Continuous Behavioural Variation Particle Swarm Optimisation (CBV-PSO)

Ed Keedwell[1], Mark Morley[1], Darren Croft[2]

[1]College of Engineering, Mathematics and Physical Sciences, University of Exeter, Harrison Building, North Park Road, Exeter, EX4 4QF

[2]Center for Research in Animal Behaviour, College of Life and Environmental Sciences, Washington Singer Laboratories, University of Exeter, Exeter, EX4 4QG

Abstract: Traditional PSO, modelled on individuals of the same species assumes homogeneity of behaviour across the swarm, despite there being considerable variation in capability of individuals within a swarm in the natural world. This paper proposes a PSO algorithm based on continuous behavioural variation, or the ability for each member of a swarm to have heterogeneous behavioural capabilities on a continuous scale. Experimentation with standard PSO test functions demonstrates that a continuously heterogeneous population is able to search more effectively than a homogeneous population for all problems of lower dimensionality and some of higher dimensionality.

Introduction

Traditional particle swarm optimisation consists of a population of particles analogous to members of a flock or swarm in the natural world, and considers each particle as a solution to a problem and the movement of these particles through decision space determines the exploration capabilities of the algorithm. PSO maintains a homogeneous population of these particles which are governed by an identical set of movement equations. Although particles adopt very different trajectories as a result of their initial random starting positions, the movement of two particles starting from the same position will be identical.

In natural flocks, individuals are often of the same species, but there exists considerable variation in the traits possessed by each individual. In much the same way as humans display disparate levels of aggression, gregariousness and inquisitiveness, so do the animals on which PSO is based [1]. Recent research has shown that this disparity of behaviour is very important in the ability of the flock to solve problems effectively, which might have profound implications for PSO. One of the key aspects is that although certain behaviour types (e.g. more adventurous individuals) might individually be better at problem solving; selecting for a group that all have adventurous traits reduces performance of the flock as a whole [2]. Therefore a flock that has a variety of behaviours leads to better performance in nature and it is this performance that motivates the work in this paper.

This paper explores a variant of PSO known as Continuous Behavioural Variation PSO (CBV-PSO) where individuals within a swarm have traits based on a continuous scale of variation as opposed to discrete behaviour groupings (e.g. ‘gregarious’ or ‘shy’ individuals). The algorithm is tested on a number of well known problems taken from the literature and results show that in a variety of circumstances, CBV-PSO outperforms a standard PSO formulation.

Previous Research

Previous research in the field of behavioural variation in PSO has been extensive (e.g. see [3]) but the work of direct relevance here is concentrated in two primary areas. The first is the use of multiple swarms intended to model different species, with disparity of behaviour in each swarm indicating membership of a particular species. Work by [4] and [5] has shown that this can be an effective method for searching large spaces, but usually there is no interaction between swarm members of the two species. This use of multiple species within a PSO run is distinct from the behavioural variation discussed here due to the lack of interaction and the highly discrete delineation between species. In contrast CBV-PSO is concerned with continuous intra-species variation, whereas speciation is naturally concerned with discrete inter-species variation.

The second method of introducing variation, and in this case, intra-species variation, is the recent work of Andries Engelbrecht [6] [7]. The methods employed in this work are similar in motivation to CBV-PSO and exploit intra-species variation via a discrete 'behaviour' pool. This means that the individuals in the swarm have a finite number of behaviours from which to select at the beginning of the optimisation. The behaviours effect the movement of the particles around the search space, leading to very different exploration and exploitation behaviours for the individuals within the swarm. A further advance on this approach allows unsuccessful particles to modify their behaviour to improve performance [6]. This method effectively performs a meta-search of the behavioural space in addition to the initial selection of behaviours and has been found to improve performance further. The discrete behaviour groups used in this work are essentially artefacts used by humans to categorise behaviors. However, animals are not simply shy, gregarious, aggressive, or passive and a large range of behaviours will exist between these extremes, particularly when discussing a population of individuals. Nevertheless, the work of Engelbrecht has shown that diversity of behaviour is a powerful improvement on PSO and has led to significant improvements in performance.

