

The Emergence and Distribution of Species in a Gradient-Based Spatially-Structured Evolutionary Algorithm

Grant Dick

Department of Information Science,
University of Otago,
Dunedin, NEW ZEALAND
Email: gdick@infoscience.otago.ac.nz

Presented at SIRC 2007 - The 19th Annual Colloquium of the Spatial Information Research Centre
University of Otago, Dunedin, New Zealand
December 6th-7th 2007

ABSTRACT

The ability to discover and maintain multiple solutions within a single run is a desirable property for evolutionary algorithms. Researchers have previously turned to many biologically-inspired methods for inspiration to produce niching evolutionary algorithms. This paper extends previous work on the Gradient-Based Spatially-Structured Evolutionary Algorithm, which attempts to embody the concept of parapatric speciation within an evolutionary algorithm. Through an comparison of the evolved population with that of an idealised, perfectly proportioned population, we show that the distribution of population members among the niches of a given problem's fitness landscape does not rely on the global properties of the landscape. Rather, the allocation of individuals to peaks relies on the relative values of neighbouring peaks with regard to their spatial relationship in the fitness landscape.

Keywords and phrases: evolutionary algorithms, population structure, niching methods, environmental gradients

1 Introduction

Traditional evolutionary algorithms (EAs) are ideal for locating a single solution in difficult multimodal landscapes. However, a simple EA will typically fail if it is required to identify multiple desirable solutions within a fitness landscape in the course of a single run. This is due to two factors: genetic drift (in the case of equal-valued optima) will converge the population onto a single point in the fitness landscape and selection (in the case of unequal peaks) will adapt the population toward the fittest peak.

Several researchers have taken the view that identifying multiple solutions is analogous to the speciation processes considered for biological populations (Goldberg & Richardson 1987, Mahfoud 1992, Davidor 1991). Speciation models often rely on some form of breeding separation between subpopulations, whether it be through assortative mate preference or by restricted mating through isolation of subpopulations. One particular method, the Gradient-Based Spatially-Structured EA (GBSSEA) draws from previous research into spatially-structured evolutionary algorithms (SSEAs) to embody the parapatric speciation concept within an EA (Dick & Whigham 2006). The GBSSEA imposes a topology on a population to separate individuals into geographically isolated groups (*demes*). This increases inbreeding and the resultant changes in genetic drift promote local genetic divergence in the population (Dick & Whigham 2005). Additionally, the GBSSEA presents a unique set of characteristics at each location in space; this encourages the demes to adapt locally to their environment, so that the overall population evolves to a diverse set of optima.

This paper extends previous work on the GBSSEA and attempts to quantify the and characterise the manner in which the optima presented by a fitness landscape are exploited. In particular, the Chi-Square-Like performance metric (Deb & Goldberg 1989) will be analysed to show that the distribution of individuals among niches depends on the local relationship between the peaks of a fitness landscape. This contrasts with the manner that most niching methods take, in which all peaks within a fitness landscape are considered simultaneously.

The remainder of this paper is structured as follows: §2 describes the GBSSEA approach used in this paper; §3 describes the Chi-Square-Like performance metric and how it measures the deviation of a given population from that of a "perfectly" partitioned population; §4 attempts to characterise the manner in which the GBSSEA

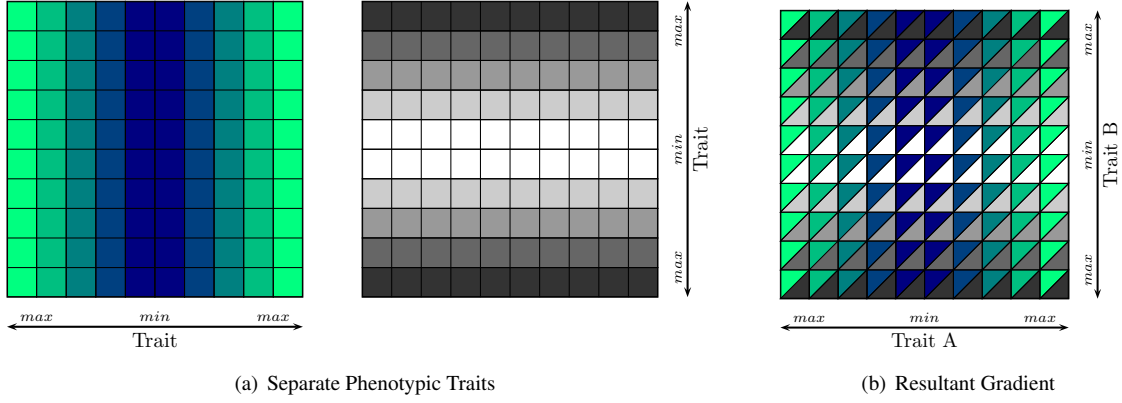


Figure 1: The generation of an environmental gradient for the GBSSEA. Note that all possible combinations of phenotypic trait (within the granularity of space) are produced.

allocates individuals to niches and the conditions that are required to ensure survival of a given optimum; finally, §5 concludes the paper with a summary of the dynamics of the GBSSEA methods’s niche allocation behaviour.

