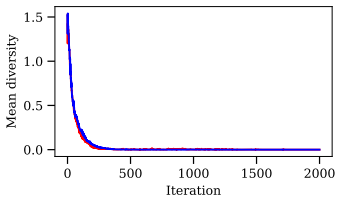

(u) VNGCPSO (red) and Lbest GCPSO (blue), Rosenbrock, 2 dimensions

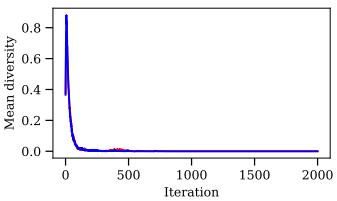

(v) VNGCPSO (red) and Lbest GCPSO (blue), Weierstrass, 2 dimensions

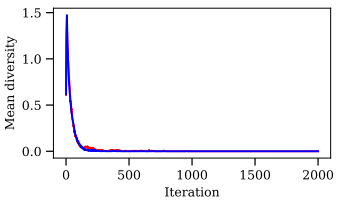

(w) VNGCPSO (red) and Lbest GCPSO (blue), Weierstrass, 5 dimensions

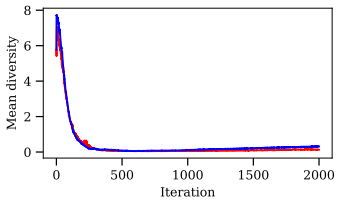

(x) VNGCPSO (red) and BBPSO (blue), Rosenbrock, 25 dimensions

Figure 3.11: (Continued from previous page) Mean diversities for PSO pairs that were correctly ranked as 0, despite contrary expectations

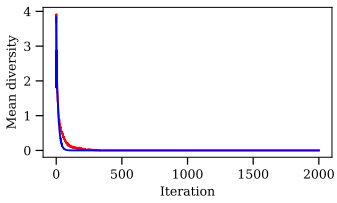

(y) BBPSO (red) and SPSO (blue), Rastrigin, 2 dimensions

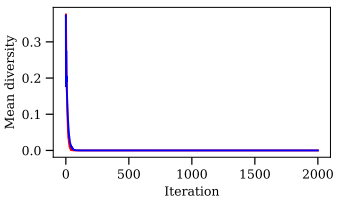

(z) MBBPSO (red) and SPSO (blue), Weierstrass, 2 dimensions

Figure 3.11: (Continued from previous page) Mean diversities for PSO pairs that were correctly ranked as 0, despite contrary expectations

3.5 DRoC Robustness with Regards to Number of Diversity Measurements

The DRoC measure is obtained from a set of instantaneous diversity measurements. The number of diversity measurements that are taken into account for a DRoC measurement may therefore influence the accuracy of the DRoC measure. On the one hand, using too few diversity measurements may provide the measure with insufficient data to produce an accurate DRoC measurement. On the other hand, with more diversity measurements used for a DRoC measurement, a larger proportion of the measurements will represent the PSO's exploitative state, and a smaller proportion of the diversity measurements will represent its explorative state. This may result in a two-piecewise linear approximation that over-represents the exploitative aspect of the PSO, and under-represents its explorative behaviour.

Similar to Sections 3.3 and 3.4, this section investigates the robustness of the DRoC measure with regards to the number of diversity measurements taken into account. The experimental procedure is given in Section 3.5.1. The results are listed and discussed in Section 3.5.2. The findings of the section are summarised in Section 3.5.3.

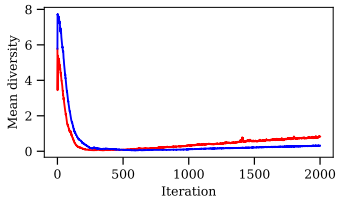

(a) Gbest PSO (red) and BBPSO (blue), Rosenbrock, 25 dimensions

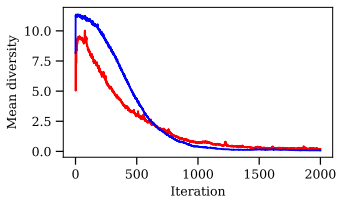

(b) Gbest PSO (red) and BBPSO (blue), Rosenbrock, 50 dimensions

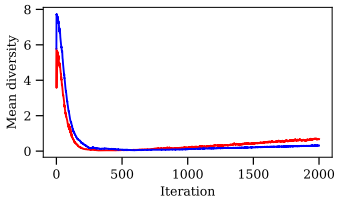

(c) Gbest GCPSO (red) and BBPSO (blue), Rosenbrock, 25 dimensions

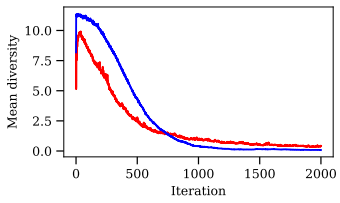

(d) Gbest GCPSO (red) and BBPSO (blue), Rosenbrock, 50 dimensions

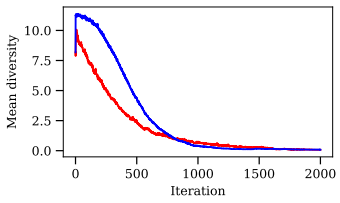

(e) VNGCPSO (red) and BBPSO (blue), Rosenbrock, 50 dimensions

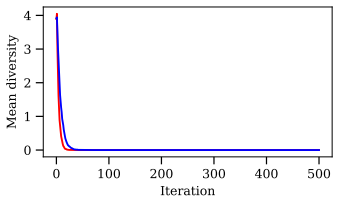

(f) BBPSO (red) and MBBPSO (blue), Spherical, 2 dimensions

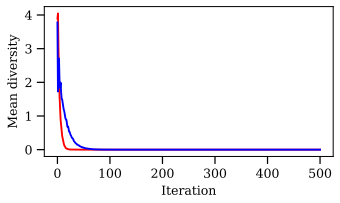

(g) BBPSO (red) and SPSO (blue), Spherical, 2 dimensions

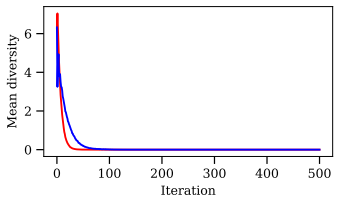

(h) BBPSO (red) and SPSO (blue), Spherical, 5 dimensions

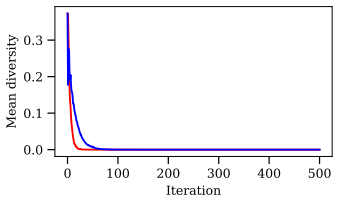

(i) BBPSO (red) and SPSO (blue), Weierstrass, 2 dimensions

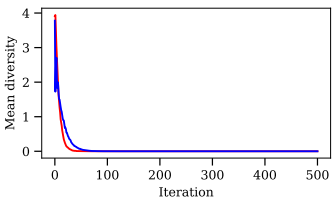

(j) MBBPSO (red) and SPSO (blue), Spherical, 2 dimensions

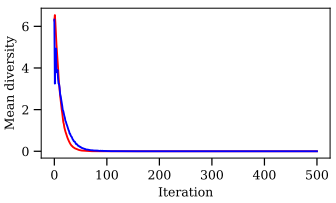

(k) MBBPSO (red) and SPSO (blue), Spherical, 5 dimensions

Figure 3.12: Mean diversities for PSO pairs that were correctly ranked as -1, despite contrary expectations

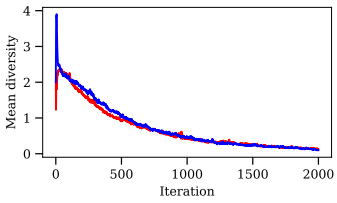

(a) Gbest PSO (red) and VNGCPSO (blue), Weierstrass, 50 dimensions, ranked 0

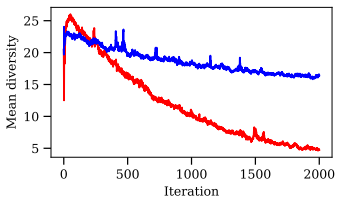

(b) Gbest PSO (red) and Lbest GCPSO (blue), Rastrigin, 50 dimensions, ranked 0

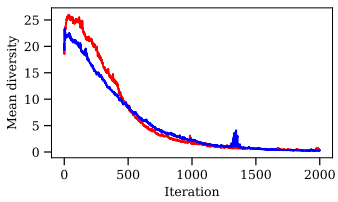

(c) VNPSO (red) and Lbest GCPSO (blue), Spherical, 50 dimensions, ranked 0

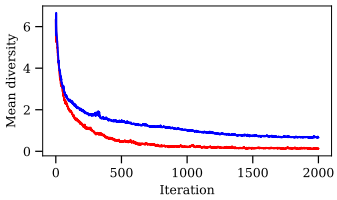

(d) VNPSO (red) and Lbest GCPSO (blue), Rastrigin, 5 dimensions, ranked 0

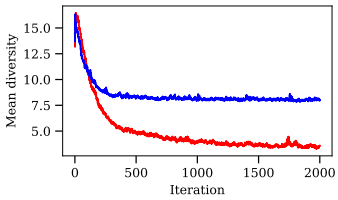

(e) VNPSO (red) and Lbest GCPSO (blue), Rastrigin, 25 dimensions, ranked 0

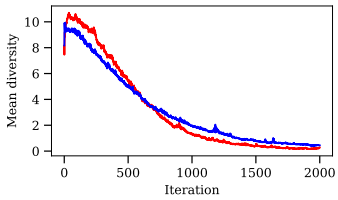

(f) VNPSO (red) and Lbest GCPSO (blue), Rosenbrock, 50 dimensions, ranked 0

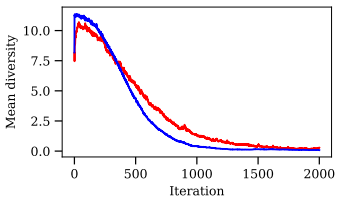

(g) VNPSO (red) and BBPSO (blue), Rosenbrock, 50 dimensions, ranked 0

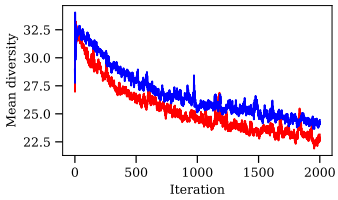

(h) VNGCPSO (red) and Lbest GCPSO (blue), Spherical, 100 dimensions, ranked 0

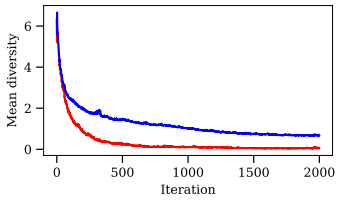

(i) VNGCPSO (red) and Lbest GCPSO (blue), Rastrigin, 5 dimensions, ranked 0

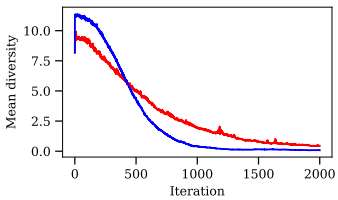

(j) Lbest GCPSO (red) and BBPSO (blue), Rosenbrock, 50 dimensions, ranked 0

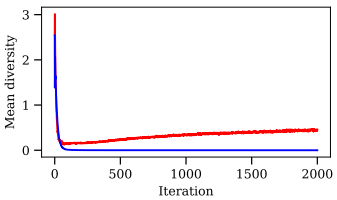

(k) BBPSO (red) and SPSO (blue), Rosenbrock, 5 dimensions, ranked -1

Figure 3.13: Mean diversities for PSO pairs where DRoC measurements did not compare as expected in various dimensionalities

3.5.1 Experimental Procedure

In order to test the DRoC measure's robustness with regards to the number of diversity measurements used, the following approach is used. Experimental configurations are generated for each PSO listed in Table 3.2, each benchmark function listed in Table 3.5, and for each of 25, 50, 75, 100, 125, 2000 and 10000 diversity measurements. For each configuration, each PSO is initialised with 25 particles at random starting positions within the search space bounds of the benchmark function in 25 dimensions. The PSO is then run for the number of iterations specified by the configuration, with a diversity measurement taken at each iteration. From these diversity measurements, a DRoC measurement is determined for the PSO. Each configuration is run 30 times.

