An Investigation of Sexual Selection as a Mechanism for Obtaining Multiple Distinct Solutions

Michael Ratford, Andrew Tuson*, and Henry Thompson

Department of Artificial Intelligence, University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, U.K.
{miker,andrewt,ht}@dai.ed.ac.uk

Abstract. The ability to obtain multiple distinct solutions in a single run is an important, though often forgotten, practical advantage of genetic algorithms, and therefore methods that enhance this ability are desirable. The approach taken here is inspired by nature: sexually reproducing organisms do not mate indiscriminately — the choice of mate has a large impact upon the fitness of the organism's offspring, and by balancing exploration and exploitation, mate (or sexual) selection can lead to speciation effects which may enhance the genetic algorithm's ability to locate multiple distinct solutions. The investigation described here shows that sexual selection was able to enhance the genetic algorithm's ability to locate and maintain multiple distinct solutions, although it can interact detrimentally with other techniques designed for the same purpose.

1 Introduction

Standard implementations of the genetic algorithm (GA) [11] select the second parent either completely at random, or on the basis of fitness. However, in nature, sexually reproducing organisms do not mate indiscriminately. The reason for this is simply that half of the offspring's genes will come from the other parent, and the choice of mate will therefore have a large impact upon their children's fitness, and hence the survival of the parent's own genes. Therefore organisms effect some form of sexual selection, one manifestation of which is to control the balance between out-breeding and inbreeding (ie. how different the parents are to each other).

Therefore, could using different criteria to select the second parent make crossover (and therefore the genetic algorithm) more effective? One possible application is the following: there are problems where obtaining more than one solution is desirable, even if it may involve a slight reduction in the quality of the solution(s) thus obtained. Examples where obtaining multiple distinct solutions would be useful in a practical situation include the following: if the problem being optimised may change (for example, when machines break down in a scheduling

* Author to whom correspondence should be addressed.

problem), if there are preferences or constraints on the problem that cannot readily be presented to the optimiser [24], or in multi-objective optimisation where the aim is to obtain a scattering of points on the trade-off surface [7].

In this context, the genetic algorithm would appear to be a good choice due to its population-based nature, an hypothesis that has been shown to be valid in some problems [21] — however any enhancement of the genetic algorithm's ability in this context would be of use. Previous work on increasing the number of solutions obtained has centred around approaches such as geographical speciation [3,15], crowding [4] and sharing [9] (a good review can be found in [23]).

It is also possible that use of mate (or sexual) selection to restrict which solutions can mate with each other may promote the development of distinct mating groups, thus increasing the number of distinct solutions in the final population. The question is: does this approach help to produce more multiple solutions in practice? To address this question, an investigation of the potential of using sexual selection to obtain multiple distinct solutions is described in this paper.

1.1 Literature Review on Sexual Selection

Work has been performed on using mate choice in genetic algorithms, with encouraging results. For example [18] showed how the use of a 'seduction function' based upon an aesthetic measure, such as unitation in the case of the royal road problem [14], to select the second parent can improve genetic algorithm performance. Work by [13] discussed the various ways that sexual selection can be used in evolutionary computation, and indicated that speciation behaviours may occur when sexual selection is used. In a different vein [6] showed that 'incest prevention', the prohibition of crossover between identical or very similar strings, can prevent premature convergence and give improved results. Also [2] has extended the idea of incest prevention somewhat by introducing the concept of 'ancestry' — in other words, solutions are prohibited from mating with, say, their grandparents, siblings, etc. Finally, more recent results [17] have shown that sexual selection appears to be a robust technique for improving genetic algorithm performance over a wide range of test problems in the genetic algorithm literature.

