



Analysing the impact of dimensionality on diversity in a multi-layered Genotype-Phenotype mapped genetic algorithm

Title	Analysing the impact of dimensionality on diversity in a multi-layered Genotype-Phenotype mapped genetic algorithm
Author(s)	Hill, Seamus;O'Riordan, Colm
Publication Date	2013-06-20
Publisher	IEEE
Repository DOI	10.1109/CEC.2013.6557589

Analysing the Impact of Dimensionality on Diversity in a Multi-layered Genotype-Phenotype mapped Genetic Algorithm

Seamus Hill

College of Engineering and Informatics
National University of Ireland Galway
Ireland
Email: seamus.hill@nuigalway.ie

Colm O’Riordan

College of Engineering and Informatics
National University of Ireland Galway
Ireland
Email: colm.oriordan@nuigalway.ie

Abstract—This paper examines the impact of changes in dimensionality on a multi-layered genotype-phenotype mapped GA. To gain an understanding of the impact we carry out a series of experiments on a number of well understood problems and compare the performance of a simple GA (SGA) to that of a multi-layered GA (MGA) to demonstrate their ability to search landscapes with varying degrees of difficulty due to changes in the dimensionality of each function. The paper also examines the impact of diversity maintenance in assisting the search and identifies the natural increase in diversity as the level of problem difficulty increases, as a result of the layered Genotype-Phenotype mapping. Initial results indicate that it may be advantageous to include a multi-layered genotype-phenotype mapping under certain circumstances.

I. INTRODUCTION

Genetic Algorithms (GAs) [1], [2], are search algorithms based on the Darwinian principal of the survival of the fittest. An initial population of individuals is created, each representing a possible solution to a given problem. The individuals in the population are allocated a fitness score based on their suitability in the environment in which they exist. The individuals are then subjected to environmental pressure and based on their fitness scores, with the propagation of fitter individuals through natural selection. However, Sewall Wright noticed that many random changes in the frequency of alleles occurring in a population were not related to selection [3]. His observations indicated that this *genetic drift* was an important component in the evolutionary process. *Neutral theory*, as proposed by Kimura [4] offered an alternative to the Darwinian view, states that the mutations involved in the evolutionary process are neither advantageous nor disadvantageous to the survival of an individual, and that most mutations are caused not by selection but rather by random genetic drift. However, Kimura pointed out that although natural selection does play a role in adaptive evolution, only a tiny fraction of DNA changes are adaptive. The vast bulk of mutations are phenotypically silent [5].

By adopting the ideas of Darwinism, simple genetic algorithms (SGAs) can be viewed as implementing the process of evolution without containing any explicit neutral mutations. In

other words, mutations are either an advantage or a disadvantage to the individual in terms of fitness, with selection then propagating the fitter individuals. As the search progresses exploration and exploitation ratios decrease as the population converges. If we are to implement a genetic algorithm (GA) based on the principles of neutral theory then neutrality needs to be introduced. Neutrality can be viewed as a situation where a number of different genotypes can represent the same phenotype.

The methodology used to examine the impact of neutrality is as follows: we ran a number of experiments over the Sphere Model function and a Sphere Model function which incorporates a changing environment. These problems were chosen as they allow us to examine the efficiency of both the SGA and the MGA over the Sphere Model and then to compare their performance when a changing environment is included. To examine how the GAs perform over varying levels of difficulty we alter the number of dimensions associated with each of the problems. In this paper the authors examine the impact of changes in dimensionality of a landscape has on a GA’s ability to optimize a problem. We also examine how the maintenance of population diversity can impact on a GA’s ability to solve a problem and how the inclusion of neutrality assists in maintaining this diversity in a natural way, without having to keep track of the population dynamics. Finally, we used a two sided paired Wilcoxon test to analyse the results. The paper is laid out as follows: Section II describes the background and outlines the motivation for conducting this study. Section III gives an overview of the multi-layered GA (MGA), while Section IV examines the test suite used in the experiments. Section V outlines the experimental results and finally we conclude and outline possible future work in Section VI.