For each PSO and benchmark function, there are multiple configurations for the various numbers of diversity measurements used. The average DRoC measurements for each of these configurations are then compared to one another in order to assess the stability of the DRoC measure when used with different numbers of diversity measurements.

3.5.2 Results

For each set of configurations that use the same PSO and benchmark functions, Figure 3.14 shows the DRoC measurements obtained with different numbers of diversity measurements. These DRoC measurements are shown over the diversity measurements at each iteration (the diversity measurements were scaled to the range of the DRoC measurements on the y-axis for easy comparison). The iterations at which the DRoC measurements were taken are indicated by vertical lines.

In Figures 3.14a–r, ab, and af–aj, when small numbers of diversity measurements are used, the DRoC measurements are relatively low. The DRoC measurements stabilise once enough diversity measurements are used. The number of iterations at which DRoC measurements stabilise is case-specific, but appears to be related to the number of iterations needed for the PSO to transition from exploration to exploitation. For example, in Figure 3.14b, the diversities of the gbest PSO stabilised after approximately 500 iterations when the PSO was solving the Rastrigin function; the DRoC measurements show a sharp increase up to approximately the same number of iterations, after which they re-

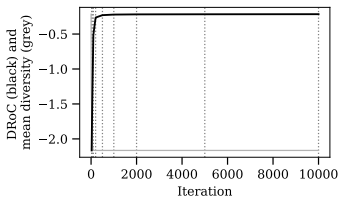

(a) Gbest PSO, Spherical

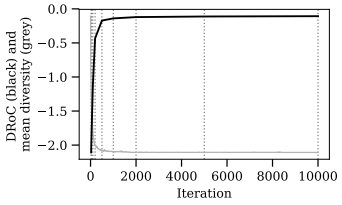

(b) Gbest PSO, Rastrigin

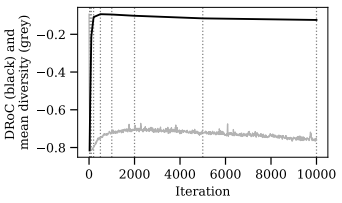

(c) Gbest PSO, Rosenbrock

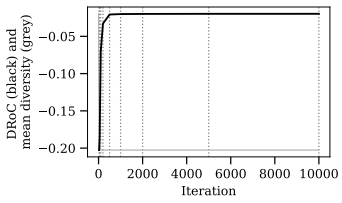

(d) Gbest PSO, Weierstrass

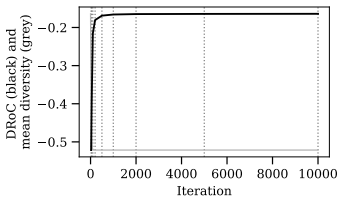

(e) VNPSO, Spherical

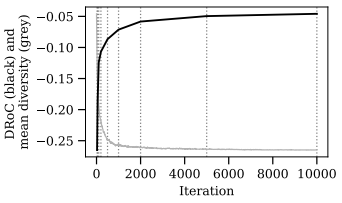

(f) VNPSO, Rastrigin

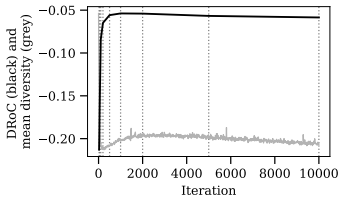

(g) VNPSO, Rosenbrock

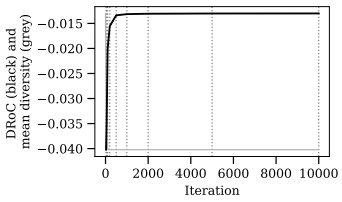

(h) VNPSO, Weierstrass

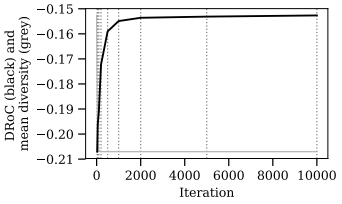

(i) Lbest PSO, Spherical

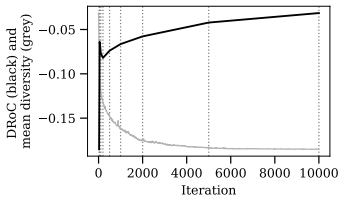

(j) Lbest PSO, Rastrigin

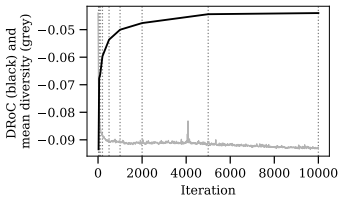

(k) Lbest PSO, Rosenbrock

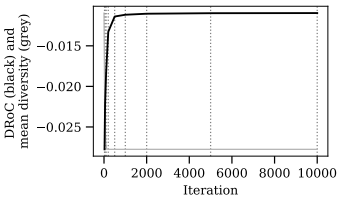

(l) Lbest PSO, Weierstrass

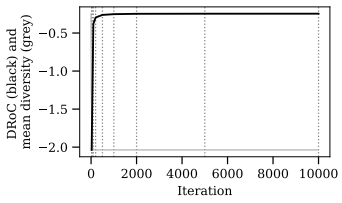

(m) Gbest GCPSO, Spherical

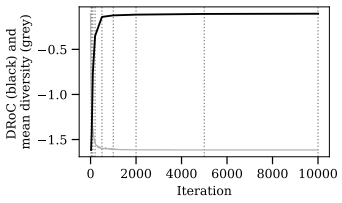

(n) Gbest GCPSO, Rastrigin

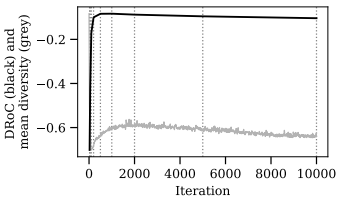

(o) Gbest GCPSO, Rosenbrock

Figure 3.14: DRoCs calculated with various numbers of iterations (black), shown over diversity measurements at each iteration (grey; not to scale) (Continued on next page)

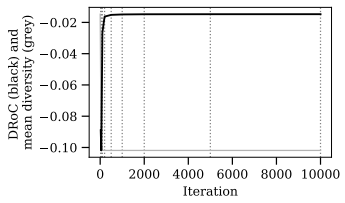

(p) Gbest GCPSO, Weierstrass

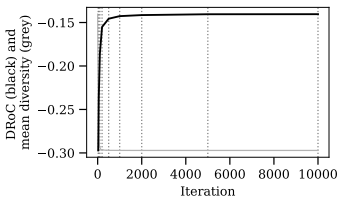

(q) VNGCPSO, Spherical

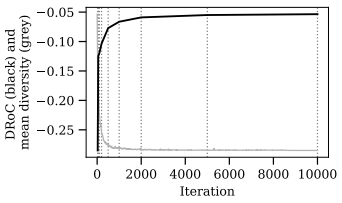

(r) VNGCPSO, Rastrigin

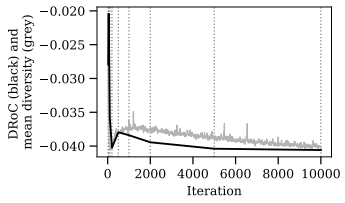

(s) VNGCPSO, Rosenbrock

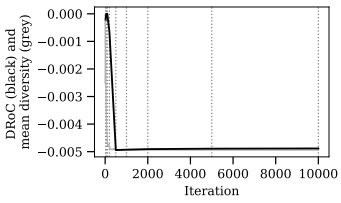

(t) VNGCPSO, Weierstrass

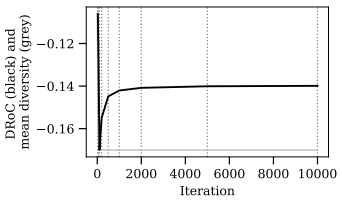

(u) Lbest GCPSO, Spherical

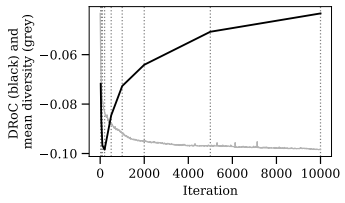

(v) Lbest GCPSO, Rastrigin

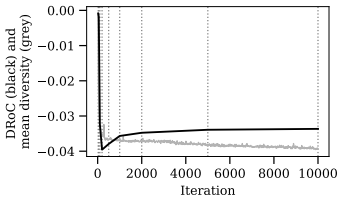

(w) Lbest GCPSO, Rosenbrock

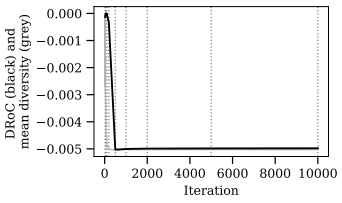

(x) Lbest GCPSO, Weierstrass

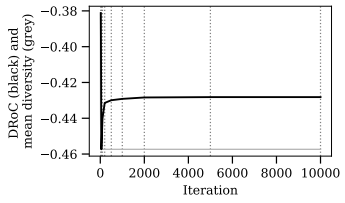

(y) BBPSO, Spherical

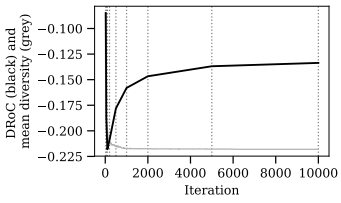

(z) BBPSO, Rastrigin

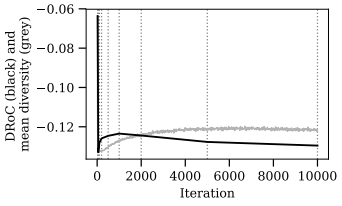

(aa) BBPSO, Rosenbrock

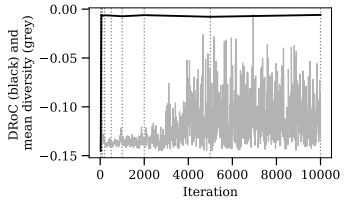

(ab) BBPSO, Weierstrass

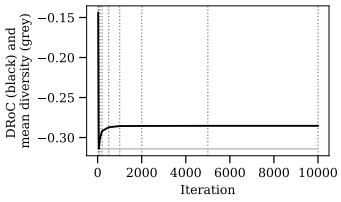

(ac) MBBPSO, Spherical

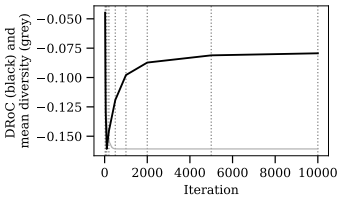

(ad) MBBPSO, Rastrigin

Figure 3.14: (Continued from previous page) DRoCs calculated with various numbers of iterations (black), shown over diversity measurements at each iteration (grey; not to scale)