2 Implementation

The publicly available genetic algorithm implementation PGA [20] was modified and used as a testbed for this investigation. When crossover was used, the first parent was selected in the usual fashion; with modifications being made in the selection of the second parent. A seduction function [18], in this case based upon the similarity, in terms of an appropriate metric, of each of the prospective second parents to the first parent was then used to select the second parent. The process is summarised by the pseudo-code below (adapted from [18]):

```
until (termination condition met) do {
     parent1=select(population);
     parent2=seduce(parent1,population);
     operate(parent1,parent2,child1,child2);
     merge(child1,child2,population);
     }
```

PGA was selected a testbed for this study as it implements a wide range of options for experimentation. For instance, the following options were used as part of this study:

- **Operators** — the standard genetic algorithm operators (two-point crossover and bit-flip mutation) for binary-coded problems were used.
- **Test Problems** — a wide range of binary-coded problems were built into PGA, and additional problems were added. All of the test problems used are listed in section 3.
- **Population Models** — PGA implements spatial and unstructured population models, which can both be used in a generational or steady-state manner.

An overview of the extensions made to PGA for the purposes of this investigation will now be given.

2.1 The Seduction Function

For this study, Hamming distance was used as the distance metric, and this was normalised to a value between 0 and 1 by dividing the raw value by the maximum hamming distance possible.

This similarity measure needs to be somehow processed to provide a measure of *seduction*; so that the higher a chromosome's seduction is in relation to another, then the better it's potential as a suitor will be. An appropriate function would give a low measure for chromosomes that were very similar or dissimilar, and a high measure for chromosomes that were fairly similar. It should also return a value between 0 and 1 to facilitate the production of a measure combining fitness and seduction.

For the sake of flexibility, additional parameters were introduced: *centre* and *depth* (the names approximating their effect on the shape of the function).

As a sophisticated curve would be computationally expensive, and a previous study [17] found it was not necessary, a simple function (Figure 1) was used. This was defined as:

$$seduction = \begin{cases} 1 - depth + \frac{distance \times depth}{centre} & distance \leq centre \\ 1 - depth \times \frac{distance - centre}{1 - centre} & distance > centre \end{cases}$$

where $distance$ was the Hamming distance and $seduction$ the seduction value.

This method seeks a compromise between the two approaches of incest prevention and assortative mating — the mating of similar solutions. This is because

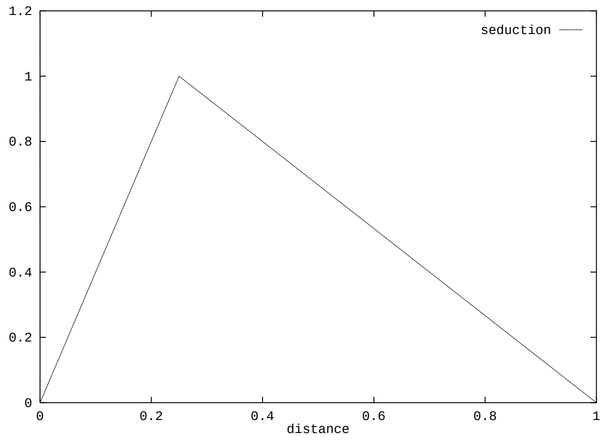

Fig. 1. The Simple Seduction Function

it would seem that which approach is best depends upon the nature of the problem and one's goal. Avoiding mating similar individuals would seem to be useful for avoiding premature convergence, but would be harmful if there is a large amount of multi-modality, in which case mating similar individuals would be a good idea.

A Dynamic Seduction Function Simulated annealing [12] uses the notion of reducing *temperature* to balance exploration and exploitation, where search at initially high temperatures is likely to be random, but as the temperature is gradually decreased random leaps from states become less and less likely.

A similar principle can be used to control the bias of the seduction function. At first chromosomes are given dissimilar mates with the aim of finding as many promising areas within the search space as possible, and as evolution progresses the seduction bias is altered so that more similar mates are chosen in order to exploit these promising areas. Another effect may be to spread an initially skewed population of chromosomes.

The tricky issue here is when to change "temperature". It would be nice to have some mathematically derived temperature schedule, but in the absence of this an ad hoc method was used of changing the curve's centre by a small step at each generation until a specified, final centre was reached.

2.2 Selecting The Mate

Now that a final measure for sexual selection has been produced, how will it be used? Calculating the "overture" from every other chromosome in the population each time an individual was selected for breeding would clearly be unmanageably time-consuming. Therefore tournament [1,8], marriage, and courtship selection were considered. Two of these are recent additions to the armoury of selection methods, and are now outlined below. The term 'Monte Carlo size' will be used to refer to the number of chromosomes competing in these selection techniques.