II. BACKGROUND

In GAs an object which forms a solution to a particular function can be referred to as a phenotype and a phenotype’s encoding within a GA is known as the genotype [6]. Both the genotype and the phenotype have their own space, which

can differ significantly. The fitness landscape for a genetic algorithm can be viewed as the visualisation of the relationship which exists between the variation operators, the candidate solutions (elements of the function's domain) and their objective or fitness values [7].

The motivation for the research described in this paper, is to gain an understanding of the relationship (if any), which exists between the level of difficulty associated with a problem and the presence within the GP-map of a level of neutrality. In order to analyse this we have selected a unimodal problem and a changing landscape problem as part of our test suite. In order to create an experimental environment which allows us to alter the levels of difficulty associated with each problem, we will examine the impact of changing the level of dimensionality associated with each of the problems. By adopting this approach as our methodology we can run a number of experiments with varying level of dimensionality, which alter the levels of difficulty. We have chosen to use 3, 15 and 30 as the number of dimensions in order to create low, medium and high levels of dimensionality.

The contribution of this paper is to implement a basic interpretation of the biological concepts of transcription and translation into the GP-map of a GA in order to introduce neutrality and to use an interpretation of a missense mutation operator in one of the GP-mapping layers, which reflects back onto the genotype. The objective of the research is to use this mechanism to naturally maintain a level of diversity within the population without having to examine the ongoing population dynamics. The paper examines how the diversity maintained in the population through the use of the GP-mapping impacts on the search over problem landscapes of different levels of difficulty.

A. Neutrality

Neutrality can be defined as a situation where following a mutation one genotype changes to another genotype, but both genotypes represent the same phenotype [4]. This implies that as neutrality is introduced, the solution space increases without increasing the genotype space. Neutral representations have appeared in a number of GAs over the past number of years. As a general rule, the introduction of neutrality into GAs can be divided into two categories. Firstly, fitness landscapes which introduce neutrality i.e. Kauffman's NK landscape [8], Barnett's NK_p landscape [9], Newman and Engelhardt's NK_q landscape [10]; and Beaudoin et al.'s ND landscape [11]. The second category, which is the focus of this paper, is the introduction of neutrality through genotype-phenotype mappings (GP-map). With this approach, neutrality is obtained implicitly rather than explicitly. Shipman [12] found neutrality to be advantageous where neutral networks (introduced by Harvey and Thompson [13] - meaning points in a search space of equal fitness), are distributed over the search space with a high degree of connectivity between them. Shaktin [14] and Shipman [15] showed that neutrality could be introduced through the use of GP-maps. Their approach of using mappings was extended by Ebner et al. [16], [17]

and outlined how high levels of mutation could be sustained by having neutral networks present. They also identified that neutral networks assist in maintaining diversity in the population, which may be advantageous in a changing environment.

B. Diversity

Premature convergence is often cited as a problem for GAs; it occurs when the population reaches a sub-optimal point and the genetic operators can no longer create offspring with higher fitness levels than their parents. This is caused by a loss of diversity in the population [18] [19]. As described by Whitley [20] and Banzhaf, [21] in order to continue to explore regions of the search space you need to maintain diversity in the population. Many methods have been proposed to maintain diversity in the population Grefenstette introduced a "*partial hypermutation*" step which was used to replace a percentage of the population ("*the replacement rate*") with randomly generated individuals. The purpose of this was to maintain enough diversity in the population to allow exploration of the search space [22]. He found this to be quite useful in changing environments. Cobb and Grefenstete used "*Random Immigrants*" to replace part of the population in each generation and also an adaptive approach called "*Triggered Hypermutation*" which increased mutation when decrease in the performance of the GA was detected. The results indicated that "*diversity represented a natural source of power in adopting to changing environments*" [23]. Again this was used in conjunction with a changing environment and advantages and disadvantages were detected. For a standard GA, high mutation rates were useful for tracking performance that changed continuously, when looking at the offline performance. However, online performance deteriorated when high rates of mutation (0.10) were present. With the Triggered Hypermutation approach, as it is adaptive, the level of diversity was increased when needed but it didn't perform well in abruptly changing environments. The Random Immigrants operated quite well when the environment was changing, but there was an overhead when the environment was stationary [23]. A micro-GA was proposed by Krishnakumar [24] which attempted to increase diversity in the population to compensate for a change detected in the environment. But the populations were small and converged quickly so constant replacement of the population was required. Bickle and Thiel defined loss of diversity to be the proportion of individuals of a population that are not selected during the selection phase. They examined fitness distributions to gain a better understanding of selection schemes and carried out a comparison of various selection schemes [25]. Motoki [26] found some results differed with that of Bickle and Thiel, although they both agreed that loss of diversity is fundamentally related to selection pressure and the arrangement of genes before selection. Motoki found tournament, and exponential ranking schemes are roughly equivalent and that with linear ranking selection the speed of evolution will only vary slightly.