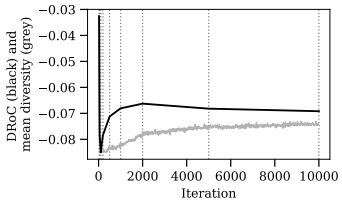

(ae) MBBPSO, Rosenbrock

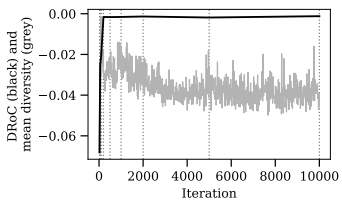

(af) MBBPSO, Weierstrass

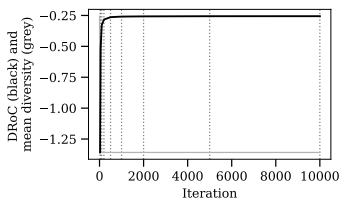

(ag) SPSO, Spherical

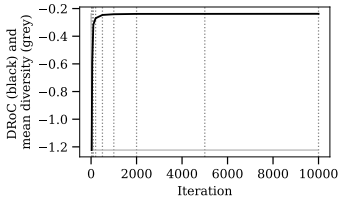

(ah) SPSO, Rastrigin

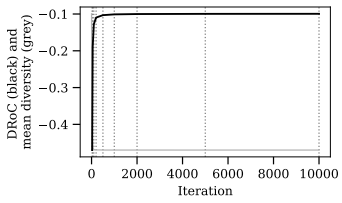

(ai) SPSO, Rosenbrock

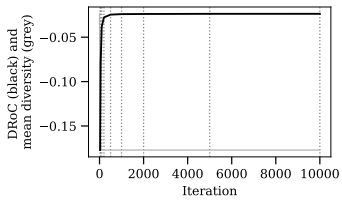

(aj) SPSO, Weierstrass

Figure 3.14: (Continued from previous page) DRoCs calculated with various numbers of iterations (black), shown over diversity measurements at each iteration (grey; not to scale)

main fairly stable at a higher measurement value. This relationship between the stability of the DRoC measurements and that of the diversity measurements is to be expected, considering that the DRoC measurement is obtained by fitting a two-piecewise linear approximation to the diversity measurements of a PSO; when fitted to a PSO which has not yet finished its transition from exploration to exploitation, the fitted lines can not reflect the rate at which the transition is completed. Thus, the DRoC measure is not robust when used with too few diversity measurements. The measure might fail to indicate the behaviour of a PSO when given fewer diversity measurements than required by the PSO to complete its transition to exploitation.

In Figures 3.14s–aa and ac–ae, rather than being relatively low, the DRoC measurements are relatively high when used with too few diversity measurements. Nonetheless, in most of these cases, the DRoC measurements stabilise once enough diversity measurements are taken into account, as with the other cases previously discussed.

In the case of Figure 3.14v, the DRoC measurements did not completely stabilise within 10000 iterations. This might simply be because the lbest PSO fails to completely reduce its diversity by the end of those iterations when solving the Rastrigin function.

Note that, in Figures 3.14ab and af, the BBPSO and the MBBPSO did not decrease their diversities over time. The DRoC measurements for these cases quickly stabilised to 0, indicating that the PSOs did not transition to exploitation.

3.5.3 Summary

This section investigated the robustness of the DRoC measure with regards to the number of diversity measurements taken into account. The DRoC measure was found to fail to indicate the search behaviour of PSOs when too few diversity measurements were used; specifically, when the PSO was still transitioning from exploration to exploitation at the iteration of the last diversity measurement taken into account, the diversity measurements provided insufficient information to result in an accurate DRoC measurement. The minimum number of iterations that should be taken into account is thus situation-specific: each PSO used with each benchmark function may take a different number of iterations to reach its exploitative phase. For the cases studied in this section, the diversity measurements from the first 500 iterations were usually sufficient.

3.6 Conclusions

This chapter described the diversity rate-of-change (DRoC) measure for quantifying the search behaviour of a PSO, and investigated its robustness with regards to three variables, namely the swarm size of the PSO, the dimensionality of the search space, and the number of diversity measurements taken into account, or equivalently, the number of iterations that the PSO was run for.

The measure was found to be robust with regards to all swarm sizes. However, two limitations of the DRoC measure were discovered. Firstly, because the DRoC measure is based on the slope of the first line of a two-piecewise approximation, it is unable to differentiate PSOs based on the iteration at which they transition from exploration to exploitation if the PSOs decrease their diversities at the same rate. Secondly, if a PSO exhibits a very smooth transition from exploration to exploitation, the DRoC measure may interpret the transition period as the exploitative phase and indicate that the transition has occurred too early.

The measure was also found to be robust with regards to dimensionality. Instances were found where, when used in high-dimensional search spaces, PSOs did not reduce their diversity as expected. This resulted in DRoC measurements that contradicted expectations. However, the DRoC measurements did succeed in representing the actual behaviour of the PSOs under these circumstances by yielding non-negative values.

Lastly, with regards to the number of diversity measurements taken into account, the DRoC measure was found to fail when informed by too few diversity measurements; 500 measurements were found to be sufficient in most cases.

When used with a sufficient number of diversity measurements, the DRoC measure can be used to quantify the behaviour of PSO algorithms in terms of exploration and exploitation. The next chapter compares DRoC measurements among PSOs solving benchmark functions exhibiting different fitness landscape characteristics in order to determine whether such characteristics predictably influence the behaviour of those PSOs.

Chapter 4

Linking Fitness Landscape Characteristics to PSO Search Behaviour

This chapter investigates the influence of the characteristics of fitness landscapes on the search behaviour of the PSOs that solve them. The previous chapter has shown that the DRoC measure robustly quantifies PSO search behaviour in terms of exploration and exploitation. In order to determine the influence of FLCs on PSO search behaviour, DRoC measurements will be compared to individual FLC measurements for PSOs solving a variety of benchmark functions that exhibit a variety of FLCs.

Section 4.1 presents a modification to the dispersion metric which improves its capability to indicate the global landscape structure of a fitness landscape. Section 4.2 investigates correlations between individual FLCs and PSO search behaviour. Conclusions are summarised in Section 4.3.

4.1 Modified Dispersion Metric

The dispersion metric (DM) introduced in Section 2.4.4 depends on two parameters, namely the size of the initial sample, n_a, and the size of the subsample of fittest solutions, n_b.

The result of a DM measurement is interpreted to indicate whether a landscape has a single funnel or multiple funnels. A positive value indicates that the landscape has multiple funnels, and a negative value that it has a single funnel. However, even for benchmark functions that are known to have multiple funnels, some DM measurements may incorrectly indicate that the function is single-funnelled.

Figure 4.1a plots multiple independent DM measurements taken on the Schwefel 2.26 benchmark function in two dimensions with various values for n_b. Schwefel 2.26 is known to have multiple funnels, so the DM measurements are expected to be positive values. However, as the figure illustrates, some negative measurements are found at virtually all values for n_b. For any choice of parameters, any single measurement could incorrectly indicate that the function has a single funnel. Similarly, Figure 4.1b shows that negative DM values are found for many n_b values taken on the Rana benchmark function in two dimensions, which is also known to be multi-funnelled.

A more robust indicator of multiple funnels may be a well-chosen statistically representative value from a set of independently sampled DM measurements for varying n_b values. The following metric is therefore proposed as an improved indicator of funnels in a landscape:

$$\mathrm{DM}_{med} = \mathrm{median}\big(\mathrm{norm}(\mathrm{disp}(S^*_{n_b})) - \mathrm{disp}_D\big)\ \forall n_b \in \{2, 3, ..., \frac{n_a}{2}\}. \tag{4.1}$$

In other words, DM_{med} is the median of the sequence of DM values obtained for subsample sizes ranging from 2 to half of the initial sample size[1].

Table 4.1 compares DM measurements to DM_{med} measurements for some functions in two dimensions that are known to be either single-funnelled or multi-funnelled. Specifically, for single-funnelled functions, the maximum DM and DM_{med} values obtained for 30 independent measurements are given, in order to display incorrect measurements that indicate that the function is multi-funnelled when it is not, if any such measurements were encountered. Similarly, for multi-funnelled functions, the minimum values for 30 independent measurements are given, in order to display any measurements that incorrectly indicate that a function has a single funnel. Incorrect values are shaded.

[1] n_b values are incremented starting from 2, because using a subsample size of 0 would result in an undefined DM value, and using a subsample size of 1 would result in a DM value of 0.

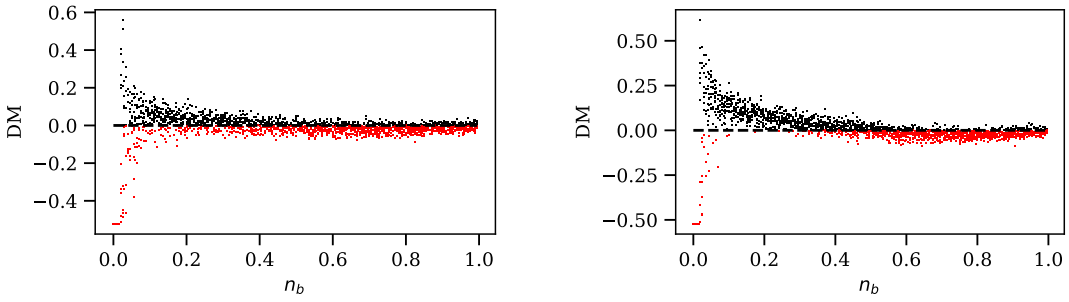

(a) DM measurements taken on Schwefel 2.26 with various values for n_b

(b) DM measurements taken on Rana with various values for n_b

Figure 4.1: DM measurements taken on benchmark functions with multiple funnels; incorrect measurements (indicating that the landscape has a single funnel) are shown in red.

The results illustrate that DM_{med} does not provide the incorrect measurements that DM sometimes does. Therefore, the DM_{med} metric will be used to measure the global landscape structure of benchmark functions in this chapter.

4.2 Linking Fitness Landscape Characteristics to Particle Swarm Optimiser Search Behaviour

This section investigates links between single FLCs and the search behaviour of PSOs. The experimental procedure used is discussed in Section 4.2.1. Results are listed in Section 4.2.2. The findings are summarised in Section 4.2.3.