Marriage Selection *Marriage* selection is a variant of tournament selection designed to soften the standard version's inherently high level of selection pressure, as the chromosome within the competing subpopulation with the highest level of fitness or attraction will always win [19]. Chromosomes are picked at random until either one is found that is more attractive than the first, or a fixed number — the Monte Carlo size — of individuals have been selected. The most attractive is then selected as the second mate.

Courtship Ritual The method outlined above depends only on the ranking of attraction. Greater control and flexibility might be provided by an operator that considered the actual, rather than relative, attraction value. Seduction functions could then be fine-tuned to provide a more targeted bias.

This consideration led to the formulation of the *courtship ritual* [16], which works by picking a member of the population at random and a random number between 0 and 1, the maximum of the attraction function. If the chosen chromosome's seduction measure is above the random number then accept it for mating, if not choose another and repeat the process. For this method the Monte Carlo size is the maximum number of chromosomes picked before the most attractive so far encountered is selected as the second parent.

2.3 Obtaining Multiple Solutions

Two other approaches to obtaining multiple solutions were taken from the literature, to see how they would interact with sexual selection.

Crowding Crowding clusters similar chromosomes only implicitly, by randomly choosing a small number of chromosomes and replacing the most similar, in terms of Hamming distance, to the child chromosome being introduced [4].

Spatial Reproduction PGA has an existing reproduction technique where the population is located on a grid, so that was used. There are, in fact, many minor variations on the theme of spatial structuring, but the general principles remain the same [15].

The population is randomly placed on a rectangular grid, again with opposite edges connected. The chromosome at any location is replaced by the child, if fitter, of two parents chosen by going on a random "walk" of a specified number of steps from that location. The parents chosen are the fittest chromosomes found on each walk. The second child replaces the chromosome "below" that location — again, only if fitter than that already there.

This process was modified so that the first parent was chosen as before, but the second parent was the most attractive encountered on the second walk with respect to the first parent.

3 The Test Problems

This investigation tested the performance of sexual selection over a range of multimodal problems in the literature, so to provide a through test of this technique.

- De Jong's test function 5 (DJ5) [4].
- Modified binary F6 (BF6) [22].
- Deb's functions (DEB1, 2, 3, and 4) [5].
- Himmelblau's function (HIMM) [10].

Attention was confined to binary coded problems as this is the most common representation used with genetic algorithms. In addition, many problems in operational research can be formulated as 0/1 IP (integer programming) problems, for which a binary representation is naturally suited.

4 Sexual Selection Alone

Space permits no more than a summary of the results obtained, given in this and the following sections. For a detailed exposition of the results, the reader is directed to [16]. It should be noted here that the results presented here are typical of several runs and as such the differences reported here are significant and not the result of a 'fortunate' run.

4.1 Unfettered Speciation

When run on DEB1 with a population of 250 chromosomes a genetic algorithm with sexual selection did no better than a standard one at maintaining chromosomes at each peak, although both found near-optimal solutions around each maximum after 1000 generations. Even with sexual selection, however, chromosomes were not maintained around each peak for much longer. This was felt to be hardly surprising given that a steady-state model was being used, as the combination of the population's unidirectional fitness movement and genetic drift made dominance by a single peak inevitable.

The discovery of multiple solutions seems to have been assisted by the large population size. What happens if the population size is reduced to 100? In this case mate selection does cause an improvement. After 200 generations both methods (with and without sexual selection) were maintaining solutions at four out of five peaks — there were no chromosomes near the maximum at 0.3 in the original population, which means that this had not been explored.