2 Gradient-Based Spatially-Structured Evolutionary Algorithms

Spatial population structure plays an important role in evolutionary algorithms; dividing the population into localised units provides a convenient paradigm for implementing parallel processing. Additionally, one can implement isolation-by-distance style speciation models into an EA through the adoption of certain topological structures (Collins & Jefferson 1991, Davidor 1991). The intent of such spatially-structured evolutionary algorithms (SSEAs) is to promote local divergence of genotypes through restricted mating and local adaptation. However, typical implementations of SSEAs forgo the necessary extensions to promote local adaptation and instead rely on genetic drift to promote divergence. Therefore, a typical SSEA is not a true implementation of the parapatric speciation concept (Mayr 1970).

Several researchers have attempted to promote true parapatric speciation into SSEAs for multiobjective optimisation (Murata, Ishibuchi & Gen 2000, Murata, Ishibuchi & Gen 2001, Kirley 2001). In all cases, local adaptation was implemented through environmental gradients that vary with location in space (Doebeli & Dieckmann 2003). Another method, the Gradient-Based Spatially-Structured EA (GBSSEA) (Dick & Whigham 2006), altered the environmental gradient concept to make it suitable to locating multiple optima in a multimodal fitness landscape. The GBSSEA extends a simple SSEA by including a set of characteristics at each location; these characteristics represent the “ideal” phenotypic traits for a given location. Individuals are compared against these phenotypic traits whenever they attempt to occupy a given location; individuals with high affinity to a location’s traits will be given preference over those who phenotype significantly differs from those presented by the location. This gives the SSEA additional power to ensure that different regions of space actively search alternative parts of the fitness landscape. An example of how a gradient that supports two phenotypic traits might be implemented within the GBSSEA is shown in Figure 1.

The important concept of the GBSSEA that separates it from other SSEA implementations is the concept of an individual’s *local fitness*. The GBSSEA uses a Gaussian function, similar to that proposed by Doebeli and Dieckmann (2003) to determine the *response* of an individual i to a given location l :

$$resp(i, l) = e^{-\frac{1}{2} \left(\frac{\Delta(p_l, p_i)}{\sigma} \right)^2}. \quad (1)$$

Here p_i and p_l represent the phenotypes of the individual and location (respectively). The operator $\Delta(p_l, p_i)$ measures the dissimilarity between the two phenotypes; for real-coded phenotypes this is simply a euclidean distance measure. Finally, σ controls the controls the rate of decay for the environmental response; high values of σ deemphasise the environmental response, while smaller values place a high emphasis on optimising toward the gradient over the problem space. Typically, the value of σ is set to 0.25. The response function is combined with the fitness of the individual with respect to the actual problem being searched to produce the individual local fitness:

$$localFitness(i, l) = f_i \cdot resp(i, l) \quad (2)$$

The algorithm for the GBSSEA is presented in Algorithm 1. One of the major attractions of the GBSSEA approach is its simplicity; other than the calculation of local fitness (lines 12–13 of the algorithm), the GBSSEA does not differ in execution from a typical SSEA.

```

input : A given problem
output: A spatially-structured population of evolved candidate solutions to the problem.

1 population ← {};
2 foreach location in space do
3   | population[location] ← initialIndividual();
4 end
5 while not done do
6   | generation ← {};
7   foreach location in space do
8     | c ← Individual currently at location;
9     | deme ← constructDeme (location);
10    | parents ← select (deme);
11    | o ← breedOffspring (parents);
12    | fo ← evaluate (o) × resp (o, location);
13    | fc ← evaluate (c) × resp (c, location);
14    | if fo > fc then
15      | | generation[location] ← o;
16    | else
17      | | generation[location] ← c;
18    | end
19  | end
20  | population ← generation;
21 end
22 return population;

```

Algorithm 1: The general sequence for a spatially-structured evolutionary algorithm that incorporates an environmental gradient.