It is against this backdrop of behavioural variation that the CBV-PSO has been developed. The algorithm has been developed to be as biologically plausible as possible, primarily with the aim of recreating the results seen in the natural world on the schooling behaviours of guppies [2]. The algorithm achieves heterogeneity through the generation of $c1$ and $c2$ coefficient values on a continuous scale to determine different behaviours for each of the particles in the initialisation stage. Once determined, these coefficient selections remain constant for each particle for the duration of the optimisation. The remainder of the PSO algorithm is then run in a standard fashion.

Method

CBV-PSO is based on a standard PSO implementation. In standard PSO implementations, each particle is identical, although they may adopt different positions within the search space. CBV modifies standard PSO by introducing variation in the extent to which the particles are influenced by the global and local best positions within the group. The algorithm therefore is run as follows:

1. Generate a random population of particles.
2. For each particle, select *c1* and *c2* coefficient values drawn from a Gaussian distribution with a specified mean and standard deviation.
3. Evaluate the population of particles
4. Use standard computation of velocity (b) to determine the new position of each particle (a):
    a. $x_{ij}[t+1] = x_{ij}[t]+v_{ij}[t]$
    b. $v_{ij}[t+1] = v_{ij}[t] +c1*r1*(pbestx_{ij}-x_{ij})+c2*r2*(gbestx_{ij-} x_{ij})$

    Where $x_{ij}[t]$ is the position of the particle *i* in dimension *j* at time *t*, $v_{ij}[t]$ is the velocity of particle in dimension *j* at time *t, gbest* is the best position obtained by the flock, *pbest* is the best position obtained by particle *i*, *c1* and *c2* are constants determined in step 1 and *r1* and *r2* are random floating point values drawn from a uniform distribution with bounds [0,1].
5. Evaluate new particle positions
6. Assign *gbest* and *pbest* positions for global best and local best flock positions respectively.
7. Iterate until stopping criterion reached (in this case, number of function evaluations).

Through this method, a wide variety of behaviours for each particle type can be achieved. For instance, the following four extremes of behaviour can be generated thus:

1. **Low c1 and low c2.** Small movements are made over the search space in each iteration. This might represent a cautious individual in the natural world.
2. **High c1 and low c2**. Particles of this type will move close to their previous known best positions and the social influence of the global best will be small. This might represent an introverted individual in the natural world.
3. **Low c1 and high c2.** These individuals will effectively hillclimb towards the current best position and are the most socially motivated. This might represent gregarious individuals in the natural world.
4. **High c1 and high c2.** These individuals will make large movements across the search space in each iteration with roughly equivalent weighting to gbest and pbest. This might represent a risk-taking individual in the natural world.

Although these four behaviours represent the extremes of behaviour, the use of a Gaussian distribution ensures that actual particle behaviours will reside on a continuum between these four points, as opposed to discrete behaviours at the end of each spectrum. This method has clear links to the underlying biology of organisms in that modern genetics (e.g. genome-wide association studies) identifies traits in behaviour for complex organisms rather than absolute behaviour types.

Experimental Setup

The CBV-PSO algorithm was coded in C++ and compared with the same PSO set-up with no behavioural variation. The algorithm was run for 1000 generations with a swarm size of 50, and for 30 repeated trials to account for the effect of the random seed. Runs were conducted on a modern machine with a Core i5 - 2.5GHz CPU and 8Gb RAM.