However after 300 generations the standard genetic algorithm had reduced its area of interest to around one peak, whereas the genetic algorithm using

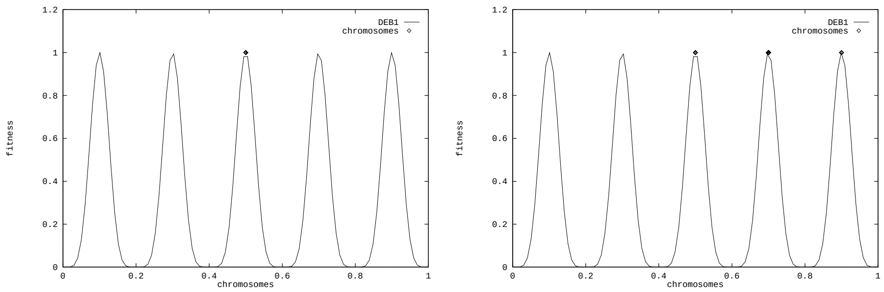

Fig. 2. Result Obtained with a Population of 100 Chromosomes, DEB1, After 300 Generations using a Standard GA and a GA with Sexual Selection

sexual selection was maintaining chromosomes around three peaks (Figure 2), and this situation remained constant until after 700 generations when only two peaks had chromosomes around them. This was good, but not perfect. Could performance be improved beyond this?

4.2 Sharpening Seduction

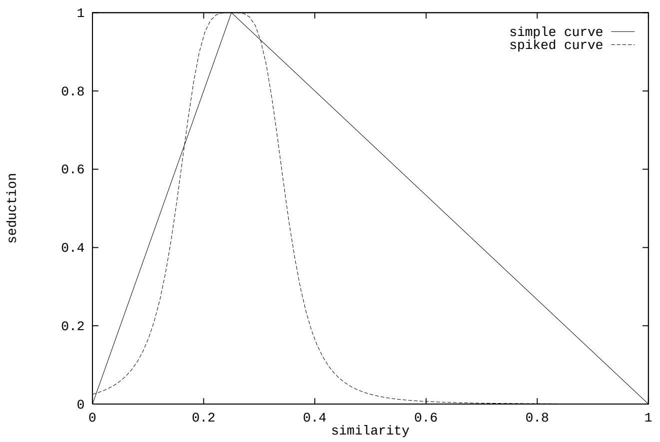

Fig. 3. A Spiked Seduction Function

Careful consideration of the particular version of mate selection implemented, however, may indicate some improvements. Using the simple seduction curve, centred on 0.25, there is a 0.5 or greater probability of accepting as a mate half of all possible chromosomes for any one chromosome — those which are a Hamming distance away of between 0.125 and 0.625 of the chromosome length. Perhaps the bias needs a little more "focus". Changing the seduction function so that those outside of a narrow band are excluded may encourage speciation. Thus a sharper seduction function, as in Figure 3, was tried on DEB1. This made no difference when marriage selection was used, but when applying courtship selection the third peak was maintained for another 100 generations. It would appear that the shape of the seduction function has some effect in this case, but the choice is not critical.

4.3 Shaking the Population

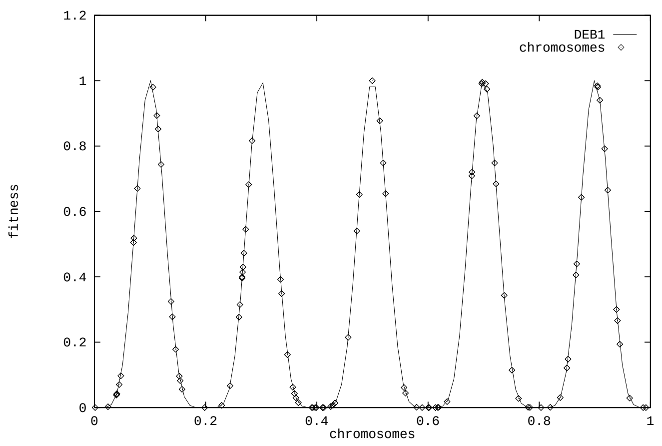

Fig. 4. An Initial Population of 100 Chromosomes, DEB1

Earlier, the use of a dynamic seduction function to attempt to redistribute, or "shake", an initially skewed distribution of chromosomes was mentioned. When the distribution of the initial population was examined (Figure 4), it was observed that the fittest chromosomes were gathered around the peaks away from 0.3, thus giving the genetic algorithm a reduced chance of finding the peak at 0.3. If, when evolution is started, diverse chromosomes are mated the whereabouts of the peaks in the search space could be established and efforts concentrated around there.