3 Expected Niche Proportions – Chi-Square-Like Performance

Deb and Goldberg (1989) introduced the notion that an EA will ideally allocate individuals to the optima in a fitness landscape in proportion to the relative sizes of the peaks. The Chi-Square-Like performance metric was introduced to measure the deviation of a given population from the ideal population; it is defined as the disparity between the observed allocation of individuals to the optima in the fitness landscape and the “perfect” allocation of individuals to peaks in proportion to their value. Under perfect conditions, we would expect a peak i to be allocated μ_i individuals, which is calculated by:

$$\mu_i = N \frac{fitness(i)}{\sum fitness(i)} \quad (3)$$

where N is the size of the population. Similarly, we would expect the variance of the allocation of individuals for a given peak to be:

$$\sigma_i^2 = \mu_i \left(1 - \frac{\mu_i}{N}\right) \quad (4)$$

This paper intends to investigate the Chi-Square-Like performance of the Gradient-Based SSEA. In particular, we focus on two simple one-dimensional test functions, each with five desirable optima (Deb & Goldberg 1989, Mahfoud 1995). These two test functions are ideal for testing such behaviours, as they are completely understood and, despite their simplicity, present difficulty to EAs when the goal is to identify and maintain all five optima. The two test problems, $M1$ and $M2$ are shown graphically in Figure 2 and are defined as:

$$\begin{aligned} M1(x) &= \sin^6(5\pi x) \\ M2(x) &= 2^{-2\left(\frac{x-0.1}{0.8}\right)^2} \cdot \sin^6(5\pi x) \end{aligned}$$




Figure 2: The two test functions used in this paper.

where $0 \leq x \leq 1$. The expected proportions and variances for each peak of $M1$ and $M2$ are listed in Tables 1 and 2, respectively. These values are subsequently used to measure the Chi-Square-Like performance of the Grid-Based SSEA.

	Peak					
	1	2	3	4	5	Non-Peak.
x_i	0.100	0.300	0.500	0.700	0.900	-
$f(x_i)$	1.000	1.000	1.000	1.000	1.000	-
μ_i/N	0.200	0.200	0.200	0.200	0.200	0.000
σ_i^2/N	0.160	0.160	0.160	0.160	0.160	0.800

Table 1: The expected distribution of individuals over the fitness landscape of $M1$.

	Peak					
	1	2	3	4	5	Non-Peak.
x_i	0.100	0.300	0.500	0.700	0.900	-
$f(x_i)$	1.000	0.917	0.707	0.459	0.250	-
μ_i/N	0.300	0.275	0.212	0.138	0.075	0.000
σ_i^2/N	0.210	0.199	0.167	0.119	0.069	0.765

Table 2: The expected distribution of individuals over the fitness landscape of $M2$.

3.1 Empirical Observations

The GBSSEA was run on $M1$ and $M2$ using a population size of 400 (equating to a 20×20 torus)¹. Each run lasted for 500 generations and the tests were repeated 100 times for each problem. In addition to measuring the Chi-Square-Like performance, the number of observed peaks was measured at the end of each generation. For a point of reference, each problem was applied to the equivalent SSEA without an environmental gradient. This allowed for a rough evaluation of the GBSSEA to be made before a more in-depth analysis took place.

The number of peaks discovered and maintained at each generation is shown in Figure 3. In both cases, the number of optima observed by the GBSSEA is significantly greater than that of a standard SSEA. In the case of $M1$, the GBSSEA is able to support all five optima in the fitness landscape. However, in the case of $M2$, although able to support more optima than a simple SSEA, the GBSSEA could only identify and maintain four out of five optima. In every run of the GBSSEA, the smallest optimum was lost from the population. The Chi-Square-Like performance for each problem is shown in Figure 4. The GBSSEA consistently attained a Chi-Square-Like performance measure of zero on $M1$, suggesting that each run was able to produce a perfectly proportioned population, with each peak of the fitness landscape being assigned the correct number of individuals. The performance of

¹It should be noted that this population size is much larger than that typically used by other EAs on these two problems. However, this population size was chosen to highlight the speciation properties of the GBSSEA, rather than to analyse the computational efficiency of the approach.

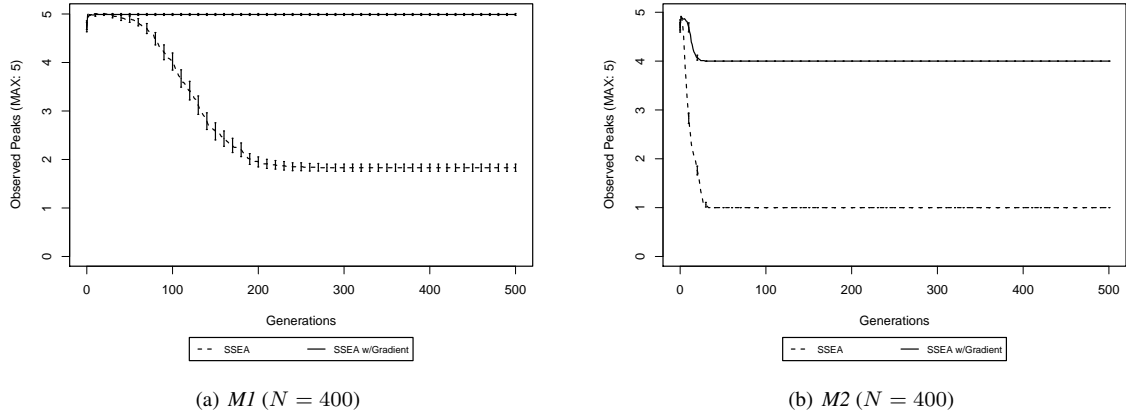


Figure 3: Discovery and preservation of niches under differing conditions of relative optima value.