CBV-PSO has been tested on four test functions taken from the literature. These are:

Absolute Value

$$f(x) = \sum_{j=1}^{n_x} |x_j|$$

Spherical

$$f(x) = \sum_{j=1}^{n_x} x_j^2$$

Griewank

$$f(x) = 1 + \frac{1}{4000}\sum_{j=1}^{n_x} x_j^2 - \prod_{j=1}^{n_x} \cos\left(\frac{x_j}{\sqrt{j}}\right)$$

Ackley

$$f(x) = -20e^{-0.2\sqrt{\frac{1}{n_x}\sum_{j=1}^{n_x} x_j^2}} - e^{\frac{1}{n_x}\sum_{j=1}^{n_x}\cos(2\pi x_j)} + 20 + e$$

An additional aspect of the PSO algorithm investigated here is the scalability of the approach to larger dimensional problems. Therefore experiments were conducted on problem sizes from 10 to 100 dimensions and are explored in the results below.

Finally, the extent to which individuals vary within a population is important and so the mean and standard deviation of the Gaussian applied to c1 and c2 is modified to determine the effect of different levels of variation on the population. The standard PSO operates with fixed c1 and c2 coefficients of 2.0.

Results

The following tables of results show the performance of the best performing individual for each algorithm and for the four test problems. The best result for each number of dimensions is shown in bold. Not all dimensions are shown in these tables for reasons of brevity.

Table 1 – Absolute Value Function, Mean Best Solution

No. of Dimensions	Standard PSO	CBV-PSO (Mean 2.0, SD 1.0)	CBV-PSO (Mean 1.0, SD 1.0)	CBV-PSO (Mean 2.0, SD 0.5)
10	9.42E-15	6.23E-15	7.72E-23	**2.68E-31**
20	7.30E-07	2.08E-05	8.81E-09	**3.91E-11**
30	0.00055294	0.00600296	5.43E-05	**1.55E-06**
40	0.0248816	0.0565303	0.00210781	**0.00020928**
50	0.163591	0.373818	0.0363336	**0.00634206**
60	0.57161	0.744932	0.170024	**0.0997118**
70	1.77105	2.85459	0.622282	**0.180297**
80	3.30997	4.75152	2.32416	**1.05955**
90	5.97625	9.64508	5.13261	**1.89295**
100	9.47515	16.092	7.08589	**4.06602**

Table 2 – Spherical Function, Mean Best Solution

No. of Dimensions	Standard PSO	CBV-PSO (Mean 2.0, SD 1.0)	CBV-PSO (Mean 1.0, SD 1.0)	CBV-PSO (Mean 2.0, SD 0.5)
10	9.40E-26	5.31E-35	**2.12E-52**	2.91E-48
20	5.21E-11	6.18E-16	**5.74E-21**	4.06E-18
30	1.74E-06	3.94E-10	**1.01E-13**	1.03E-11
40	0.00089097	4.11E-05	**1.13E-07**	1.02E-06
50	0.0185017	0.00015653	0.00017071	**4.89E-05**
60	0.114694	0.00436014	0.00325725	**0.0028324**
70	0.490711	0.0327966	0.0737765	**0.0200073**
80	1.53618	0.172011	0.258165	**0.163199**
90	3.4506	0.604844	1.38218	**0.329333**
100	6.7252	1.94237	2.52391	**1.15571**

Table 3 – Griewank Function, Mean Best Solution

No. of Dimensions	Standard PSO	CBV-PSO (Mean 2.0, SD 1.0)	CBV-PSO (Mean 1.0, SD 1.0)	CBV-PSO (Mean 2.0, SD 0.5)
10	0.163126	**0.106684**	0.154858	0.12956
20	0.0146919	**0.0103463**	0.0114139	0.016249

30	0.0103463	0.00856124	**0.00752455**	0.00853706
40	**0.00462193**	0.0051753	0.0096526	0.005418
50	**0.00472468**	0.0064188	0.00923136	0.00484678
60	**0.00596461**	0.00648561	0.00738712	0.00665521
70	0.0131454	0.00890112	0.0142913	**0.00551836**
80	0.0275156	0.0173526	0.027529	**0.00733312**
90	0.0508129	0.019459	0.0431736	**0.0119852**
100	0.0822002	0.0465551	0.0563486	**0.0201065**