4.2.1 Experimental Procedure

The following procedure was used to study links between FLCs and PSO search behaviour in the context of the DRoC measure. Various FLCs were measured for each function in a suite of benchmark functions. DRoC measurements were taken for a set of PSOs solving those benchmark functions. For each PSO, the influence of each measured FLC on the search behaviour of that PSO was investigated by comparing the FLC measurements taken on the various benchmark functions against the PSO's DRoC measurements for

Table 4.1: Worst DM and DM_{med} measurements of 30 independent runs on various benchmark functions

Benchmark function	Underlying modality	Worst DM	Worst DM_{med}
Ackley	Unimodal	-0.266	-0.227
Beale	Unimodal	-0.122	-0.169
Bohachevksy	Unimodal	-0.277	-0.241
Griewank	Unimodal	-0.288	-0.247
Levy 13	Unimodal	-0.301	-0.240
Michalewicz	Unimodal	-0.013	-0.092
Quadric	Unimodal	-0.254	-0.202
Quartic	Unimodal	-0.281	-0.245
Rana	Multimodal	-0.020	0.050
Rastrigin	Unimodal	-0.141	-0.152
Rosenbrock	Unimodal	-0.140	-0.167
Salomon	Unimodal	-0.277	-0.239
Schwefel 2.22	Unimodal	-0.223	-0.200
Schwefel 2.26	Multimodal	-0.157	0.010
Six Hump Camel Back	Unimodal	-0.284	-0.240
Skew Rastrigin	Unimodal	-0.231	-0.230
Spherical	Unimodal	-0.250	-0.246
Step	Unimodal	-0.278	-0.249
Weierstrass	Unimodal	-0.211	-0.190
Zakharov	Unimodal	-0.223	-0.187

those functions.

The FLC metrics included in this study, along with their parameters, are as follows:

Ruggedness The $FEM_{0.1}$ metric was used with a step size upper bound of $0.1 \times (x_{max} - x_{min})$, and the $FEM_{0.01}$ metric was used with a step size upper bound of $0.01 \times (x_{max} - x_{min})$, where $(x_{max} - x_{min})$ is the domain range of the fitness function. Both metrics were calculated from 1000 steps for each starting zone, with the number of starting zones equal to the dimensionality of the fitness function.

Neutrality The PN and LSN metrics were each used with a step size upper bound of $0.02 \times (x_{max} - x_{min})$, calculated from 200 steps for each starting zone, with the number of starting zones equal to the dimensionality of the fitness function. For the PN and LSN metrics, points in a random walk were considered to be neutral relative to one another when their fitness values were within 1×10^{-8} of one another.

Gradients The G_{avg} and G_{dev} metrics were each used with a fixed step size of $0.001 \times (x_{max} - x_{min}) \times D$, calculated from 1000 steps for each starting zone, with one starting zone used for each dimension of the fitness function.

Global landscape structure The DM_{med} metric was used with n_a, the initial sample size, set to 100 positions, and with n_b, the size of the subsample of fittest positions, ranging from 2 to 50 positions.

Deception The FDC metric was used with a sample size of 1000.

Searchability The FCI_{soc} and FCI_{cog} metrics were each used with a swarm size of 500 particles. The $FCI_{\bar{\sigma}}$ metric was computed from 30 measurements of each of the FCI_{soc} and FCI_{soc} metrics.

These FLC metrics were determined for each benchmark function listed in Table 4.2. Most benchmark functions were used in 5 dimensions, with the exception of certain functions that are only defined in 2 dimensions. Each FLC metric was calculated for each benchmark function 30 times, giving 30 measurements.

Table 3.2 lists the PSOs included in this study. Each of these PSOs were used with a swarm size of 25 particles and run for 2000 iterations, after which a DRoC measurement was determined. Each PSO was run on each benchmark function 30 times, giving 30 DRoC measurements.

4.2.2 Results

For each FLC metric, Table 4.3 shows the Spearman correlation coefficient [56] between the FLC metric values for each benchmark function, and the DRoC values obtained for PSOs on those benchmark functions.

Table 4.2: Benchmark functions studied in this section

Function name	Domain
Ackley [74]	$\mathbf{x} \in [-32, 32]^5$
Alpine [53]	$\mathbf{x} \in [-10, 10]^5$
Beale [44]	$\mathbf{x} \in [-4.5, 4.5]^2$
Bohachevsky [21] (generalised)	$\mathbf{x} \in [-15, 15]^5$
Egg Holder [45]	$\mathbf{x} \in [-512, 512]^2$
Goldstein-Price [74]	$\mathbf{x} \in [-2, 2]^2$
Griewank [74]	$\mathbf{x} \in [-600, 600]^5$
Levy 13 [45] (generalised)	$\mathbf{x} \in [-10, 10]^5$
Michalewicz [44]	$\mathbf{x} \in [0, \pi]^5$
Pathological [53]	$\mathbf{x} \in [-100, 100]^5$
Quadric (Schwefel 1.2) [74]	$\mathbf{x} \in [-100, 100]^5$
Quartic [74]	$\mathbf{x} \in [-1.28, 1.28]^5$
Rana [52]	$\mathbf{x} \in [-512, 512]^5$
Rastrigin [74]	$\mathbf{x} \in [-5.12, 5.12]^5$
Rosenbrock [74] (generalised)	$\mathbf{x} \in [-2.048, 2.048]^5$
Salomon [52]	$\mathbf{x} \in [-100, 100]^5$
Schwefel 2.22 [74]	$\mathbf{x} \in [-10, 10]^5$
Schwefel 2.26 [74]	$\mathbf{x} \in [-500, 500]^5$
Six-hump Camel Back [74]	$\mathbf{x} \in [-5, 5]^2$
Skew Rastrigin [21]	$\mathbf{x} \in [-5, 5]^5$
Spherical [8]	$\mathbf{x} \in [-100, 100]^5$
Step [74]	$\mathbf{x} \in [-20, 20]^5$
Weierstrass [44]	$\mathbf{x} \in [-0.5, 0.5,]^5$
Zakharov [44]	$\mathbf{x} \in [-5, 10]^5$

Two groups of strong correlations between FLC metrics and DRoC measurements are apparent: FCI_{soc} shows a strong negative correlation with the DRoC measurements of all PSOs, peaking for the lbest PSO with a correlation coefficient of -0.507; and $\text{FCI}_{\bar{\sigma}}$ shows a strong positive correlation with the DRoC measurements of all PSOs, peaking

Table 4.3: Spearman's correlation coefficient between FLC metrics and DRoC values

FLC metric	Gbest PSO	VNPSO	Lbest PSO	Gbest GCPSO	VNGCPSO
DM_{med}	0.290	0.352	0.404	0.296	0.309
LSN	0.027	0.034	0.027	0.029	0.032
FDC	-0.308	-0.358	-0.408	-0.311	-0.323
$FCI_{\bar{\sigma}}$	0.472	0.504	0.470	0.466	0.493
$FEM_{0.01}$	-0.291	-0.289	-0.345	-0.285	-0.292
$FEM_{0.1}$	0.069	0.092	0.097	0.059	0.068
G_{avg}	0.042	0.067	0.127	0.035	0.054
FCI_{soc}	-0.392	-0.439	-0.507	-0.409	-0.419
G_{dev}	0.290	0.321	0.384	0.298	0.319
PN	0.025	0.032	0.024	0.026	0.029
FCI_{cog}	-0.090	-0.131	-0.133	-0.087	-0.110

FLC metric	Lbest GCPSO	BBPSO	MBBPSO	SPSO
DM_{med}	0.353	0.355	0.270	0.206
LSN	0.025	0.022	0.036	0.018
FDC	-0.371	-0.355	-0.293	-0.233
$FCI_{\bar{\sigma}}$	0.476	0.503	0.370	0.409
$FEM_{0.01}$	-0.361	-0.280	-0.361	-0.306
$FEM_{0.1}$	0.055	0.070	-0.068	0.010
G_{avg}	0.097	0.047	-0.073	0.002
FCI_{soc}	-0.459	-0.438	-0.388	-0.335
G_{dev}	0.366	0.315	0.245	0.261
PN	0.022	0.020	0.034	0.016
FCI_{cog}	-0.097	-0.124	-0.016	-0.020

for the VNPSO with a correlation coefficient of 0.504. The FDC metric exhibits a moderate negative correlation with the DRoC measurements of the lbest PSO, as well as other PSOs to a lesser extent. Similarly, the DM_{med} exhibits a moderate positive correlation with the DRoC measurements of the lbest PSO. The $FEM_{0.01}$ and G_{dev} metrics also respectively exhibit a moderate negative and positive correlation. The remaining

FLC metrics correlated weakly with the DRoC metric. The correlations between the neutrality metrics, namely PN and LSN, and the DRoC metric were weak, but consistently positive. Very weak correlations were found between the G_{avg} measurements and the DRoC measurements.

Recall that the FCI_{soc} metric indicates the searchability of a landscape with respect to social-only PSO updates, where higher values indicate more searchable landscapes, and that low values for the DRoC measure indicate fast convergence (large negative slope), whereas high DRoC values indicates indicates slow convergence (small negative slope). The strong negative correlation between this metric and the DRoC measurements suggests that more searchable landscapes can be expected to correspond with a faster transition from exploration to exploitation in PSOs. This is to be expected: in the general sense, the searchability of a landscape with regards to PSOs simply refers to the ability of a particles of a PSO algorithm to move to fitter positions. In more searchable landscapes, fitter regions will typically be less dispersed than in less searchable landscapes. As such, the rate at which a PSO reduces its diversity can be expected to be positively influenced by more searchable landscapes.

Similarly, higher $FCI_{\bar{\sigma}}$ values indicate less searchable landscapes, based on the unpredictability of the FCI_{soc} and FCI_{cog} metrics. Therefore, the positive correlation between $FCI_{\bar{\sigma}}$ and DRoC measurements is to be expected for the same reasons that FCI_{soc} and DRoC were found to be correlated negatively: in more searchable landscapes, PSOs are able to converge on fitter regions at a faster rate.

It is notable that, while the FCI_{cog} measurements did also correlate negatively with the DRoC measurements, the correlation was relatively weak compared to the correlation between the FCI_{soc} measurements and the DRoC measurements, with correlation coefficients ranging from -0.133 to only -0.016. The FCI_{soc} and FCI_{cog} metrics both indicate the searchability of a landscape with regards to PSO algorithms, respectively based on the improvements in fitness obtained by social-only and cognitive-only PSOs. In a cognitive-only PSO, the particles of the swarm effectively act as individual hill-climbers, uninfluenced by the normal co-attraction effects between the particles of a more typical PSO. Thus, rather than indicating that a landscape is more searchable by PSOs in general, higher FCI_{cog} measurements may be interpreted to indicate that, on average,

individual regions of the search space are more searchable. Based on this interpretation, it is not expected that the metric should correlate strongly with the rate at which PSOs reduce their diversity, as is supported by the weaker correlations between the FCI_{cog} and DRoC measurements.

All of the benchmark functions used in this section are minimisation problems; therefore, higher FDC values indicate landscapes with higher-quality and less deceptive information to guide the search towards a fitter position. The negative correlation between the FDC metric and the DRoC measure is to be expected: in less deceptive landscapes, the particles of a PSO are guided towards a single fittest region in a more uniform fashion. Such landscapes facilitate convergence; therefore, the particles of PSOs in such landscapes can be expected to reduce their diversity at faster rates.

