# Why use Elitism and Sharing in a Multi-Objective Genetic Algorithm?

**Robin C. Purshouse and Peter J. Fleming**
Department of Automatic Control and Systems Engineering
University of Sheffield, UK.
{r.purshouse, p.fleming}@sheffield.ac.uk

## Abstract

Elitism and sharing are two mechanisms that are believed to improve the performance of a multi-objective evolutionary algorithm (MOEA). Using a new empirical inquiry framework, this paper studies the effect of elitism and sharing design choices using a benchmark suite of two-criterion problems. Performance is assessed, via known metrics, in terms of both closeness to the true Pareto-optimal front and diversity across the front. Randomisation methods are employed to determine significant differences in performance. Informative visualisation of results is achieved using the attainment surface concept. Elitism is found to offer a consistent improvement in terms of both closeness and diversity, thus confirming results from other studies. Sharing can be beneficial, but can also prove surprisingly ineffective. Evidence presented herein suggests that parameter-less schemes are more robust than their parameter-based equivalents (including those with automatic tuning). A multi-objective genetic algorithm (MOGA) combining both elitism and parameter-less sharing is shown to offer high performance across the test suite.

## 1 INTRODUCTION

Evolutionary multi-criterion optimisation (EMO) practitioners are faced with a number of design choices beyond those encountered in a standard evolutionary algorithm (EA). Suitable strategies for elitism and sharing can significantly improve optimiser performance. This paper presents new evidence and understanding concerning elitism and sharing that will help practitioners to make informed choices. Through the application of tractable algorithm modifications and a rigorous experimental framework, the effect of MOEA component-level choices can be more clearly exposed.

An *EMO empirical inquiry framework* is introduced in Section 2. The dual performance metrics of closeness and diversity are measured using the generational distance and spread metrics respectively. Statistical comparisons are then made using randomisation testing. Information-rich visualisations of the identified trade-off surfaces are obtained using attainment surfaces. The analysis is based on the two-criterion set of test problems proposed by Zitzler *et al* [2000].

The performance of a baseline MOGA optimiser is established in Section 3. The effects of elitism and sharing are then considered with reference to this baseline. An elitist strategy, based on Zitzler's [1999] *universal elitism*, is developed in Section 4. Sharing methodologies for the promotion of diversity are discussed in Section 5. A new parameter-less technique, formulated as an accompaniment to Pareto-based ranking, is compared to the standard parameter-based approach. In Section 6, a high-performance MOGA incorporating both elitism and parameter-less sharing is investigated.

## 2 EMO INQUIRY FRAMEWORK

### 2.1 TEST SUITE

The established set of test problems developed by Zitzler *et al* [2000] (ZDT) is used in this study. The suite consists of six, tractable, two-criterion functions, with varying characteristics as summarised in Table 1.

Table 1: Test function characteristics

| NAME | ATTRIBUTES |
|---|---|
| ZDT-1 | Convex front |
| ZDT-2 | Non-convex front |
| ZDT-3 | Non-contiguous convex front |
| ZDT-4 | Many local fronts, single global convex front |
| ZDT-5 | Deceptive problem, convex front |
| ZDT-6 | Non-uniform distribution, non-convex front |

### 2.2 MEASURING PERFORMANCE

Performance of a MOEA can be decomposed into two criteria:

- **Closeness** – the nearness of the obtained non-dominated solutions to the true front.
- **Diversity** – the coverage of the trade-off surface by the obtained solutions.

The ideal outcome, in test cases of this type, is a final population with a uniform distribution of globally non-dominated solutions spread across the entire trade-off surface. Various performance metrics have been proposed to measure closeness, diversity, and in some cases both together. Some metrics require that the global trade-off surface is known and can be sampled (straightforward in the ZDT cases), whilst others involve a purely relative comparison of two results sets. This study utilises three known performance metrics: *generational distance* to measure accuracy, *spread* to measure diversity, and *attainment surfaces* to provide visualisation.

- **Generational distance** – an average of the Euclidean distances between each obtained solution and the nearest point on the true front [Veldhuizen, 1999].
- **Spread** – the sum of the differences between nearest neighbour distances and the mean of all such distances, coupled with a term to account for the extent of the obtained front [Deb *et al*, 2000].
- **Attainment surface** – the boundary in criterion-space that separates the region that is dominated by the obtained solutions from that which is non-dominated [Fonseca and Fleming, 1996].

The superposition of multiple attainment surfaces can be treated statistically and also provides a rich qualitative indication of performance. A typical plot is shown later in Figure 1. The heavy line indicates the 50%-attainment surface (akin to the median), the thinner lines show the 25% and 75% surfaces (quartiles), and the dotted lines describe the 0% and 100% surfaces. Thus location, dispersion, and skewness information can be obtained in a similar manner to the box plot [Cleveland, 1993].