Table 4 – Ackley Function, Mean Best Solution

No. of Dimensions	**Standard PSO**	**CBV-PSO (Mean 2.0, SD 1.0)**	**CBV-PSO (Mean 1.0, SD 1.0)**	**CBV-PSO (Mean 2.0, SD 0.5)**
10	1.12E-13	3.40E-18	3.05E-18	**2.01E-18**
20	3.76E-06	0.00034368	0.148285	**3.80E-10**
30	**0.00122449**	0.565017	1.23208	0.268053
40	**0.07018**	1.67546	2.27732	1.18875
50	**0.286271**	2.07832	2.75871	1.93871
60	**0.879147**	2.53535	3.21985	2.35003
70	**1.27674**	3.19615	3.90879	2.7222
80	**1.85029**	3.50581	4.25278	2.98451
90	**2.1071**	3.92109	4.71472	3.44805
100	**2.49497**	4.17211	5.51884	3.58679

Clearly the CBV-PSO improves the search procedure at low dimensionalities for all problems. However, the performance is superseded by the standard PSO for high dimensionalties on the simpler problems and for a larger part of the Ackley function.

To better determine performance, a comparison across all problems and dimensionalities would be beneficial. However, determining the mean performance for all dimensionalities is difficult as the results would clearly be skewed to the larger values present in higher dimensionalities. Therefore in the following plots, a ranking system is used. Each algorithm is ranked from 1 (the best) to 4 (the worst) for each problem type and dimensionality. The mean ranks are then used to determine performance for the algorithms. This is performed for the average performance of algorithm (Figure 1) and the best solutions generated across all runs (Figure 2).

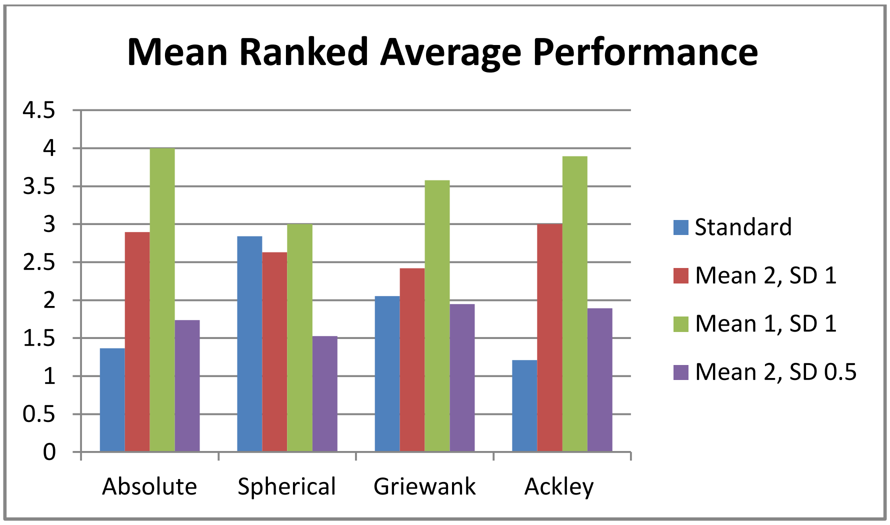

Figure 1 – Mean Ranked average best performance for all algorithms

Figure 1 shows the mean rank of the average performance for all algorithms across dimensionalities and from this it is clear that the CBV-PSO with a mean of 2 and a standard deviation of 0.5 is superior to standard PSO on the Spherical and Griewank functions and competitive on the remaining functions. CBV-PSO with a mean and standard deviation of 1 is universally poor, indicating perhaps that there are too many individuals with small c1 and c2 values leading to poor exploration of the search space. The results for the CBV-PSO, Mean 2, SD 1 are inbetween these two results.