From the positive correlation between the DM_{med} and DRoC measurements, it appears that single-funnelled landscapes, or landscapes with an underlying unimodal global structure, allow PSOs to reduce their diversities at a faster rate. Contrarily, in multi-funnelled landscapes with underlying multimodal global structures, the PSOs reduced their diversities at a slower rate. Consider that, in single-funnelled landscapes, particles are likely to be attracted towards the same global optimum, as opposed to multi-funnelled landscapes, where particles from different starting regions are more likely to be attracted towards the optima of their regional funnels. It is thus to be expected that single-funnelled landscapes would correspond with the faster rate at which particles reduced their diversities in these landscapes.

It is interesting to note that the $FEM_{0.01}$ measurements correlated negatively to the DRoC measurements. Intuitively, landscapes that are more rugged at the micro level, indicated by higher $FEM_{0.01}$ measurements, may be expected to inhibit the search capabilities of PSOs, resulting in a slower reduction of diversity. Instead, the negative correlations found between the $FEM_{0.01}$ and DRoC measurements indicate that faster reduction in diversity corresponded to more rugged landscapes at the micro level.

The positive correlation between G_{dev} and DRoC measurements indicate that PSOs reduced their diversities at slower rates in landscapes with more variation in their gradients. In other words, for landscapes with steep gradients in some regions and shallow gradients in others, PSOs reduced their diversities at a slower rate than they did for

landscapes with more uniform gradients. The Spearman correlation between the G_{dev} and FCI_{soc} measurements taken in this study is -0.672; this suggests that, of the benchmark functions included in this study, those with more uniform gradients are also more searchable.

The correlations between the DRoC measurements and both the PN and LSN metrics were weak, ranging from 0.036 to only 0.016. These FLCs each indicate the degree of neutrality in a landscape. This suggests that more neutral landscapes may result in a slower reduction of diversity in PSOs, but that the effect is not significant.

The correlation between the DRoC measurements and the $FEM_{0.1}$ measurements were mostly positive but weak, with the strongest correlation being found for the lbest PSO at 0.097. A negative correlation of -0.068 was found for the MBBPSO. This suggests that landscapes with more ruggedness at the macro level correspond with slightly slower convergence in most PSOs, but the effect is not significant.

No discernible correlation was found between the G_{avg} measurements and the DRoC measurements.

4.2.3 Summary

In this section, links between individual FLC metrics and DRoC measurements were investigated by studying the correlation coefficients between those metrics and the corresponding DRoC measurements.

It was found that PSOs reduced their diversity at a faster rate when used on landscapes that were more searchable with regards to PSOs. Specifically, higher FCI_{soc} measurements and lower $FCI_{\bar{\sigma}}$ measurements were associated with faster transitions from exploration to exploitation. Faster reduction in PSO diversity was also associated with less deceptive landscapes, indicated by higher FDC measurements, with single-funnelled landscapes, indicated by negatively-valued DM_{med} measurements, and with landscapes with less variation in gradients, indicated by lower G_{dev} measurements. These FLCs are associated with the difficulty of benchmark functions, and one could conclude that landscapes that are easier to solve in terms of searchability, deception, global landscape structure and gradient features would result in faster convergence in the PSOs solving them.

However, more rugged landscapes also corresponded to faster reduction in diversity in PSOs.

No strong correlations were found between the rate at which PSOs reduced their diversities and the neutrality of the landscapes, as indicated by either PN or LSN measurements. Nor was G_{avg} indicative of PSO search behaviour.

4.3 Conclusions

In this chapter, the link between FLCs and PSO search behaviour was investigated.

The DM metric, which indicates the underlying structure of a landscape, was shown to be occasionally inaccurate. The metric depends on a parameter, n_b, which determines the number of fitter positions to take the dispersion of; however, the optimal choice for this parameter is problem-specific. Additionally, for any choice of n_b, the metric may still yield results that are not indicative of the true structure of the landscape. A modified metric, namely DM_{med}, was proposed to provide more reliable results. The DM_{med} metric, which is obtained as the median of DM values obtained for various n_b-values between 2 and $\frac{n_a}{2}$, was shown to provide more accurate measurements of a landscape's global structure in cases where the unadapted DM metric failed.

The link between single FLCs and PSO search behaviour was investigated by computing correlation coefficients between each FLC metric and DRoC measurements on a suite of benchmark functions. A number of possible relationships between FLC metrics and PSO search behaviour were found. PSOs tended to reduce their diversities at a faster rate when solving landscapes that were indicated to be more searchable by PSOs, landscapes that were less deceptive, single-funnelled landscapes, and landscapes with less variation in their gradients. Neutrality and gradients were not strong indicators of PSO behaviour. More ruggedness in landscapes was associated with faster convergence, not with slower convergence as one might expect.

Chapter 5

Conclusions

The research question posed at the beginning of this dissertation was whether links could be found between the characteristics of optimisation problems and the behaviour of the search algorithms solving those problems. Specifically, the aim was to find correlations between each of a set of fitness landscape characteristics (FLCs) and the search behaviour of a specific class of search algorithms, namely particle swarm optimisers (PSOs).

Chapter 2 discussed the background concepts of the study. The class of particle swarm optimisation algorithms was explained, as well as each specific variation of the algorithm that was used in this study; those being the gbest PSO, the lbest PSO, the Von Neumann PSO (VNPSO), three variations of the guaranteed convergence PSOs (gbest GCPSO, lbest GCPSO, and VNGCPSO), the barebones PSO (BBPSO) and the modified barebones PSO (MBBPSO), the social-only PSO (SPSO), and the cognitive-only PSO (CPSO). These PSOs, except for the CPSO, formed the suite of PSOs that were studied in terms of the change in their behaviour when applied to different problems.

The behaviour of a PSO at a particular instant can be characterised by measuring the diversity of its swarm; this notion of diversity was discussed in Chapter 2, and a metric for quantifying a swarm's diversity, namely the average distance around the swarm centre, was explained. Diversity measurements can be interpreted as follows: When a swarm is relatively diverse, it is in an explorative state. Likewise, when a swarm is relatively concentrated within a small region of the search space, it is in an exploitative state. This metric therefore serves as a suitable indicator of a swarm's behaviour at a

particular instant. However, in order to obtain a profile of a PSO's behaviour over the course of its execution, single diversity measurements are insufficient; a metric is required to give an indication of the overall trend of diversity measurements.

The notion of fitness landscape characteristics (FLCs) was also discussed Chapter 2. A suite of metrics for quantifying specific FLCs was given, namely: the first entropic measures (FEMs) for quantifying ruggedness of a landscape at the macro- and micro-level; the proportion of neutral structures (PN) and the longest subsequence of neutral structures (LSN) encountered during a random walk, for quantifying neutrality in a landscape; the average (G_{avg}) and the standard deviation (G_{dev}) of gradients encountered during a random walk, for quantifying the gradients in a landscape; the dispersion metric (DM) for detecting the presence of multiple funnels in a landscape; the fitness-distance correlation metric (FDC) for quantifying the availability and quality of information in a landscape to guide search towards an optimal solution, or the degree to which a landscape might be deceptive to a search algorithm; and fitness cloud index metrics (FCI_{soc} and FCI_{cog}) which quantify the searchability of a landscape with regards to social-only and cognitive-only PSOs, as well as the fitness cloud mean standard deviation ($FCI_{\bar{\sigma}}$), which quantifies the unpredictability of the FCI metrics. These metrics have been selected to cover a wide range of FLCs, and because each metric requires no prior knowledge of the search space, is suitable for continuous search spaces, and produces a single, scalar-valued output. These are the metrics that were ultimately linked to the behaviour of PSOs.

When the notion of diversity was discussed, it was explained that individual diversity measurements are insufficient for comparing the behaviours of PSOs among one another. A novel metric is required to quantify the behaviour of PSOs. Such a metric was presented in Chapter 3. The metric, called the diversity rate-of-change measure (DRoC), is obtained by taking a series of diversity measurements while a PSO is executing, then fitting a two-piece linear approximation to those diversity measurements, and finally yielding the slope of the first line. The DRoC measure indicates the rate at which a swarm has transitioned from an explorative state, during which its particles were highly diverse, to an exploitative state, during which its particles were concentrated in a small region of the search space. The metric generally produces a negative real value. The

magnitude of such a measurement indicates the rate at which the transition from exploration to exploitation was completed: negative results with a higher magnitude indicate a faster rate of transition.

Because the DRoC measurement is determined based on a series of diversity measurements that are obtained during the execution of a PSO algorithm, its suitability may be expected to depend upon certain variables present during the PSO's execution. Diversity measurements are dependent on the individual particles in a swarm, and so the size of the swarm may influence the DRoC measurements obtained using those diversity measurements. Similarly, diversity measurements depend on the dimensionality of the search space; this dimensionality may have an influence on DRoC measurements. Lastly, the number of diversity measurements taken into account when determining a DRoC measurement may influence the DRoC measurement.

In order to determine whether the DRoC metric is a suitable indicator of the search behaviour of PSOs in terms of exploration and exploitation, the remainder of Chapter 3 served to test the metric's robustness with regards to these variables. The DRoC measure was found to be mostly robust with regards to the size of the PSO's swarm, the dimensionality of the search space, and the number of diversity measurements that it takes into account, with the following limitations. Firstly, the measure does not indicate the earliest iteration at which a swarm can be considered to have transitioned from exploration to exploitation; instead, it indicates the rate at which this transition occurs. The distinction is subtle, but implies that for two PSOs that decrease their diversities at roughly the same rate, but with one PSO smoothing its diversity out at an earlier iteration, the DRoC measure might not be able to differentiate between the two PSOs. Secondly, it was found that the DRoC measure sometimes failed to fit its two-piecewise approximation in a sensible fashion for PSOs exhibiting a very smooth, consistent decrease in diversity. Where the measure would be expected to indicate a relatively slow transition to exploitation, it would sometimes indicate a relatively fast transition instead. Thirdly, for search spaces with very high dimensionalities, PSO swarms often failed to reduce their diversities altogether. In these instances, the DRoC measure did not produce negative results as expected. However, the non-negative DRoC measurements may be interpreted to indicate this failure of a PSO to transition from exploration

to exploitation. Lastly, in order to indicate the rate at which a PSO transitions from exploration to exploitation, the DRoC must be determined from diversity measurements taken at least up until the point when the transition is complete. If the measure is calculated while the PSO is still transitioning, the diversity measurements were found to provide insufficient information to inform the measurement. The number of diversity measurements required is thus problem-specific; however, for most cases, the diversity measurements from the first 500 iterations were sufficient.

In Chapter 4, the metric for indicating the presence of multiple funnels in a landscape, DM, was shown to be an unreliable metric. A modified metric was presented. The modified metric, called the DM_{med}, is obtained by taking several DM measurements, and then taking the median of those measurements. The DM_{med} metric was shown to be a more suitable indicator of a landscape's global structure.

Using the suite of FLCs introduced in Chapter 2 (with DM replaced by the new DM_{med}), as well as the suite of PSOs introduced in Chapter 2, Chapter 4 then addressed the research question by studying correlations between FLCs and the search behaviour of PSOs.