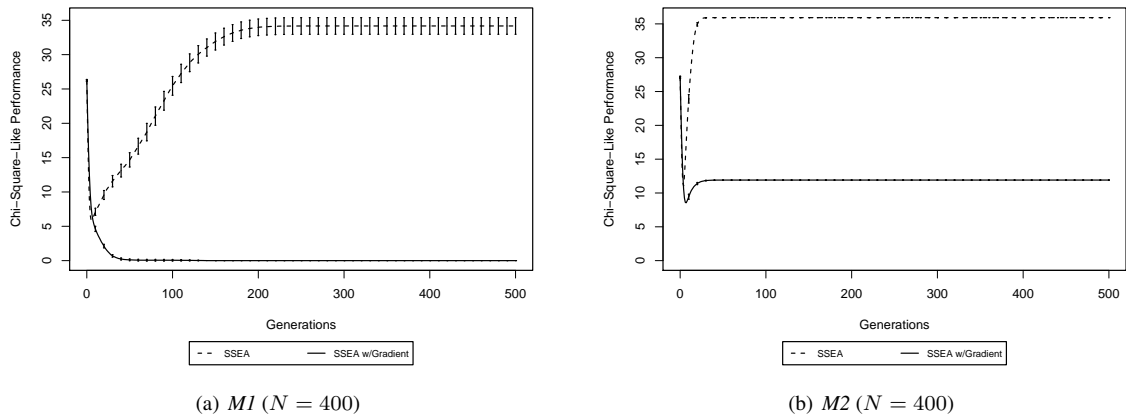


Figure 4: Discovery and preservation of niches under differing conditions of relative optima value.

the GBSSEA on $M2$ is significantly different, however; although there is an initial improvement in Chi-Square-Like performance, the overall performance measure plateaus after approximately fifty generations to a value of 11.91 ± 0.01 . This suggests that the GBSSEA exhibits significant deviations from an ideal population when searching the $M2$ fitness landscape.

A final visual analysis can be made by examining the actual distribution of individuals among peaks over the course of a run. Figure 5 shows the assignment of individuals to specific peaks over time for the $M1$ problem using a simple SSEA. A similar series of visualisations for the GBSSEA is shown in Figure 6. In each plot, a shaded cell indicates the presence of an individual that has discovered a peak in the fitness landscape. A white cell indicates an individual that has not discovered any peak. Both SSEA approaches have formed significant blocks of individuals representing the same peak. However, in the case of the simple SSEA, there is no consistency in these runs; the population will take on another configuration in subsequent runs. Likewise, the SSEA is slowly losing clusters of peaks through genetic drift. The GBSSEA exhibits none of these problems; the emergence of species in the population is consistent and repeatable, while the probability of losing species over time has reduced significantly. Additionally, we can see from these snapshots that the distribution of individuals among peaks is even, which confirms the perfect Chi-Square-Like performance observed in Figure 4.

The population plots for the $M2$ problem show a similar trend. The simple SSEA (as shown in Figure 7) quickly develops clusters of similar individuals, however these clusters are irregular in consistency and rapidly disappear with time. Although the GBSSEA cannot sustain all five optima in the population (as shown in Figure 8), it can maintain four out of five peaks in a robust and repeatable manner. However, these plots also give an insight into the relatively high Chi-Square-Like performance of the GBSSEA on $M2$. Aside from the fact that only four optima are present, we can see that the fittest peak (represented as the darkest shade) is under-represented. Instead, the second and fourth-largest peaks are the most represented in the population.




Figure 5: Emergence of niches over time for the *MI* problem using a simple SSEA.