This would support the use of a dynamic seduction function. However, when the experiments were performed on DEB1, the change made little difference.

4.4 A Closer Look

Closer examination of the initial population (Table 1) shows the three fittest members close to the peaks at 0.5, 0.7 and 0.9, and these eventually came to dominate the population, however much the parameters were adjusted. In an attempt to avoid this, the search was made more vigorous and "undirected", by increasing the mutation rate and using a non-elitist generational selection. However, when this was tried the solution quality deteriorated unacceptably.

Rank	Decoded Solution	Fitness
1	0.499765	0.9999591
2	0.697472	0.9952811
3	0.696505	0.9909944
4	0.703496	0.9909896
5	0.904606	0.9844043
6	0.905171	0.9803775
7	0.105179	0.9803169
8	0.705977	0.9738675
9	0.909122	0.9400626
10	0.112306	0.8933300

Table 1. Fittest of Initial Population, DEB1

Even if some method was found that works for this particular problem (DEB1), experience suggests that it will be fragile, not working for slightly different initial populations or problems. As suspected, genetic drift was a problem. Thus schemes were introduced to overcome this.

5 Crowding

Crowding was introduced: the three random chromosomes were chosen and the closest to the child chromosome replaced. Unfortunately when tried on DEB1, solution quality suffered with no advantage in terms of multiple solutions. The technique was then changed from a "blind" to an elitist version, so that the closest of three was found and then replaced only if it was less fit than the potential replacement. This went some way to solving the problem of poor solution quality, but did nothing to aid the maintenance of diverse solutions — after 400 generations all solutions were clustered around a single peak.

6 Spatial Selection

Does the use for a spatial population model enhance sexual selection's ability to maintain multiple solutions? To investigate this the steady-state version of spatial and mate selection was used, in order to facilitate comparison with the previous methods used — all other parameters were those used for "standard" mate selection.

6.1 Does it Work?

The results for DEB1 were encouraging. Chromosomes were kept at all five peaks until after 700 generations when there were only four, 800 generations when there were only three, and then the peaks were gradually lost. Sexual selection on its own, however, could only manage three peaks until 700 generations, so this result is a definite improvement (albeit on only one problem). Using the "spiked" seduction function gave no advantage — chromosomes were kept at five peaks for 600 generations.

Examination of the solutions at 600 generations uncovered the optimal solutions at 0.3 and 0.9. At the other peaks, solutions had fitness of greater than 0.99, within 1% of the optimum.

However, the generational version seemed to produce slightly better results. Did it do as well at producing multiple solutions? The answer appeared to be yes. It held all peaks for only five generations, but solutions are again found within 1% of the optima for each peak. Four peaks, however, were maintained for thirty generations.

To check to see if the algorithm was working as expected, a map of the chromosomes was obtained for each generation. A cluster of chromosomes "belonging" to a peak might be expected. This is partly true, but not to any great extent. This may have something to do with the fact that random walks of 5 steps were taken; rather large given the size of the population.

What happens when this figure is reduced to 3? There may be a better "grouping", but that may only be because the number of peaks reduced far more rapidly. What if the size of the population is increased to 225, keeping the walk length of 5? There was still no geographical coalescence between chromosomes gathering around the same peak. Perhaps a form of this technique emphasising local competition, such as Davidor's grid system [3], would demonstrate the expected phenomenon. Consideration should also be given to the fact that the figures show chromosomes belonging to a peak in phenotype space, whereas the genetic algorithm operates in genotype space.

7 Expanding The Study

But is sexual selection effective at finding multiple solutions for other test functions? The study was expanded to test this.

The results for these functions (DEB2, DEB3, DEB4 and HIMM) show that sexual selection helps in finding multiple solutions.