Some promising and intuitive correlations were found. For more searchable landscapes, indicated both by higher FCI_{soc} values and lower $\mathrm{FCI}_{\bar{\sigma}}$ values, the PSOs were found to reduce their diversities at a faster rate. Higher FCI_{cog} values also correlated with faster reduction of diversity, but the correlation was relatively weak; this is possibly because the FCI_{cog} metric indicates the searchability of a landscape with regards to the cognitive-only PSO, which does not exhibit the same convergent behaviour found with the other PSOs studied in this chapter. For landscapes that offer higher quality information to guide search algorithms towards optima, indicated by higher FDC values, the PSOs were also found to transition to exploitation at a faster rate, while lower FDC values, which indicate more deceptive landscapes, correlated with slower convergence in the PSOs. Faster transition to exploitation was observed in PSOs for landscapes with simpler, single-funnelled structures, indicated by lower DM_{med} values, and landscapes with multiple funnels, indicated by higher DM_{med} values, correlated with slower convergence in the PSOs. In landscapes with less variation in their gradients, indicated by lower G_{dev} values, the PSOs were found to converge at a faster rate.

Some of the resulting correlations were surprising. Landscapes with more ruggedness at the micro level, indicated by higher $FEM_{0.01}$ values, may be expected to correlate with a slower transition to exploitation due to the increased complexity in the landscapes. Instead, these landscapes correlated negatively with the DRoC measure, indicating that the PSOs transitioned to exploitation at a faster rate.

The remaining correlations were weak, indicating that the FLCs in question did not act as strong indicators of PSO search behaviour. For more neutral landscapes, indicated by higher PN and LSN values, the PSOs were found to reduce their diversities at a slightly slower rate, though the correlation was not significant. For landscapes with more ruggedness at the macro level, indicated by higher $FEM_{0.1}$ values, most PSOs were found to reduce their diversities at a slightly slower rate, though the correlation was not significant, and an exception was found for the MMPSO, which was found to reduce its diversities at a slightly faster rate. No correlation was found between the average steepness of gradients in a landscape, given by G_{avg}, and the rate at which the PSOs converged.

For most of the findings, it seems that when the FLC values indicated a landscape to be easier to optimise, the PSOs were able to converge on a promising region at a faster rate. This was the case for landscapes that could be considered easier to solve due to being more searchable, being less deceptive, having simpler underlying structures, and having less variation in gradients. The exceptions were that landscapes did not correlate with faster convergence when those landscapes exhibited less ruggedness, less neutrality, or shallower gradients.

These findings may form a basis upon which an algorithm selection framework can be founded by providing a deeper understanding of the link between optimisation problems and search algorithm behaviour. Being able to predict how a search algorithm will behave on an optimisation problem may eventually allow optimisation practitioners to choose an optimal algorithm for any given problem, based on a quick analysis of the characteristics of the problem.

The DRoC measure, which was used to quantify the search behaviour of PSOs, exhibited certain limitations. The measure did not indicate the approximate iteration at which a swarm transitions from exploration to exploitation; instead, it indicated the rate

at which this transition occurred in terms of decreased swarm diversity. The measure also failed in some test cases where a swarm decreased its diversity at very gradual rate, and in some test cases where a swarm showed a trend of increasing diversity after the initial transition to exploitation.

Future work is needed to address these limitations. For this study, the DRoC measure was determined as the slope of the first line of a two-piecewise linear approximation. Other metrics from such an approximation may prove to be successful indicators of algorithm behaviour; for example, the x-coordinate of the intersection between the two lines could be an indicator of the iteration at which the transition to exploitation is completed, and could thus indicate the duration of the algorithm's explorative phase. The limitations of the measure might also be overcome by generalising the two-piecewise approximation to other regression techniques, such as polynomial regression.

The scope of this research was limited in certain aspects with regards to the research question. The characteristics of optimisation problems were only studied in terms of FLCs, and only in terms of a specific suite of metrics to indicate particular FLCs. The optimisation problems considered were all continuous, with a single global optimum, and did not change dynamically. The search algorithms studied were limited to a suite of PSOs that are known to exhibit convergent behaviour. The search behaviour of the PSOs was measured only in terms of the rate at which they transitioned from exploration to exploitation.

Future research may focus on expanding upon the above limitations. Optimisation problems may be characterised in terms of other FLC metrics not included in this study, such as low-level exploratory landscape metrics [42]. Aside from PSOs, many other search techniques remain to be studied, including other population-based algorithms such as genetic algorithms [22], ant colony optimisation [10], and firefly optimisation [69], as well as a plethora of other search techniques, such as gradient descent, tabu search [19], and simulated annealing [30]. For any given search technique, the search behaviour of the algorithm may be quantified in various ways, possibly shedding more light on the relationship between the problems being solved and the algorithms solving them.

Bibliography

[1] B. O. Arani, P. Mirzabeygi, and M. S. Panahi. An improved PSO algorithm with a territorial diversity-preserving scheme and enhanced exploration–exploitation balance. *Swarm and Evolutionary Computation*, 11:1–15, 2013.

[2] A. Auger and O. Teytaud. Continuous lunches are free plus the design of optimal optimization algorithms. *Algorithmica*, 57(1):121–146, 2010.

[3] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys (CSUR)*, 35(3):268–308, 2003.

[4] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm intelligence: from natural to artificial systems*. Number 1. Oxford university press, 1999.

[5] P. Bosman and A. P. Engelbrecht. Diversity rate of change measurement for particle swarm optimisers. In M. Dorigo, M. Birattari, S. Garnier, H. Hamann, M. Montes de Oca, C. Solnon, and T. Stützle, editors, *Swarm Intelligence*, pages 86–97, Cham, 2014. Springer International Publishing.

[6] X. Chen and Y. Li. A modified PSO structure resulting in high exploration ability with convergence guaranteed. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(5):1271–1289, 2007.

[7] M. Clerc. From theory to practice in particle swarm optimization. In *Handbook of Swarm Intelligence*, pages 3–36. Springer, 2011.

[8] K. A. De Jong. *Analysis of the behavior of a class of genetic adaptive systems*. PhD thesis, Computer and Communication Sciences, University of Michigan, Ann Arbor, 1975.

[9] Y. Delice, E. K. Aydoğan, U. Özcan, and M. S. İlkay. A modified particle swarm optimization algorithm to mixed-model two-sided assembly line balancing. *Journal of Intelligent Manufacturing*, 28(1):23–36, 2017.

[10] M. Dorigo and G. Di Caro. Ant colony optimization: a new meta-heuristic. In *Proceedings of the IEEE Congress on Evolutionary Computation*, volume 2, pages 1470–1477. IEEE, 1999.

[11] R. Eberhart and J. Kennedy. A new optimizer using particle swarm theory. In *Proceedings of the Sixth International Symposium on Micro Machine and Human Science*, pages 39–43, Oct 1995.

[12] M. Ebner, P. Langguth, J. Albert, M. Shackleton, and R. Shipman. On neutral networks and evolvability. In *Proceedings of the IEEE Congress on Evolutionary Computation*, volume 1, pages 1–8 vol. 1, 2001.

[13] A. P. Engelbrecht. Scalability of a heterogeneous particle swarm optimizer. In *Proceedings of the IEEE Symposium on Swarm Intelligence*, pages 1–8, April 2011.

[14] I. Fister Jr, M. Perc, K. Ljubič, S. M. Kamal, A. Iglesias, and I. Fister. Particle swarm optimization for automatic creation of complex graphic characters. *Chaos, Solitons & Fractals*, 73:29–35, 2015.

[15] C. A. Floudas and C. E. Gounaris. A review of recent advances in global optimization. *Journal of Global Optimization*, 45(1):3–38, 2009.

[16] A. H. Gandomi and A. H. Alavi. Krill herd: a new bio-inspired optimization algorithm. *Communications in Nonlinear Science and Numerical Simulation*, 17(12):4831–4845, 2012.

[17] E. García-Gonzalo and J. Fernández-Martínez. A brief historical review of particle swarm optimization PSO. *Journal of Bioinformatics and Intelligent Control*, 1(1):3–16, 2012.

[18] S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.

[19] F. Glover. Future paths for integer programming and links to artificial intelligence. *Computers & operations research*, 13(5):533–549, 1986.

[20] M. K. Gupta, P. Sood, and V. S. Sharma. Machining parameters optimization of titanium alloy using response surface methodology and particle swarm optimization under minimum-quantity lubrication environment. *Materials and Manufacturing Processes*, 31(13):1671–1682, 2016.

[21] N. Hansen and S. Kern. Evaluating the CMA evolution strategy on multimodal test functions. In *Proceedings of the International Conference on Parallel Problem Solving from Nature*, pages 282–291. Springer, 2004.

[22] J. H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press, 1992.

[23] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python. `http://www.scipy.org/` [Online; accessed 2018-07-29], 2001–.

[24] T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 184–192. Morgan Kaufmann, 1995.

[25] A. R. Jordehi. Enhanced leader PSO (ELPSO): a new PSO variant for solving global optimisation problems. *Applied Soft Computing*, 26:401–417, 2015.

[26] J. Kennedy. The particle swarm: social adaptation of knowledge. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, pages 303–308, Apr 1997.

[27] J. Kennedy. Bare bones particle swarms. In *Proceedings of the IEEE Swarm Intelligence Symposium*, pages 80–87, April 2003.

[28] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of the IEEE International Conference on Neural Networks*, volume 4, pages 1942–1948 vol.4, Nov 1995.

[29] J. Kennedy and R. Mendes. Population structure and particle swarm performance. In *Proceedings of the Congress on Evolutionary Computation*, volume 2, pages 1671–1676, 2002.

[30] A. Khachaturyan, S. Semenovskaya, and B. Vainstein. A statistical-thermodynamic approach to determination of structure amplitude phases. *Soviet Physics, Crystallography*, 24(5):519–524, 1979.

[31] K. Krishnanand and D. Ghose. Glowworm swarm optimisation: a new method for optimising multi-modal functions. *International Journal of Computational Intelligence Studies*, 1(1):93–119, 2009.

[32] D. A. Levinthal. Adaptation on rugged landscapes. *Management Science*, 43(7):934–950, 1997.

[33] M. Lunacek and D. Whitley. The dispersion metric and the CMA evolution strategy. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 477–484, New York, NY, USA, 2006. ACM.

[34] K. M. Malan. *Characterising continuous optimisation problems for particle swarm optimisation performance prediction*. PhD thesis, University of Pretoria, 2014.

[35] K. M. Malan and A. P. Engelbrecht. Quantifying ruggedness of continuous landscapes using entropy. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1440–1447. IEEE, 2009.

[36] K. M. Malan and A. P. Engelbrecht. Ruggedness, funnels and gradients in fitness landscapes and the effect on PSO performance. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 963–970. IEEE, 2013.

[37] K. M. Malan and A. P. Engelbrecht. Characterising the searchability of continuous optimisation problems for PSO. *Swarm Intelligence*, 8(4):275–302, 2014.

[38] K. M. Malan and A. P. Engelbrecht. A progressive random walk algorithm for sampling continuous fitness landscapes. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 2507–2514, July 2014.

[39] R. F. Malik, T. A. Rahman, S. Z. M. Hashim, and R. Ngah. New particle swarm optimizer with sigmoid increasing inertia weight. *International Journal of Computer Science and Security*, 1(2):35–44, 2007.