When examining diversity with a SGA, genotypic and phenotypic diversity can be viewed as one and the same.

However, with the MGA, due to the nature of the GP-map, genotypic and phenotypic diversity are viewed in isolation as there exists a many-to-one relationship between the genotype and the phenotype. In this paper we measure both genotypic diversity (G) and phenotypic diversity (P) because of the many-to-one mapping which exists between G and P for the MGA. We calculate G and P using Hamming Average as outlined in [27]. In theory the dynamics of the MGA population should be quite different to that of the SGA because of the relationship between G and P . With a traditional GA the initial hamming distance (h) is $l/2$ and moves towards $h = 0$ [27]; the MGA may differ in diversity at G level.

III. MULTI-LAYERED GP-MAP

The primary inspiration for the multi-layered GA can be found in the biological processes of *transcription* and *translation*. At a very basic level, the biological process of transcription involves the copying of information stored in DNA into an RNA molecule, which is complementary to one strand of the DNA. The process of translation then converts the RNA, using a predefined translation table, to manufacture proteins by joining amino acids. These proteins can be viewed as a manifestation of the genetic code contained within DNA and act as organic catalysts in anatomy.

The multi-layered GA (see Figure 1) includes a layered genotype-phenotype map which adopts a basic interpretation of the transcription and translation processes and allows for a basic implementation of a missense mutation operator. The genotype consists of “1”s and “0”s and the first stage of transcription is to convert the genotype into a string of characters from the alphabet $A, C, G,$ and T (which attempts to represent the biological concept of a template strand). To achieve this we use the following mappings; “00” represents A , “01” represents C , “10” represents G and “11” represents T . Following this we create an interpretation of a coding strand where $A \rightarrow T, C \rightarrow G, G \rightarrow C$ and $T \rightarrow A$. The final phase of the transcription stage maps $T \rightarrow U, G \rightarrow G, C \rightarrow C$ and $A \rightarrow A$ and creates an RNA sequence. The translation stage compares the RNA sequence to a translation table which is randomly generated at initialisation to create a mapping from the RNA sequence into a series of basic interpretations of amino acids, called *phenes*, which are then combined to create the phenotype.

By including a basic interpretation of transcription and translation into the genotype-phenotype mapping we introduce neutrality, which allows for a many-to-one relationship between the genotype and the phenotype. The missense mutation operator is interpreted and implemented as follows. Once the process of *transcription* has taken place, missense mutation occurs at a given rate of probability. If a missense mutation takes place then one of the RNA bases, is flipped to another and the *translation* stage is carried out. It should be noted that unlike traditional point mutation operators, the probability of a missense mutation taking place is on each single codon, which

in this case is a collection of four characters. Following this the translation phase takes place. For a more detailed overview of the MGA see [28].




Fig. 1: Overview of Multi-layered GA (MGA)

IV. TEST SUITE

In this paper we compare the performance of a SGA and a MGA over a number of well known problems. In order to allow for a comparable performance we have used the following values for each of the experiments; crossover rate $P_c = 0.70$, mutation Rate $P_m = 1/l$, where l is the length of the chromosome. We also use a tournament selection scheme with a tournament size of 4. Each of the experiments has a population size of 400 and the Sphere Model experiments run for 2000 generations, the changing environment experiments run for 4000. The probability of missense mutation P_{mm} was set to 0.03, this figure was chosen as being suitable, having carried out a number of trial runs to briefly examine its impact and kept constant throughout the search.