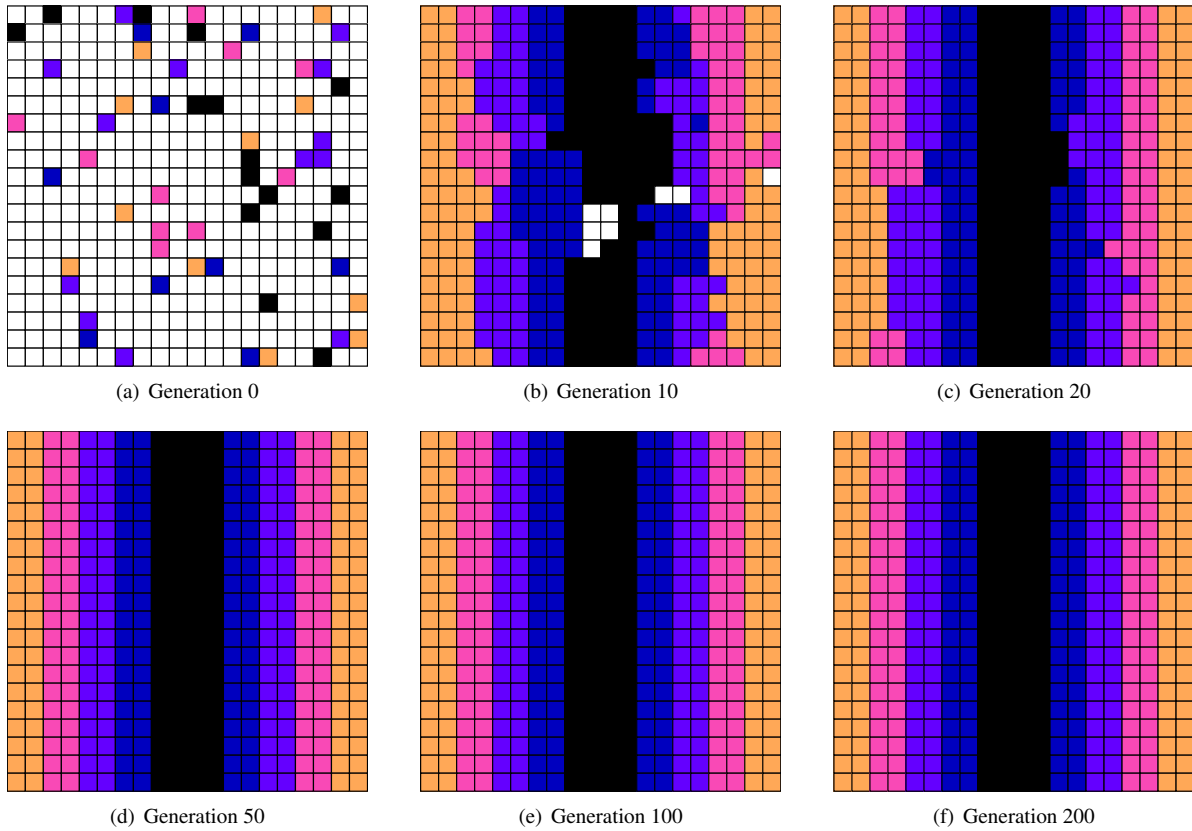


Figure 6: Emergence of niches over time for the *MI* problem using a Gradient-Based SSEA.