[40] P. Marrow. Evolvability: Evolution, computation, biology. In *Proceedings of the Genetic and Evolutionary Computation Conference Workshop Program*, pages 30–33, 1999.

[41] X.-B. Meng, X. Z. Gao, L. Lu, Y. Liu, and H. Zhang. A new bio-inspired optimisation algorithm: Bird swarm algorithm. *Journal of Experimental & Theoretical Artificial Intelligence*, 28(4):673–687, 2016.

[42] O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, and G. Rudolph. Exploratory landscape analysis. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, pages 829–836. ACM, 2011.

[43] S. Mirjalili and A. Lewis. The whale optimization algorithm. *Advances in Engineering Software*, 95:51–67, 2016.

[44] S. K. Mishra. Performance of repulsive particle swarm method in global optimization of some important test functions: A fortran program. Technical report, Social Science Research Network, August 2006.

[45] S. K. Mishra. Some new test functions for global optimization and performance of repulsive particle swarm method. Technical Report 2718, University Library of Munich, Germany, August 2006.

[46] E. T. Oldewage. *The perils of particle swarm optimization in high dimensional problem spaces*. PhD thesis, University of Pretoria, 2018.

[47] O. Olorunda and A. P. Engelbrecht. Measuring exploration/exploitation in particle swarms using swarm diversity. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1128–1134. IEEE, 2008.

[48] R. Palmer. Optimization on rugged landscapes. In A. S. Perelson, editor, *Molecular Evolution on Rugged Landscapes: Protein, RNA, and the Immune System*, volume 9. CRC Press, 2018.

[49] S. Pandey, L. Wu, S. M. Guru, and R. Buyya. A particle swarm optimization-based heuristic for scheduling workflow applications in cloud computing environments. In *Proceedings of the IEEE International Conference on Advanced Information Networking and Applications*, pages 400–407. IEEE, 2010.

[50] R. S. Parpinelli and H. S. Lopes. New inspirations in swarm intelligence: a survey. *International Journal of Bio-Inspired Computation*, 3(1):1–16, 2011.

[51] E. S. Peer, F. van den Bergh, and A. P. Engelbrecht. Using neighbourhoods with the guaranteed convergence PSO. In *Proceedings of the IEEE Swarm Intelligence Symposium*, pages 235–242. IEEE, 2003.

[52] K. V. Price, R. M. Storn, and J. A. Lampinen. Unconstrained unimodal test functions. In *Differential Evolution A Practical Approach to Global Optimization*, pages 514–533. Springer-Verlag, 2005.

[53] S. Rahnamayan, H. R. Tizhoosh, and M. M. Salama. A novel population initialization method for accelerating evolutionary algorithms. *Computers & Mathematics with Applications*, 53(10):1605–1614, 2007.

[54] J. R. Rice. The algorithm selection problem. In *Advances in computers*, volume 15, pages 65–118. Elsevier, 1976.

[55] Y. Shi and R. Eberhart. A modified particle swarm optimizer. In *Proceedings of the IEEE International Conference on Evolutionary Computation Proceedings*, pages 69–73. IEEE, 1998.

[56] C. Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.

[57] I. C. Trelea. The particle swarm optimization algorithm: convergence analysis and parameter selection. *Information Processing Letters*, 85(6):317–325, 2003.

[58] W. A. van Aardt, A. S. Bosman, and K. M. Malan. Characterising neutrality in neural network error landscapes. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1374–1381. IEEE, 2017.

[59] F. Van Den Bergh. *An analysis of particle swarm optimizers*. PhD thesis, University of Pretoria South Africa, 2001.

[60] F. van den Bergh and A. P. Engelbrecht. A new locally convergent particle swarm optimizer. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, volume 3, pages 94–99, 2002.

[61] F. Van Den Bergh and A. P. Engelbrecht. A study of particle swarm optimization particle trajectories. *Information Sciences*, 176(8):937–971, 2006.

[62] F. Van den Bergh and A. P. Engelbrecht. A convergence proof for the particle swarm optimiser. *Fundamenta Informaticae*, 105(4):341–374, 2010.

[63] V. K. Vassilev, T. C. Fogarty, and J. F. Miller. Smoothness, ruggedness and neutrality of fitness landscapes: from theory to application. In *Advances in Evolutionary Computing*, pages 3–44. Springer, 2003.

[64] S. Verel, P. Collard, and M. Clergue. Where are bottlenecks in NK fitness landscapes? In *Proceedings of the IEEE Congress on Evolutionary Computation*, volume 1, pages 273–280. IEEE, 2003.

[65] M. Villalobos-Arias, C. A. C. Coello, and O. Hernández-Lerma. Asymptotic convergence of metaheuristics for multiobjective optimization problems. *Soft Computing*, 10(11):1001–1005, 2006.

[66] D. H. Wolpert and W. G. Macready. Coevolutionary free lunches. *IEEE Transactions on Evolutionary Computation*, 9(6):721–735, 2005.

[67] D. H. Wolpert, W. G. Macready, et al. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82, 1997.

[68] S. Wright. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. In *Proceedings of the Sixth International Congress on Genetics*, pages 356—-366, 1932.

[69] X.-S. Yang. Firefly algorithms for multimodal optimization. In *Proceedings of the International Symposium on Stochastic Algorithms*, pages 169–178. Springer, 2009.

[70] X.-S. Yang. *Nature-inspired metaheuristic algorithms*. Luniver press, 2010.

[71] X.-S. Yang. A new metaheuristic bat-inspired algorithm. In *Nature Inspired Cooperative Strategies for Optimization*, pages 65–74. Springer, 2010.

[72] X.-S. Yang and S. Deb. Cuckoo search via Lévy flights. In *Proceedings of the World Congress on Nature & Biologically Inspired Computing*, pages 210–214. IEEE, 2009.

[73] X.-S. Yang and X. He. Firefly algorithm: Recent advances and applications. *International Journal of Swarm Intelligence*, 2013.

[74] X. Yao, Y. Liu, and G. Lin. Evolutionary programming made faster. *IEEE Transactions on Evolutionary Computation*, 3(2):82–102, 1999.

[75] L. Zhang, H. Yu, and S. Hu. A new approach to improve particle swarm optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 134–139. Springer, 2003.

Appendix A

Symbols

A.1 Chapter 2

$\mathbf{x}$	A solution vector
$\mathbf{x}^*$	Globally optimal solution vector
$f(\mathbf{x})$	Fitness of candidate solution $\mathbf{x}$
$\mathcal{S}$	Search space
$\mathcal{N}$	Neighbourhood
n_t	Number of steps in a walk
ε	Error parameter
$S(\varepsilon)$	String of symbols
s_i	Symbol i in a string of symbols
$H(\varepsilon)$	Measure of entropy
ε^*	Information stability
$\text{FEM}_{0.1}$	First entropic measure of macro ruggedness
$\text{FEM}_{0.01}$	First entropic measure of micro ruggedness
n_n	Number of neutral structures in a walk
$\|W\|$	Total number of structures in a walk
w_n	Neutral-only subsequence of a walk
g_t	Gradient at step t
$\text{PN}(\varepsilon)$	Proportion of neutral structures in a walk

$\mathrm{LSN}(\varepsilon)$	Longest subsequence of neutral structures in a walk
u	Step size
G_{avg}	Average estimated gradient within a walk
G_{dev}	Standard deviation of gradient measures from the mean
n_a	Sample size
S_{n_a}	Sample of n_a positions
n_b	Sub-sample size
$S^*_{n_b}$	Subset of the n_b fittest positions of S_{n_a}
$\mathrm{disp}(S)$	Dispersion of S
DM	Dispersion metric
d^*_i	Distance from position i to x^*
FDC_m	Fitness distance correlation measure
$\mathbf{x}'_i$	Updated solution vector at index i
n_v	Number of valid solutions after updates
FCI_{cog}	Fitness cloud index based on cognitive updates
FCI_{soc}	Fitness cloud index based on social updates
$\mathrm{FCI}_{\bar{\sigma}}$	Fitness cloud index mean standard deviation
n_s	Number of particles
D	Dimensionality of the search space
$\mathbf{x}_i(t)$	Position of particle i at time step t
$\mathbf{v}_i(t)$	Velocity of particle i at time step t
w	Inertia weight constant
c_1	Cognitive acceleration constant
c_2	Social acceleration constant
$\mathbf{y}_i(t)$	Personal best position of particle i at time step t
$\hat{\mathbf{y}}(t)$	Global best position at time step t
$\hat{\mathbf{y}}_i(t)$	Neighbourhood best position of particle i at time step t
τ	Index of particle with fittest personal best
$\rho_i(t)$	Local search bounding box scaling factor for particle i at time step t

A.2 Chapter 3

$\mathcal{D}$	Swarm diversity
x_{ik}	k-th dimensional component of position i
$y(t) \sim \mathcal{D}(t)$	Two-piecewise linear approximation of $\mathcal{D}$
a_j	y intersection of the j-th line segment in piecewise linear approximation
b_j	Gradient of the j-th line segment in two-piecewise linear approximation
t'	x offset of the breakpoint in two-piecewise linear approximation
LSE	Least squares error

A.3 Chapter 4

DM_{med}	Median disperion metric
ι	Incrementation value used in DM_{med}

Appendix B

Complete Results from DRoC Robustness Experiments

The robustness of the DRoC measure was tested in Chapter 3 by ranking DRoC measurements against one another for pairs of PSO algorithms. Results were only listed and discussed for pairs of algorithms for which expectations regarding their differing search behaviour had been defined. This appendix lists the full results with comparisons among all pairs of PSO algorithms.

B.1 Robustness with Regards to Swarm Size

Table B.1: Complete ranks of DRoCs for PSO pairs with various swarm sizes

Swarm size:	5	10	25	50	75	100	500
Gbest PSO, Lbest PSO							
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, VNPSO							

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	0	-1	-1	-1	-1	-1	-1
Weierstrass	0	0	-1	-1	-1	-1	-1
Gbest PSO, Gbest GCPSO							
Spherical	0	0	0	0	0	0	0
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	0	-1	0	0	0	-1	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, Lbest GCPSO							
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, VNGCPSO							
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, BBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	1	0	1	0	0	0	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, MBBPSO							
Spherical	0	0	1	1	1	1	1
Rastrigin	-1	-1	0	0	0	-1	-1
Rosenbrock	0	-1	-1	-1	-1	-1	-1

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, SPSO							
Spherical	1	1	1	1	1	1	0
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	0	0	-1	-1	-1
Weierstrass	1	1	1	1	1	0	0
Lbest PSO, VNPSO							
Spherical	0	0	0	1	1	1	1
Rastrigin	0	0	0	-1	-1	-1	-1
Rosenbrock	0	0	1	1	1	1	1
Weierstrass	0	0	1	1	1	1	1
Lbest PSO, Gbest GCPSO							
Spherical	0	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	0	1	1	1	1	1	1
Weierstrass	-1	0	1	1	1	1	1
Lbest PSO, Lbest GCPSO							
Spherical	-1	0	0	0	-1	-1	-1
Rastrigin	0	0	0	0	1	0	1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Lbest PSO, VNGCPSO							
Spherical	0	0	0	0	0	0	-1
Rastrigin	0	1	0	0	0	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Lbest PSO, BBPSO							
Spherical	1	1	1	1	1	1	1