A. Sphere Model [29]

The Sphere Model [29] is relatively easy to optimise as it is continuous, convex and unimodal. This function is normally used to measure the efficiency of a particular algorithm. For this paper the authors altered the number of dimensions (n) associated with the Sphere Model in order to vary the level of difficulty and to examine the impact of dimensionality on the algorithms. Experiments on the Sphere Model were carried out with $n = 3, n = 15$ and $n = 30$. The details of the Sphere Model are as follows;

$$f_1(\vec{x}) = \sum_{i=1}^2 x_i^2$$

$$\text{Limits} - 5.12 \leq x_i \leq 5.11$$

$$\min(f_1) = f_1(0, \dots, 0) = 0$$

B. Sphere model, changing environment [30]

The changing environment experiments have at its heart the sphere model as outlined in [29], however the idea is to allow the GAs to search the landscape defined by the sphere model and to then change the function values after 1500 generations, so that the landscape also changes. The aim of this set of experiments is to examine how both the SGA and MGA cope in a changing landscape environment with the dimensionality again being, $n = 3$, $n = 15$ and $n = 30$. The details of the Sphere Model Changing Environment are as follows;

$$f_2(x(t)) = \begin{cases} \sum_{i=1}^n x_i^2(t) & : t \bmod a \text{ even} \\ \sum_{i=1}^n (x_i - b)^2 & : t \bmod a \text{ odd} \end{cases}$$

$$-5.12 \leq x_i \leq 5.11$$

$$a = 1500 \text{ generations} \quad ; \quad b = 4$$

$$\min(f_2) = \begin{cases} f_2(0, \dots, 0) & : t \bmod a \text{ even} \\ f_2(b, \dots, b) & : t \bmod a \text{ odd} \end{cases} = 0$$

V. EXPERIMENT FINDINGS

A. Sphere model experiments

The first experiment was carried out using the Sphere Model and the results of the experiments are outlined in Table I. The table illustrates the percentage of times the global optimum is located and the average number of runs required to locate the global optimum, averaged over the number of successful runs over the varying levels of dimensions. The figures appear to indicate that the number of dimensions contained in the problem has an impact on the performance of the algorithms over the Sphere Model function.

When the number of dimensions n is set to 3 the problem is extremely easy for the SGA and the MGA, with both achieving 100% success in locating the global optimum. The SGA locates the optimum on average after only 5 generations and the MGA takes on average 22 generations. One possible reason for this is that an adequate level of diversity in the population exists early in the search and that due to the relative ease of the problem with $n = 3$, the SGA on average, locates the global before the MGA. The level of diversity in the population for both the SGA and the MGA is shown in Figure 2. It should be noted that G and P for the SGA are one and the same, whereas G and P for the MGA differ due to the nature of the GP-map, and this is shown where G for the MGA differs significantly from P . Also keeping track of the off-line (average best) performance and on-line (average) performance is useful as an indicator of whether the balance between exploration and exploitation is being maintained. The recorded off-line and on-line performances of the SGA and the MGA were quite similar as the low degree of dimensionality meant that the problem was relatively easy for both GAs and the global optimum was located very quickly.

F1 - Sphere Model Experiments		
Number of Dimensions $n = 3$		
GA Description	SGA	MGA
Optimum Located	100%	100%
Avg. No. Generations Required	5	22
Number of Dimensions $n = 15$		
GA Description	SGA	MGA
Optimum Located	100%	100%
Avg. No. Generations Required	1014	359
Number of Dimensions $n = 30$		
GA Description	SGA	MGA
Optimum Located	80%	100%
Avg. No. Generations Required	1681	1007

TABLE I: F1: Sphere Model Experiments




Fig. 2: Sphere Model Population Diversity $n = 3$ - SGA & MGA

When the number of dimensions increases to $n = 15$, the level of difficulty increases for both GAs, as shown in Table I. The level of difficulty can be seen in the average number of generations taken to locate the global optimum. In this experiment, both the GAs obtained a 100% success rate. However, the SGA needed an average of 1014, while the MGA needed an average of 359 generations, illustrating that the increase in dimensionality has increased the level of difficulty for both GAs, but the effect has been more pronounced for the SGA. This may be due to the lack of diversity within the population as illustrated outlined in Figure 3. When we examine Figure 3 we can see that the population for the SGA converges quite early in the search and this may account for the average number of generations required to locate the global optimum as the level of difficulty has increased due

to the increase in the number of dimensions, but the level of diversity within the population decreases as the search progresses through the generations. We also noted that as the level of difficulty increases, so too does G for the MGA, this illustrates the natural introduction of diversity as difficulty is increased. With regard to the off-line and on-line performance of both the SGA and the MGA they were relatively similar when $n = 15$ and that the balance between exploration and exploitation is maintained.