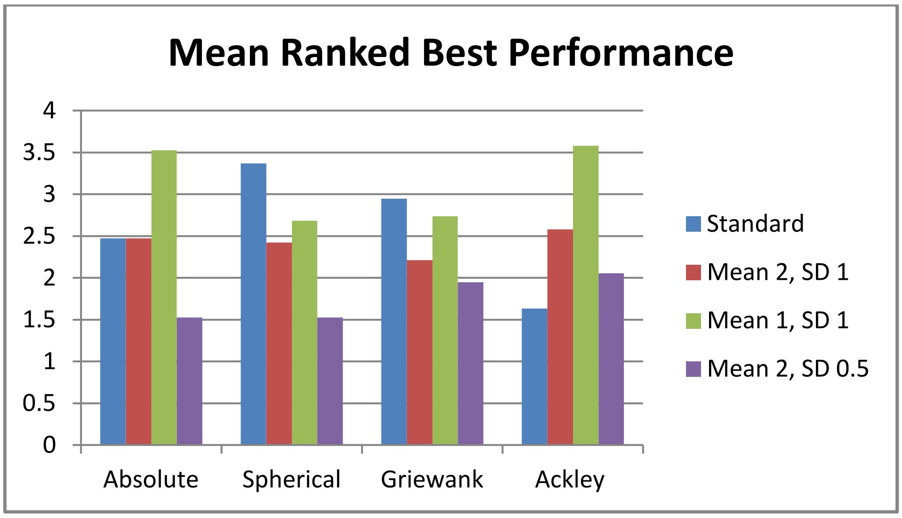

Figure 2 – Mean ranked best performance for all algorithms

Figure 2 shows that the best performing algorithm on three of the four problems is CBV-PSO with a mean of 2 and standard deviation of 0.5. It is apparent that the standard PSO performs very poorly on this function for lower dimensionalities, leading to a number of '4' ranked solutions in this category. This is perhaps best shown with a radar plot as shown in Figure 3.

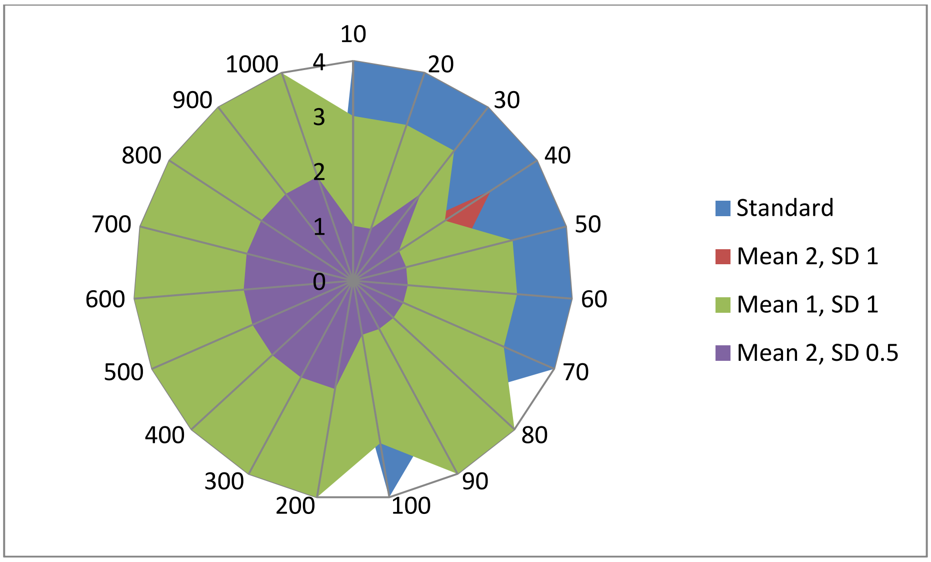

Figure 3 – Radar plot of ranks for all algorithms for the Absolute function (dimensionality is indicated around the circumference of the plot, outer ranks are worse).

Figure 3 shows that the standard PSO performs comparatively badly up to 70-100 dimensions whereas the CBV-PSO-2-0.5 performs well at this point and never has a rank worse than 2 for all dimensionalities. The CBV-PSO could be said to be more robust over problem dimensionalities for this problem type.