For DEB3, chromosomes were kept around all five peaks at 10 generations, and then three peaks at 15. A steady-state non-spatial genetic algorithm maintained all peaks for 300 generations, after which two peaks were lost. It is not

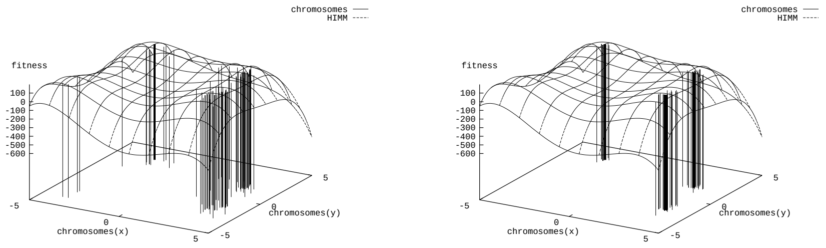

Fig. 5. Chromosomes After 5 and 10 Generations, HIMM, Generational GA with Mate Selection

easy to exactly compare the two, but spatial selection will make $15 \times 20 = 3000$ evaluations, while the steady-state genetic algorithm makes $300 \times 2 = 600$ evaluations. Of course, not all of the evaluated chromosomes will be placed in the population in the former case. For a more direct comparison, five peaks were maintained at 10 generations by a generational genetic algorithm with mate selection, but solution quality was worse.

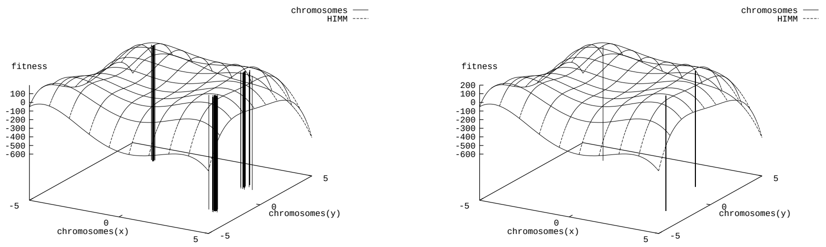

Fig. 6. Chromosomes After 15 and 120 Generations, HIMM, Generational GA with Mate Selection

Similar effects were seen for HIMM (Figures 5 and 6). Chromosomes remained around all four peaks for five generations, and around three peaks until after 120 generations. The steady-state non-spatial genetic algorithm maintained all peaks after 200 generations, after which peaks were lost until after 500 generations the entire population were clustered around the same peak.

Fitness function	Evaluations per Optimum Sexual. Selection	Spatial + Sexual Selection
DJ5	7513	120433
BF6	20363	61567
DEB1	1809	75292
DEB2	5773	114461
DEB3	3940	115783
DEB4	9883	155575
HIMM	63900	691309

Table 2. Average Evaluations to an Optimum

7.1 Spatial Selection Revisited

The number of evaluations the combined spatial and sexual selection scheme took to find each optimum was far greater than that for sexual selection only (Table 2). Even allowing for the possibility that the former is finding more solutions per run, it is difficult to wholeheartedly promote its use. Repeated iterations of a standard genetic algorithm with mate selection would seem to be a more efficient alternative.

8 Conclusion

This study has examined the ability of sexual selection to obtain multiple distinct solutions. Sexual selection was found to be able to maintain solutions at multiple optima longer than a standard genetic algorithm, especially as the population size was reduced. Experiments were then performed to see if there was any interaction between sexual selection and other techniques that are used to help obtain multiple solutions.

The combination of sexual selection with geographically structured populations was found to work better than either of the techniques alone, though the additional time required makes it questionable whether the additional effort is rewarded. Using sexual selection with crowding lead to poor quality solutions with no advantage in terms of multiple solutions. This is especially encouraging when it is noted that earlier work with this form of sexual selection [17] showed that the quality of solutions obtained was also improved for a wide range of test problems.

It appears that sexual selection is a potentially useful addition to the armoury of techniques that can enhance the genetic algorithm's ability to locate and maintain multiple optima, though care should be taken when it is combined with other such techniques. Finally, the results here show that there is still some useful mileage in applying biological metaphors to optimisation.

9 Acknowledgements

Thanks to the EPSRC for their support of Michael Ratford and Andrew Tuson via studentships with references 95411068 and 95306458.