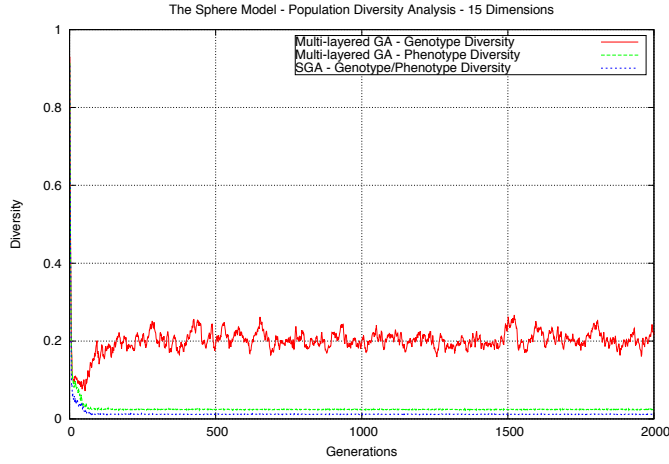


Fig. 3: Sphere Model Population Diversity $n = 15$ - SGA & MGA

In the final Sphere Model experiment, the number of dimensions were increased to $n = 30$. The effect of this can be seen in Table I and Figure 4. This increase in the number of dimensions has had a significant impact on both of the GAs, but again the SGA’s performance shows the largest drop in performance. The SGA only succeeded in locating the global optimum 80% of the time, while the MGA continues with a 100% success rate. The other interesting result here can be found in the average number of generations required to locate the optimum, with the SGA requiring, on average 1681 generations while the MGA needed, on average, 1007 generations. This is a significant improvement in performance on the part of the MGA when compared to that of the SGA. The results indicate that the search of a landscape with increased dimensionality has increased, may be assisted by maintaining an element of diversity in the population.

Figure 4 may indicate that this improvement in performance is due to the implicit level of diversity being maintained in the population of the MGA. The Figure illustrates that for the SGA, diversity decreases relatively quickly. With the MGA although the level of P diminishes, G is significantly higher for the MGA. This implicitly-maintained diversity is a direct result of the GP-map and the many-to-one relationship which exists between the genotype and the phenotype in the MGA.

B. Sphere model, changing environment experiments

The second set of experiments were carried out in a changing environment with the results outlined in Table II.

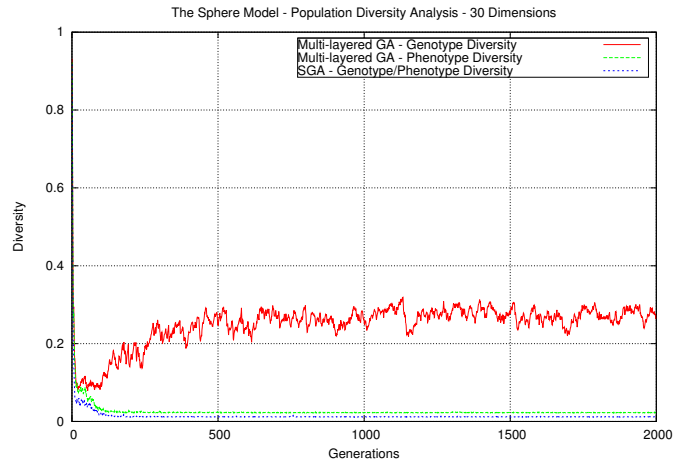


Fig. 4: Sphere Model Population Diversity $n = 30$ - SGA & MGA

We can see that when the dimension level n is set to 3, the problem landscape is relatively easy for both the SGA and the MGA, with both algorithms succeeding 100% of the time. However, the MGA discovers the global optimum in the changing landscape in an average of 1548 generations which is a significant improvement in performance over the SGA, with an average of 2522 generations. Figure 5 illustrates the population diversity for both the SGA and the MGA where $n = 3$. The diversity within the population is relatively similar to that of the first Sphere Model experiment where $n = 3$, but main difference is that when the environment changes the diversity in the population also changes as both algorithms attempt to locate the new global optimum.