Comparison with Previous Work

The direct competitor to this algorithm is the discrete heterogeneous algorithm (HPSO) presented in [6]. Comparisons are made with the static heterogeneous PSO described in this paper as this is the direct competitor to CBV-PSO. The dynamic DHPSO has additional online learning of behaviour. From Table 4 and 5 it would appear that the standard PSO and CBV-PSO (Mean 2, SD 0.5) implementation used here improves on the performance demonstrated by HPSO on both functions.

Table 5 – Comparison with HPSO on Griewank Function

Dimensions	**HPSO**	**CBV-PSO-2-0.5**
10	**0.0782**	0.130
30	0.0407	**0.0085**
50	0.154	**0.0048**

100	3.61	**0.02**

Table 6 – Comparison with HPSO on Ackley Function

Dimensions	HPSO	CBV-PSO-2-0.5
10	3.99E-15	**2.01E-18**
30	1.20	**0.27**
50	2.87	**1.94**
100	**2.87**	3.58

Conclusion

A continuous behavioural variation PSO has been developed and tested on 4 optimisation problems taken from the literature. The algorithm has been shown to outperform a standard PSO formulation on these problems, particularly for lower dimensional versions of these problems. A broad comparison with HPSO is not conclusive, but does appear to show that the continuous approach has benefits over the discrete behavioural selection approach. A more in-depth comparison of these approaches is now warranted.

A feature of this algorithm is that the mean performance is not affected as much as the best performance of the algorithm which suggests that the effect of producing a population with a variety of characteristics does aid the performance of one individual, as is shown in real flocks [2]. The variation of the mean and standard deviation of the c1 and c2 coefficients clearly has an effect on the performance of the algorithm. In the experiments conducted here, it is apparent that a larger mean and smaller standard deviation produce the best results. It is hypothesised that the larger standard deviations and smaller means lead to a greater number of individuals that have small c1 and c2 values, producing slow-moving populations that cannot discover good solutions in the 1000 generation timeframe. This is particularly the case where the mean and standard deviation of 1.0 are used.

It is not particularly clear why the beneficial effect of continuous heterogeneity should be so apparent at small dimensions and break down in a number of cases at higher dimensions. One hypothesis might be that for the larger problems, there is less of an opportunity for the c1 and c2 parameters to influence the search, but this does not explain why performance would dip below that of the standard PSO.

Further Work

A further, more comprehensive comparison should be made with the discrete variation systems proposed by Engelbrecht and a further investigation is planned into the settings of c1 and c2 and the extent of behavioural variation within species.

Acknowledgements

This work was funded by the “Bridging the Gaps: Exeter Science Exchange” project funded by the EPSRC (EP/I001433/1 – PI Prof. David Butler).

References

Bibliography

[1] D. Croft, J. Krause, S. Darden, I. Ramnarine, J. Faria and R. James, "Behavioural trait assortment in a social network: patterns and implications," *Behavioral Ecology and Sociobiology,* vol. 63, pp. 1495-1503, 2009.

[2] J. Dyer, D. Croft, L. Morrell and K. J, "Shoal composition determines foraging success in the guppy," *Behavioral Ecology,* pp. 165-171, 2009.

[3] M. Montes de Oca, J. Pena, T. Stutzle, C. Pinciroli and M. Dorigo, "Heterogeneous Particle Swarm Optimizers," IRIDIA – Technical Report Series, Brussels, 2009.

[4] M. Iwamatsu, "Multi-Species Particle Swarm Optimizer for Multimodal Function Optimization," *IEICE TRANSACTIONS on Information and System,* vol. E89, no. 3, pp. 1181-1187, 2006.

[5] C.-k. Chow and H.-t. Tsui, "Autonomous agent response learning by a multi-species particle swarm optimization," in *Congress on Evolutionary Computation. CEC2004.* , 2004.

[6] A. Engelbrecht, "Heterogeneous Particle Swarm Optimisation," in *ANTS 2010*, Brussels, 2010.

[7] A. Engelbrecht, "Scalability of A Heterogeneous Particle Swarm," in *IEEE SSCI 2011*, Paris, France, 2011.