References

1. A. Brindle. *Genetic Algorithms for Function Optimization.* PhD thesis, University of Alberta, 1981.
2. R. Craighurst and W. Martin. Enhancing GA Performance through Crossover-Prohibitions Based on Ancestry. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, 1995.
3. Y. Davidor. A naturally occurring niche and species phenomenon: the model and first results. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms.* Morgan Kaufmann, 1991.
4. K. A. De Jong. *An Analysis of the Behavior of a Class of Genetic Adaptive Systems.* PhD thesis, University of Michigan, 1975.
5. K. Deb. Genetic algorithms in multimodal function optimization. Master's thesis, University of Alabama, Department of Engineering Mechanics, 1989. TCGA Report No. 89002.
6. L. J. Eshelman and J. D. Schaffer. Preventing premature convergence in genetic algorithms by preventing incest. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms.* Morgan Kaufmann, 1991.
7. C. M. Fonseca and P. J. Fleming. An Overview of Evolutionary Algorithms in Multiobjective Optimization. *Evolutionary Computation*, 3(1):1–16, 1995.
8. D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. In G. J. E. Rawlins, editor, *Foundations of Genetic Algorithms*, pages 69–93. Morgan Kaufmann, 1991.
9. David E. Goldberg and Jon Richardson. Genetic algorithms with sharing for multimodal function optimisation. In J. J. Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms*, Hillsdale, NJ, 1987. Lawrence Erlbaum Associates.
10. D. M. Himmelblau. *Applied Nonlinear Programming.* McGraw-Hill, 1972.
11. J. H. Holland. *Adaptation in Natural and Artificial Systems.* Ann Arbor: The University of Michigan Press, 1975.
12. S. Kirkpatrick, C.D. Gelatt, Jr., and M.P. Vecchi. Optimization by Simulated Annealing. *Science*, 220:671–680, 1983.
13. G. F. Miller. Exploiting mate choice in evolutionary computation: Sexual selection as a process of search, optimization, and diversification. In T. C. Fogarty, editor, *Proceedings of the 1994 AISB Workshop on Evolutionary Computing.* Springer-Verlag, 1994.
14. M. Mitchell, S. Forrest, and J. H. Holland. The Royal Road for Genetic Algorithms: Fitness Landscapes and GA Performance. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, Cambridge, MA, 1991. MIT Press.
15. H. Mühlenbein. Evolution in time and space — the parallel genetic algorithm. In G. Rawlins, editor, *Foundations of Genetic Algorithms.* Morgan Kaufmann, 1991.

16. M. Ratford. The Single Chromosome's Guide To Dating. Master's thesis, Department of Artificial Intelligence, University of Edinburgh, 1996.
17. M. Ratford, A. Tuson, and H. Thompson. The Single Chromosome's Guide To Dating. In *Proceedings of the International Conference on Artificial Neural Networks and Genetic Algorithms*, 1997.
18. E. Ronald. When Selection Meets Seduction. In Larry J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, San Francisco, CA, 1995. Morgan Kaufmann.
19. P. Ross. Personal communication, 1996.
20. P. Ross and G. H. Ballinger. *PGA - Parallel Genetic Algorithm Testbed*. Department of Artificial Intelligence, University of Edinburgh, 1993.
21. P. Ross and D. Corne. Comparing Genetic Algorithms, Simulated Annealing, and Stochastic Hillclimbing on Timetabling Problems. In T. C. Fogarty, editor, *Evolutionary Computing; AISB Workshop, Sheffield 1995, Selected Papers. Springer-Verlag Lecture Notes in Computer Science 993*, 1995.
22. J. D. Schaffer, R. A. Caruna, L. J. Eshelman, and R. Das. A study of control parameters affecting online performance of genetic algorithms for function optimization. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann, 1989.
23. P.A. Turner. Genetic algorithms and multiple distinct solutions. Master's thesis, Department of Artificial Intelligence, University of Edinburgh, 1994.
24. A. Tuson, P. Ross, and T. Duncan. On Interactive Neighbourhood Search Schedulers. In *the 6th Workshop of the UK Planning and Scheduling Special Interest Group*, 1997.