F2 - Sphere Model Changing Environment Experiments		
Number of Dimensions $n = 3$		
GA Description	<i>SGA</i>	<i>MGA</i>
Optimum Located	100%	100%
Avg. No. Generations Required	2522	1543
Number of Dimensions $n = 15$		
GA Description	<i>SGA</i>	<i>MGA</i>
Optimum Located	20%	100%
Avg. No. Generations Required	2939	2182
Number of Dimensions $n = 30$		
GA Description	<i>SGA</i>	<i>MGA</i>
Optimum Located	0%	100%
Avg. No. Generations Required	<i>N/A</i>	3557

TABLE II: F2: Sphere Model Changing Environment Experiments




Fig. 5: Sphere Model Changing Environment Population Diversity $n = 3$ - SGA & MGA

The impact of the changing environment becomes noticeable when we examine the on-line and off-line performance of the SGA in Figure 6. Although we saw that the SGA had little or no difficulty solving the Sphere Model problem when $n = 3$, it is apparent that the change in the environment has had little effect the on-line results of both the SGA and the MGA 7, this may be because of the relative ease of the problem.

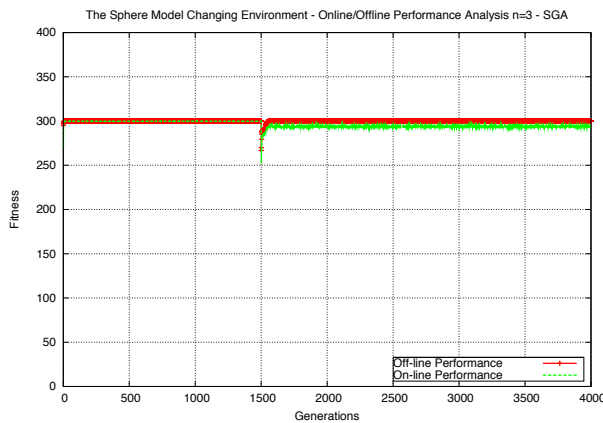


Fig. 6: Sphere Model Changing Environment Online & Offline Performance $n = 3$ - SGA

In the second set of changing landscapes experiments we set the number of dimensions to $n = 15$. We can see from Table II, that the SGA is finding it difficult to cope with the changing environment when problem difficulty increases due to the increase in the level of dimensionality. The SGA only manages to succeed in locating the global optimum 20% of the time and when it did locate it the average number of generations required was 2939. The MGA, on the other hand, was able to locate the global optimum 100% of the time and



Fig. 7: Sphere Model Changing Environment Online & Offline Performance $n = 3$ - MGA

the average number of generations required was 2182. Again the population diversity as illustrated in Figure 8, contains a similar pattern to that of the Sphere Model experiment where $n = 15$. The trend of the difference between the G and P increasing as difficulty increases continues and the change in diversity is magnified as the dimensionality has increased from 3 to 15.




Fig. 8: Sphere Model Changing Environment Population Diversity $n = 15$ - SGA & MGA

As the number of dimensions increase, the impact on the on-line and off-line values for both of the SGA and the MGA is significant (Figures 9 and 10 respectively). These results illustrate that at the beginning of the search both the on-line and off-line performance for the SGA and MGA are quite similar, but once the environment changes the result on the on-line performance are more noticeable. The on-line fitness level of the MGA falls below that of the SGA, but it still maintains the ability to search the space more effectively than the SGA. As the parameters are held constant

the only difference between the performance before the change at generation 1500 and the rest of the search is the greater presence of diversity in the population of the MGA due to neutrality in the GA-map.



Fig. 9: Sphere Model Changing Environment Online & Offline Performance $n = 15$ - SGA




Fig. 10: Sphere Model Changing Environment Online & Offline Performance $n = 15$ - MGA

For the final changing environment experiment the level of dimensionality n was increased to 30. As this increased the level of difficulty we can see that the SGA was unable to locate the global optimum. The MGA, however, did locate the new global optimum 100% of the time, with an average number of generations of 3189 required (see Table II). When we examine the population diversity shown in Figure 11 we can see a similar pattern to the previous experiments, with the additional feature of the time required for the population diversity to revert back to normal levels after the landscape changing, increasing due to the increase in dimensionality.