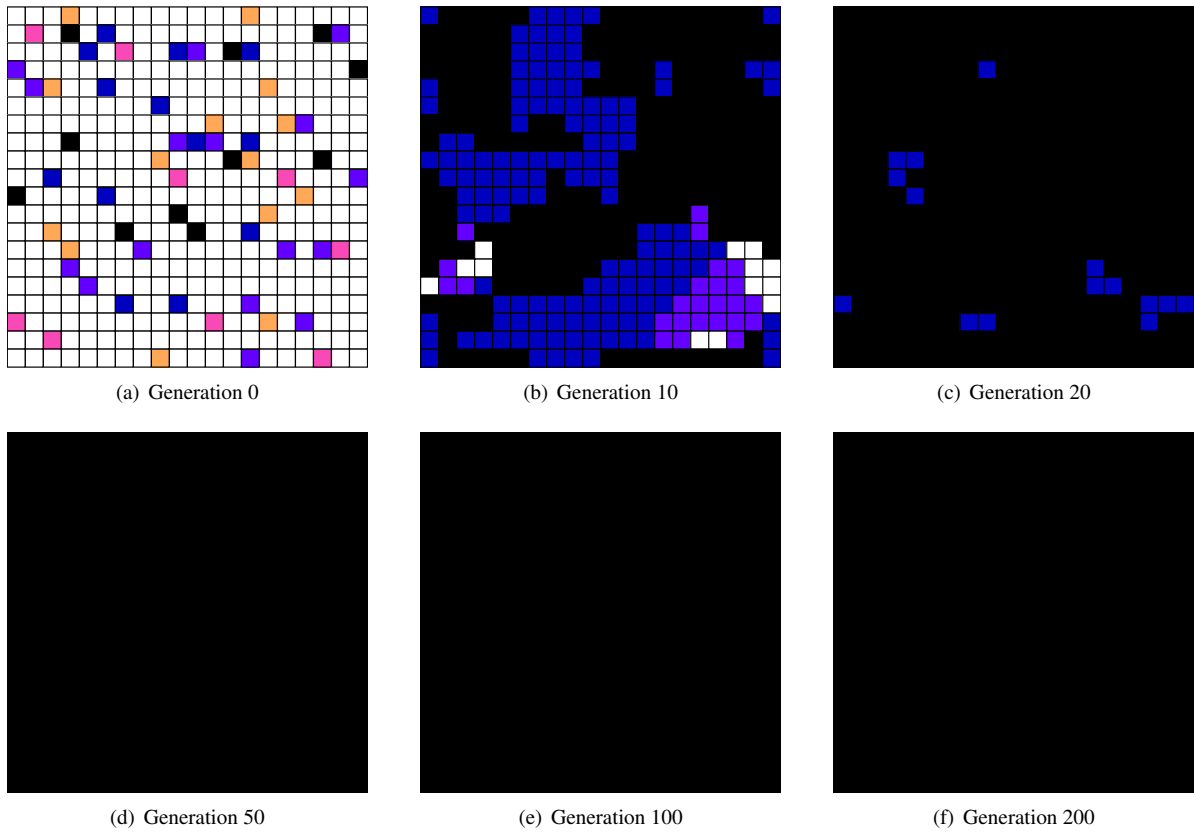


Figure 7: Emergence of niches over time for the $M2$ problem using a simple SSEA.




Figure 8: Emergence of niches over time for the $M2$ problem using a Gradient-Based SSEA.

4 The Niche-Allocation of Individuals in GBSSEAs

The previous section identified some interesting properties with regards to the behaviour of the GBSSEA method. In particular, the snapshots of the population over time revealed that the niching properties of the GBSSEA produced distinct bands of “species” within the population. The boundaries between two niches are of particular interest; this suggests that a single fittest peak cannot “take over” the entire population as there exists at least one location in space where it no longer represents the fittest solution. This raises the question - is it possible to determine a location in space in which two competing optima possess equivalent local fitness?

We can use Equation (2) to help answer this question. Ultimately, we are trying to find a location l such that the relationship:

$$\begin{aligned} localFitness(x, l) &= localFitness(y, l), \quad \text{or} \\ f_x \cdot e^{-\frac{1}{2} \left(\frac{\Delta(p_l, p_x)}{\sigma} \right)^2} &= f_y \cdot e^{-\frac{1}{2} \left(\frac{\Delta(p_l, p_y)}{\sigma} \right)^2} \end{aligned} \quad (5)$$

holds for any two optima x and y . Rearranging this equation to solve for p_l gives:

$$p_l = \frac{2\sigma^2 \log_e \left(\frac{f_x}{f_y} \right) + p_y^2 - p_x^2}{2(p_y - p_x)}. \quad (6)$$

This equation can be used to formulate a table of locations where any two peaks in the fitness landscape become equivalent, which is presented in Table 3. From this table, we can see that that the fitness of peaks four and five are equivalent at locations where the ideal phenotype = 0.990. This explains why the GBSSEA was unable to maintain all five optima within the population. The method used in this paper to generate the gradient is unable to produce a location in space with this value. In fact, in order to provide enough granularity in the gradient so that a location with this ideal phenotype is created, a minimum population size of 10000 individuals (a 100×100 torus) is required.

		Peak x					
		p_x	0.100	0.300	0.500	0.700	0.900
Peak y	p_y	fitness	1.000	0.917	0.707	0.459	0.250
	0.100	1.000	-	0.227	0.354	0.481	0.608
	0.300	0.917	0.227	-	0.481	0.608	0.735
	0.500	0.707	0.354	0.481	-	0.735	0.862
	0.700	0.459	0.481	0.608	0.735	-	0.990
	0.900	0.250	0.608	0.735	0.862	0.990	-

Table 3: The equivalence points between peaks for the $M2$ function.


The runs for $M2$ were repeated, this time with a larger population size of 10000 individuals. Once again, the number of observed peaks, the Chi-Square-Like performance and snapshots of individual runs were recorded. As shown in Figure 9, the larger population size allows the GBSSEA to maintain all five optima. The Chi-Square-Like performance still indicates a significant deviation from an “ideal” population, however. The overall behaviour of the GBSSEA is confirmed by the population snapshots (Figure 10); the fifth optimum resides in a single small row on each side of the population. The deviation from perfect Chi-Square-Like performance can be explained by the fact that, as with the smaller population size, the fittest peak is under-represented, while the second and third-fittest peak represent the majority of the population’s individuals.

4.1 Expected Niche Survival


Equation (6) can be used to determine the equivalence point of any two optima in the fitness landscape. In the $M2$ problem, the fifth optimum was effectively being crowded out of the population; as it lies close to the edge of possible phenotypes there was no locations presented by the gradient that it could effectively reside as the fittest representative. The fitness of the fifth optimum in $M2$ is one quarter that of the global optimum. In a real-world problem, it is unlikely that we would want to maintain such an optimum, however it is important to note this behaviour as a potential limitation of the GBSSEA approach.

A related problem to that presented by the GBSSEA is the situation in which a lesser-valued optimum lies between two optima of greater value. For example consider the simple three-peak function that is shown in Figure 11 and defined as:

$$PW(x, \alpha_1, \alpha_2, \alpha_3) = \alpha_i [1 - |6x - 2i + 1|]$$




(a) Number of peaks discovered




(b) Chi-Square-Like performance

Figure 9: Discovery and preservation of niches for the $M2$ problem using 100×100 population.