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Lbest PSO, MBBPSO							
Spherical	0	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	0	0	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Lbest PSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
VNPSO, Gbest GCPSO							
Spherical	0	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	0	1	1	1	1	1	1
Weierstrass	-1	-1	1	1	1	1	1
VNPSO, Lbest GCPSO							
Spherical	-1	0	-1	-1	-1	-1	-1
Rastrigin	0	0	0	1	1	1	1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
VNPSO, VNGCPSO							
Spherical	0	0	-1	-1	-1	-1	-1
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
VNPSO, BBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
VNPSO, MBBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	0	0	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
VNPSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
Gbest GCPSO, Lbest GCPSO							
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest GCPSO, VNGCPSO							
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest GCPSO, BBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Rosenbrock	1	1	1	0	0	1	0
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest GCPSO, MBBPSO							
Spherical	0	1	1	1	1	1	1
Rastrigin	-1	-1	0	-1	-1	0	-1
Rosenbrock	0	0	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest GCPSO, SPSO							
Spherical	1	1	0	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	0	0	-1	0	-1
Weierstrass	1	1	1	1	1	1	0
Lbest GCPSO, VNGCPSO							
Spherical	0	0	0	0	0	0	1
Rastrigin	0	1	0	-1	-1	-1	-1
Rosenbrock	0	0	1	1	1	1	1
Weierstrass	0	0	0	0	0	0	1
Lbest GCPSO, BBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	-1	-1	-1	-1	0	0
Lbest GCPSO, MBBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	0
Lbest GCPSO, SPSO							

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
VNGCPSO, BBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	-1	-1	-1	-1	0	0
VNGCPSO, MBBPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	0	0	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	-1	-1	-1	-1	-1	-1	0
VNGCPSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
BBPSO, MBBPSO							
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	0	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	0	-1	0	0	0	0
BBPSO, SPSO							
Spherical	0	-1	-1	-1	-1	-1	-1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	0	0	-1	0	-1	-1	-1

Continued on next page ...

Table B.1 (continued)

Swarm size:	5	10	25	50	75	100	500
Weierstrass	1	1	1	1	1	1	1
MBBPSO, SPSO							
Spherical	1	0	-1	0	-1	-1	-1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1

B.2 Robustness with Regards to Dimensionality

Table B.2: Complete ranks of DRoCs for PSO pairs in various dimensionalities

Dimensionality:	2	5	25	50	100	500	1000
Gbest PSO, Lbest PSO							
Spherical	-1	-1	-1	-1	-1	0	-1
Rastrigin	-1	-1	-1	0	-1	0	-1
Rosenbrock	-1	-1	-1	-1	0	-1	0
Weierstrass	-1	-1	-1	-1	0	0	-1
Gbest PSO, VNPSO							
Spherical	-1	-1	-1	-1	0	0	0
Rastrigin	-1	-1	-1	0	0	0	0
Rosenbrock	-1	-1	-1	-1	0	0	0
Weierstrass	-1	-1	-1	-1	0	0	0
Gbest PSO, Gbest GCPSO							
Spherical	0	0	0	0	-1	0	0
Rastrigin	0	0	0	0	0	0	0
Rosenbrock	0	0	-1	0	0	0	0
Weierstrass	-1	-1	0	0	0	0	0
Gbest PSO, Lbest GCPSO							

Continued on next page ...

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Spherical	-1	-1	-1	-1	-1	-1	-1
Rastrigin	-1	-1	-1	0	-1	0	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	-1	-1	-1
Gbest PSO, VNGCPSO							
Spherical	-1	-1	-1	1	0	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	-1	1	-1	-1	0
Weierstrass	-1	-1	-1	0	0	0	-1
Gbest PSO, BBPSO							
Spherical	1	1	0	1	0	0	0
Rastrigin	1	1	0	1	0	0	0
Rosenbrock	1	1	-1	-1	0	-1	0
Weierstrass	1	-1	-1	-1	1	1	1
Gbest PSO, MBBPSO							
Spherical	1	1	0	1	1	-1	-1
Rastrigin	0	0	-1	1	-1	0	-1
Rosenbrock	0	-1	-1	1	-1	-1	-1
Weierstrass	1	-1	-1	-1	1	1	1
Gbest PSO, SPSO							
Spherical	0	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	0	1	1	1	1	1
Weierstrass	0	1	1	1	1	1	1
Lbest PSO, VNPSO							
Spherical	0	0	1	1	0	0	1
Rastrigin	1	0	1	0	1	0	1
Rosenbrock	0	1	1	1	0	1	0

Continued on next page ...

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Weierstrass	0	1	1	0	0	0	1
Lbest PSO, Gbest GCPSO							
Spherical	1	1	1	1	0	0	1
Rastrigin	1	1	1	1	1	0	1
Rosenbrock	1	1	1	1	-1	1	0
Weierstrass	0	1	1	1	0	0	1
Lbest PSO, Lbest GCPSO							
Spherical	0	0	1	1	0	-1	-1
Rastrigin	0	0	1	1	0	-1	-1
Rosenbrock	0	-1	0	1	-1	-1	-1
Weierstrass	-1	-1	-1	0	0	-1	0
Lbest PSO, VNGCPSO							
Spherical	0	0	1	1	0	-1	-1
Rastrigin	1	0	1	1	-1	-1	-1
Rosenbrock	0	-1	1	1	0	-1	0
Weierstrass	-1	-1	1	1	0	-1	-1
Lbest PSO, BBPSO							
Spherical	1	1	1	1	0	0	1
Rastrigin	1	1	1	1	0	0	1
Rosenbrock	1	1	1	1	0	0	0
Weierstrass	1	-1	-1	1	1	1	1
Lbest PSO, MBBPSO							
Spherical	1	1	1	1	1	-1	-1
Rastrigin	1	1	1	1	-1	-1	-1
Rosenbrock	1	1	1	1	-1	0	0
Weierstrass	1	-1	-1	1	1	1	1
Lbest PSO, SPSO							
Spherical	1	1	1	1	1	1	1

Continued on next page ...

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
VNPSO, Gbest GCPSO							
Spherical	1	1	1	1	0	0	0
Rastrigin	1	1	1	1	0	0	0
Rosenbrock	1	1	1	1	-1	0	0
Weierstrass	0	1	1	1	0	0	0
VNPSO, Lbest GCPSO							
Spherical	-1	-1	0	0	0	-1	-1
Rastrigin	-1	0	0	0	-1	-1	-1
Rosenbrock	0	-1	-1	0	-1	-1	-1
Weierstrass	-1	-1	-1	0	-1	-1	-1
VNPSO, VNGCPSO							
Spherical	-1	-1	1	1	0	-1	-1
Rastrigin	0	0	1	1	-1	-1	-1
Rosenbrock	0	-1	1	1	0	-1	0
Weierstrass	-1	-1	-1	1	0	0	-1
VNPSO, BBPSO							
Spherical	1	1	1	1	0	0	0
Rastrigin	1	1	1	1	0	0	0
Rosenbrock	1	1	0	0	0	0	0
Weierstrass	1	-1	-1	0	1	1	1
VNPSO, MBBPSO							
Spherical	1	1	1	1	1	-1	-1
Rastrigin	1	1	1	1	-1	-1	-1
Rosenbrock	1	1	1	1	-1	0	0
Weierstrass	1	-1	-1	1	1	1	1

Continued on next page . . .

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
VNPSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
Gbest GCPSO, Lbest GCPSO							
Spherical	-1	-1	-1	-1	0	-1	-1
Rastrigin	-1	-1	-1	-1	-1	-1	-1
Rosenbrock	-1	-1	-1	-1	-1	-1	-1
Weierstrass	-1	-1	-1	-1	0	-1	-1
Gbest GCPSO, VNGCPSO							
Spherical	-1	-1	-1	0	0	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	-1	0	0	-1	0
Weierstrass	-1	-1	-1	0	0	0	-1
Gbest GCPSO, BBPSO							
Spherical	1	1	0	1	0	0	0
Rastrigin	1	1	0	1	-1	0	0
Rosenbrock	0	1	-1	-1	0	0	0
Weierstrass	1	-1	-1	0	1	1	1
Gbest GCPSO, MBBPSO							
Spherical	1	1	-1	1	1	-1	-1
Rastrigin	0	0	-1	1	-1	-1	-1
Rosenbrock	0	-1	-1	1	0	-1	-1
Weierstrass	1	-1	-1	0	1	1	1
Gbest GCPSO, SPSO							
Spherical	0	0	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1

Continued on next page ...

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Rosenbrock	1	0	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
Lbest GCPSO, VNGCPSO							
Spherical	0	0	1	1	0	0	-1
Rastrigin	0	0	1	1	0	-1	0
Rosenbrock	0	1	1	1	0	0	0
Weierstrass	0	0	1	1	0	0	0
Lbest GCPSO, BBPSO							
Spherical	1	1	1	1	0	1	1
Rastrigin	1	1	1	1	0	1	1
Rosenbrock	1	1	1	0	1	1	1
Weierstrass	1	-1	-1	1	1	1	1
Lbest GCPSO, MBBPSO							
Spherical	1	1	1	1	1	1	0
Rastrigin	1	1	1	1	-1	0	0
Rosenbrock	1	1	1	1	0	1	0
Weierstrass	1	-1	-1	1	1	1	1
Lbest GCPSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
VNGCPSO, BBPSO							
Spherical	1	1	1	1	0	1	1
Rastrigin	1	1	1	1	0	1	1
Rosenbrock	1	1	0	-1	0	1	0
Weierstrass	1	-1	-1	-1	1	1	1
VNGCPSO, MBBPSO							

Continued on next page . . .

Table B.2 (continued)

Dimensionality:	2	5	25	50	100	500	1000
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	0	1	0	1	0
Rosenbrock	1	1	1	1	0	1	0
Weierstrass	1	-1	-1	-1	1	1	1
VNGCPSO, SPSO							
Spherical	1	1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	1	1	1	1	1	1	1
BBPSO, MBBPSO							
Spherical	-1	-1	0	1	1	-1	-1
Rastrigin	-1	-1	-1	1	-1	-1	-1
Rosenbrock	-1	-1	1	1	-1	0	-1
Weierstrass	-1	-1	1	1	0	0	0
BBPSO, SPSO							
Spherical	-1	-1	1	1	1	1	1
Rastrigin	0	1	1	1	1	1	1
Rosenbrock	1	-1	1	1	1	1	1
Weierstrass	-1	1	1	1	1	1	1
MBBPSO, SPSO							
Spherical	-1	-1	1	1	1	1	1
Rastrigin	1	1	1	1	1	1	1
Rosenbrock	1	1	1	1	1	1	1
Weierstrass	0	1	1	1	1	0	1

Appendix C

Derived Publications

The following publications are derived from this study:

1. P. Bosman and A. P. Engelbrecht. Diversity rate of change measurement for particle swarm optimisers. In Dorigo, Marco and Birattari, Mauro and Garnier, Simon and Hamann, Heiko and Montes de Oca, Marco and Solnon, Christine and Stützle, Thomas, editors, *Swarm Intelligence*, pages 86–97, Cham, 2014. Springer International Publishing.