Fig. 11: Sphere Model Changing Environment Population Diversity $n = 30$ - SGA & MGA

For both the SGA (see Figure 12) and the MGA (see 13) the impact of the change in the environment when dimensionality is set to 30 is interesting. The off-line performance for both GAs appears to almost recover to where it was before the change in the environment. However, the impact on the on-line performance has been so severe that the figures have fallen into negative values for both GAs, with the SGA failing to locate the global optimum.

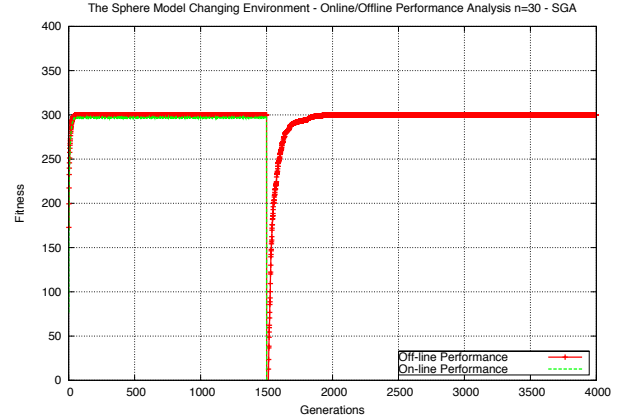


Fig. 12: Sphere Model Changing Environment Online & Offline Performance $n = 30$ - SGA

C. Statistical Analysis

A two sided paired Wilcoxon test was carried out on the results of each experiments to access whether the population means differ. The results were shown to be statistically significant with a P value of $P < 2.2e - 16$.

VI. CONCLUSION & FUTURE WORK

The experiment we conducted in order to determine if a relationship existed between the level of difficulty within a




Fig. 13: Sphere Model Changing Environment Online & Offline Performance $n = 30$ - MGA

problem and the presence within the GP-map of a level of neutrality. The results appear to indicate that the presence of a degree of neutrality within the GP-map maintains a level of diversity within the population and that this diversity can assist in searching the landscape of particular problems. Overall, the results concur with previous research showing the neutrality appears to be beneficial particularly in changing environments. However, what the MGA offers is a novel way to maintain, in a natural way, a level of diversity within the population through the many-to-one mapping between the genotype and the phenotype. Which was illustrated where the difference between G and P increased as the level of difficulty increased. This approach may be useful as there is no requirement to monitor the population dynamics and the readjust the parameter settings.

Future work includes the following:

- Additional experiments on an expanded test suite.
- Further analysis of the use of operators in the various layers of the GP-map.
- Comparison of the MGA to other diversity maintaining methods.

REFERENCES

[1] J. H. Holland, *Adaptation in natural artificial systems*. Ann Arbor: University of Michigan Press, 1975.

[2] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley Publishing Company, Inc., 1989.

[3] S. Wright, "The Role of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," in *Proceedings of the Sixth International Congress on Genetics*, D. Jones, Ed., vol. 1, 1932, pp. 356–366.

[4] M. Kimura, "Evolutionary Rate at the Molecular Level," *Nature*, vol. 217, no. 1, pp. 624–626, 1968.

[5] —, *The Neutral Theory of Molecular Evolution*. Cambridge University Press, 1983.

[6] A. E. Eiben and J. E. Smith, "Introduction to evolutionary computing." Springer, 2003.

[7] M. Mitchell, *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1996.

[8] S. A. Kauffman, *Origins of Order*. Oxford University Press, Oxford., 1993.

[9] L. Barnett, "Ruggedness and Neutrality – the NKp family of Fitness Landscapes," in *Artificial Life VI: Proc. of the Sixth Int. Conf. on Artificial Life*, C. Adami, R. K. Belew, H. Kitano, and C. Taylor, Eds. Cambridge, MA: The MIT Press, 1998, pp. 18–27.

[10] M. E. J. Newman and R. Engelhardt, "Effects of neutral selection on the evolution of molecular species," Santa Fe Institute, Tech. Rep. Working Papers 98-01-001, 1998.

[11] W. Beaudoin, S. Verel, P. Collard, and C. Escazut, "Deceptiveness and neutrality the nd family of fitness landscapes." in *GECCO'06*, 2006, pp. 507–514.