(a) Generation 0




(b) Generation 10




(c) Generation 20



(d) Generation 50



(e) Generation 100



(f) Generation 200

Figure 10: Emergence of niches over time for the $M2$ problem using a Gradient-Based SSEA and a 100×100 population size.




Figure 11: A “piecewise” tunable three-peak function. The parameters α_1 , α_2 and α_3 control the relative heights of the peaks.

where $0 \leq x \leq 1$ and $i = \lceil 3x \rceil$. This function has three optima and the height of each peak is tunable through a parameter α_i . In Figure 11, the values of α_1 and α_3 are set to 1.00 and 0.75 respectively. Using these values, it is important to know for which values of α_2 we can expect the GBSSEA to maintain all three optima within a single population.

As configured in Figure 11, the equivalence point for peaks one and three is found at a location with an ideal phenotype of 0.527 and the local fitness of each peak at this location is 0.354. Therefore, in order to ensure the survival of the second peak, the local fitness of the middle peak at this location must be greater than 0.354. Given that the fitness of the second peak is equal to α_2 , a suitable value for this parameter is when the condition

$$\alpha_2 > \frac{0.354}{resp(0.5, 0.527)}$$

holds. This means, to ensure survival of the second optimum, the value of α_2 must be at least 0.357 or greater. To test this, the GBSSEA was run on two configurations of the three-peak function; once with a value of $\alpha_2 = 0.300$, and with $\alpha_2 = 0.400$. The population plots over time for randomly selected runs are shown in Figures 12 and 13. As can be seen, the GBSSEA is unable to maintain the second peak when α_2 is set to 0.300. However, it is able to maintain a reasonable number of individuals on the second peak when $\alpha_2 = 0.400$.

5 Conclusion

The Gradient-Based Spatially-Structured Evolutionary Algorithm attempts to model parapatric speciation within an EA to encourage the simultaneous location of multiple optima in the course of a single run. This paper has demonstrated that the GBSSEA does not follow an expected pattern for allocation of individuals to the niches of a given fitness landscape. Rather, it determines *equivalence points* between peaks and uses these as boundaries to distribute individuals among the peaks of the problem space. Under certain conditions, lesser-valued optima will be dropped in favour of more *locally desirable* optima. In order to ensure that all optima are preserved, a suitably fine-grained gradient is required; with the implementation of the GBSSEA used in this paper, this is only possible by using excessively large population sizes. For real-world problems, this means that the GBSSEA cannot be relied upon to discover every optimum in a given fitness landscape. However, we can safely assume that the solutions that it does find are the most interesting and robust solutions in the fitness landscape.

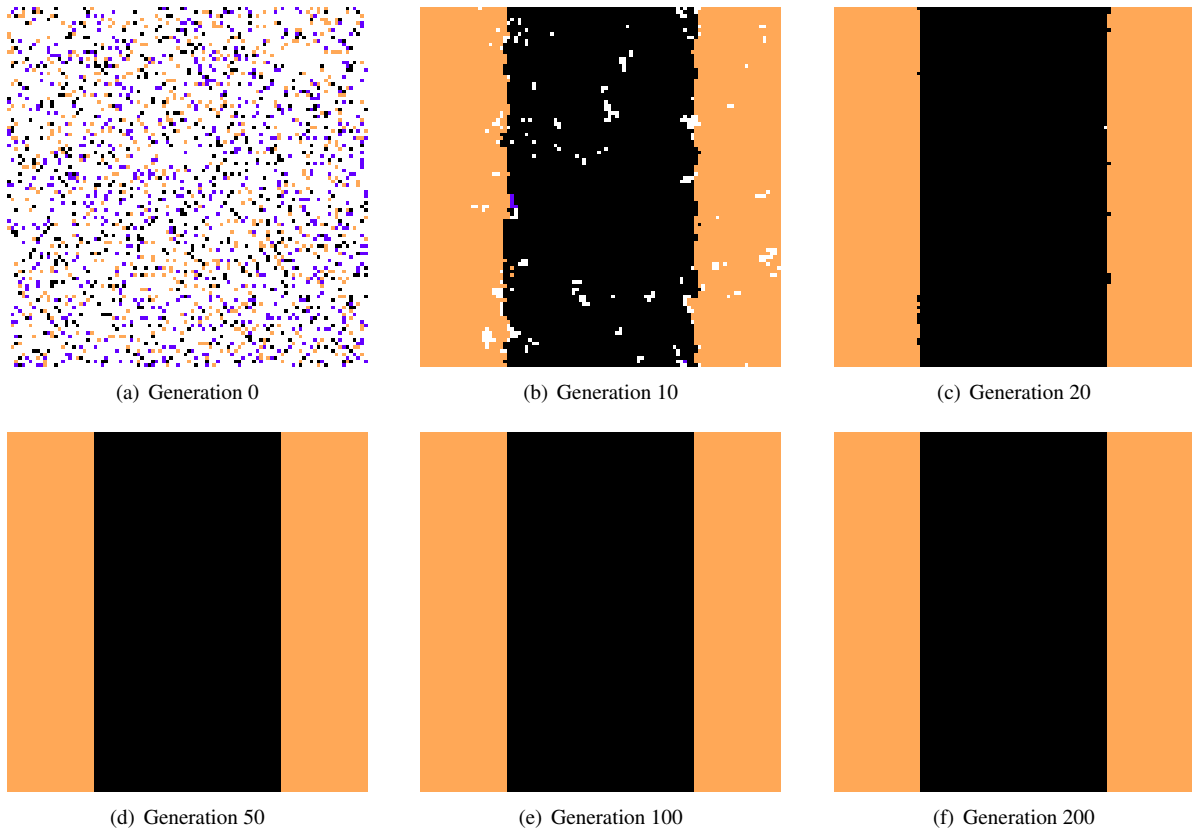


Figure 12: Emergence of niches over time with the three-peak problem ($\alpha_2 = 0.300$).




Figure 13: Emergence of niches over time with the three-peak problem ($\alpha_2 = 0.400$).

References

- Collins, R. J. & Jefferson, D. R. (1991). "Selection in Massively Parallel Genetic Algorithms" In R. K. Belew & L. B. Booker (eds), *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers San Mateo, CA.
- Davidor, Y. (1991). "A Naturally Occuring Niche & Species Phenomenon: The Model and First Results" In R. K. Belew & L. B. Booker (eds), *Proceedings of the Fourth International Conference on Genetic Algorithms (ICGA '91)*. Morgan Kaufmann Publishers San Mateo, California pp. 257–263.
- Deb, K. & Goldberg, D. E. (1989). "An Investigation of Niche and species Formation in Genetic Function Optimization" In J. D. Schaffer (ed.), *Proc. of the Third Int. Conf. on Genetic Algorithms*. Morgan Kaufmann San Mateo, CA pp. 42–50.
- Dick, G. & Whigham, P. A. (2005). "The behaviour of genetic drift in a spatially-structured evolutionary algorithm" In D. Corne, Z. Michalewicz, B. McKay, G. Eiben, D. Fogel, C. Fonseca, G. Greenwood, G. Raidl, K. C. Tan & A. Zalzala (eds), *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*. Vol. 2 IEEE Press Edinburgh, Scotland, UK pp. 1855–1860.
- Dick, G. & Whigham, P. A. (2006). "Multimodal Optimisation with Structured Populations and Local Environments" In T.-D. Wang, X. Li, S.-H. Chen, X. Wang, H. A. Abbass, H. Iba, G. Chen & X. Yao (eds), *SEAL*. Vol. 4247 of *Lecture Notes in Computer Science* Springer. pp. 505–512.
- Doebeli, M. & Dieckmann, U. (2003). "Speciation along environmental gradients." *Nature*. **421**(6920): 259–264.
- Goldberg, D. E. & Richardson, J. (1987). "Genetic algorithms with sharing for multi-modal function optimisation" *Proc of the 2nd Int. Conf. on Genetic Algorithms and Their Applications..* pp. 41–49.
- Kirley, M. (2001). "MEA: A metapopulation evolutionary algorithm for multi-objective optimisation problems" *Proceedings of the 2001 IEEE Conference on Evolutionary Computation*. IEEE Press Seoul, Korea.
- Mahfoud, S. W. (1992). "Crowding and preselection revisited" In R. Männer & B. Manderick (eds), *Parallel problem solving from nature 2*. North-Holland Amsterdam pp. 27–36.
- Mahfoud, S. W. (1995). "A Comparison of Parallel and Sequential Niching Methods" In L. Eshelman (ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms*. Morgan Kaufmann San Francisco, CA pp. 136–143.
- Mayr, E. (1970). *Populations, species and evolution; an abridgment of Animal species and evolution*. Harvard University Press.
- Murata, T., Ishibuchi, H. & Gen, M. (2000). "Cellular Genetic Local Search for Multi-Objective Optimization" In D. Whitley, D. Goldberg, E. Cantu-Paz, L. Spector, I. Parmee & H.-G. Beyer (eds), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*. Morgan Kaufmann Las Vegas, Nevada, USA pp. 307–314.
- Murata, T., Ishibuchi, H. & Gen, M. (2001). "Specification of Genetic Search Directions in Cellular Multi-objective Genetic Algorithms" In E. Zitzler, K. Deb, L. Thiele, C. A. C. Coello & D. Corne (eds), *First International Conference on Evolutionary Multi-Criterion Optimization*. Springer-Verlag. Lecture Notes in Computer Science No. 1993. pp. 82–95.