[12] R. Shipman, "Genetic Redundancy: Desirable or Problematic for Evolutionary Adaption," in *Proceedings of the 4th international Conference on Artificial Neural Networks and Genetic Algorithms (ICANNGA '99)*, A. Dobnikar, N. Steele, D. W. Pearson, and R. F. Albrecht, Eds. Berlin: Springer, 1999, pp. 337–344.

[13] I. Harvey and A. Thompson, "Through the labyrinth evolution finds a way: A silicon ridge," in *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware*, vol. 1259. Springer Verlag, 1996, pp. 406–422.

[14] M. A. Shackleton, R. Shipman, and M. Ebner, "An investigation of redundant genotype-phenotype mappings and their role in evolutionary search," in *Proceedings of the International Congress on Evolutionary Computation (CEC 2000)*. IEEE Press, 2000, pp. 493–500.

[15] R. Shipman, M. Shackleton, and I. Harvey, "The use of neutral genotype-phenotype mappings for improved evolutionary search," *BT Technology Journal*, vol. 18, pp. 103–111, October 2000.

[16] M. Ebner, M. Shackleton, and R. Shipman, "How neutral networks influence evolvability," *Complex.*, vol. 7, no. 2, pp. 19–33, 2001.

[17] M. Ebner, P. Langguth, J. Albert, M. Shackleton, and R. Shipman, "On neutral networks and evolvability," in *IEEE Congress on Evolutionary Computation (CEC)*. IEEE Press, 2001.

[18] L. Davis, in *Handbook of Genetic Algorithms*. New York: Van Nostrand Reinhold, 1991.

[19] D. Fogel, "An introduction to simulated evolutionary optimization," *IEEE Transactions on Neural Networks*, vol. 5, no. 1, pp. 3–14, January 1994.

[20] D. Whitley, "The genitor algorithm and selection pressure: Why rank-based allocation of reproductive trials is best." in *Proceedings of the Third International Conference on Genetic Algorithms*, J. D. Schaffer, Ed. San Mateo, CA.: Morgan Kaufmann, 1989, pp. 116–121.

[21] W. B. et al., "Genetic programming an introduction." San Francisco, California.: Morgan Kaufmann, 1988.

[22] J. Grefenstette, "Genetic algorithms for changing environments," in *Parallel Problem Solving from Nature 2*. Elsevier, 1992, pp. 137–144.

[23] J. J. Grefenstette and H. G. Cobb, "Genetic algorithms for tracking changing environments." in *Proc. of the 5th Int. Conf. on Genetic Algorithms and their Applications*. Morgan Kaufmann, 1993, pp. 523–530.

[24] K. Krishnakumar, "Micro-genetic algorithms for stationary and non-stationary function optimization." in *SPIE, Intelligent Control and Adaptive Systems*, 1989, pp. 289–296.

[25] T. Blickle and L. Thiele, "A comparison of selection schemes used in evolutionary algorithms." *Evolutionary Computation*, vol. 4, no. 4, pp. 361–394, 1997.

[26] T. Motoki, "Calculating the expected loss of diversity of selection schemes," *Evol. Comput.*, vol. 10, no. 4, pp. 397–422, Dec. 2002.

[27] S. J. Louis and G. J. E. Rawlins, "Predicting convergence time for genetic algorithms," in *Foundations of Genetic Algorithms 2*. Morgan Kaufmann, 1993, pp. 141–161.

[28] S. Hill and C. O'Riordan, "Examining the use of a non-trivial fixed genotype-phenotype mapping in genetic algorithms to induce phenotypic variability over deceptive uncertain landscapes," in *Proceedings of the 2011 Congress of Evolutionary Computation (CEC 2011)*. New Orleans, USA, 5-8 June 2011.

[29] K. A. De Jong, "An analysis of the behavior of a class of genetic adaptive systems," Ph.D. dissertation, University of Michigan, Ann Arbor, 1975, dissertation Abstracts International 36(10), 5140B; UMI 76-9381.

[30] F. Hoffmeister and T. Bck, "Genetic self-learning," in *Proceedings of the 1st European Conference on Artificial Life*. The MIT Press, 1992, pp. 227–235.