### 2.3 ANALYSING PERFORMANCE

Upon completion of a single run of a specific MOEA configuration on a particular problem, three sets of non-dominated criterion vectors (and associated solutions) are obtained, namely:

- **final population** – the non-dominated vectors in the final population of the algorithm,
- **on-line archive** – the final elite set of vectors, and
- **off-line archive** – the complete set of non-dominated vectors identified by the algorithm.

The first of these sets is used for analysis and comparison purposes in this study since it provides the most appropriate measure of the on-line trade-off surface *maintenance* capabilities of an algorithm.

An evolutionary algorithm is a stochastic process and, thus, multiple runs (samples) are required in order to infer reliable conclusions as to its performance. Hence, 35 runs have been conducted for each MOEA configuration when applied to a particular test problem. The performance of the algorithm is expressed in the resulting distributions of generational distance and spread. A statistical comparison of two configurations is then possible through the use of a test statistic.

In this study, the mean difference between two generational distance (or, alternatively, spread) distributions is taken as the test statistic. The significance of this observed result is then assessed using *randomisation testing*. This is a simple, yet effective, technique that does not rely on any assumptions concerning the attributes of the underlying processes, unlike conventional statistical methods [Manly, 1991]. The central premise of the method is that, if the observed result has arisen by chance, then this value will not appear unusual in a distribution of results obtained through many random relabellings of the samples. The randomisation method proceeds as follows:

1. Compute the difference between the means of the samples for each algorithm: this is the observed difference.
2. Randomly reallocate half of all samples to one algorithm and half to the other. Compute the difference between the means as before.
3. Repeat Step 2 until 5000 randomised differences have been generated, and construct a distribution of these values.
4. If the observed value is within the central 99% of the distribution, then accept the null hypothesis. Otherwise consider the alternative hypotheses. This is a two-tailed test at the 1%-level.

The null hypothesis is that the observed value has arisen through chance and so there is no performance difference between the two configurations. The alternative hypotheses are that the difference is unlikely to have arisen through chance and that one configuration has outperformed the other (depending on which side of the distribution the observed difference falls, and the direction in which the difference has been calculated).

Note that the observed value is included as one of the random relabellings since, if the null hypothesis is true, then this value is one of the possible randomisation results. 5000 randomisations is regarded as an acceptable quantity for a test at the 1%-level [Manly, 1991].

The results of randomisation testing are simple to visualise, as shown by the example in Figure 3. The randomised results are described by the grey histogram, whilst the observed result is depicted as a filled black circle. Each row shows the performance on a particular test function (from ZDT-1 at the top, to ZDT-6 at the bottom). The left-hand column indicates the relative performance regarding closeness, and the right-hand column shows the corresponding difference in diversity.

## 3 BASELINE MOGA

### 3.1 DESCRIPTION

The baseline optimiser used in this study has been developed according to the holistic design principles championed by Michalewicz and Fogel [2000] and has previously been shown to be effective at solving the ZDT test problems [Purshouse and Fleming, 2001]. A summary of the algorithm is provided in Table 2.

The multi-criterion performance of a solution is scalarised using Fonseca and Fleming's [1993] Pareto-based ranking procedure. A solution is ranked according to the number of solutions in the population that are *preferred* to it. If the entire Pareto-optimal front is to be identified, the preference relation collapses to a test for Pareto dominance.

Table 2: Baseline configuration

| EMO COMPONENT | STRATEGY |
|---|---|
| *GENERAL*<br>Population size<br>Total generations | <br>100<br>250 |
| *ELITISM* | None |
| *EVALUATION* | [1] Fonseca and Fleming [1993] Pareto-based ranking.<br>[2] Linear fitness assignment with rank-wise averaging.<br>[3] No modification of fitness to account for population density. |
| *SELECTION* | Stochastic universal sampling |
| *REPRESENTATION*<br>Real parameter functions<br><br>Binary function | <br>Concatenation of real number decision variables. Accuracy bounded by machine precision.<br>Binary string, 80 bits in length. |
| *OPERATORS*<br>For real representations<br><br>For binary representations | <br>[1] Naïve crossover<br>Probability = 0.8.<br>[2] Gaussian mutation (initial search power of 40% of variable range; sigmoidal scaling set to 15; feasibility requirement of one standard deviation).<br>Probability = Expected value of 1 phenotype per chromosome.<br>[1] Single-point binary crossover.<br>Probability = 0.8.<br>[2] Simple bit-flipping mutation.<br>Probability = 1/80. |

When ranking is complete, initial fitness values can be prescribed. The population is sorted according to rank and fitnesses are assigned by interpolating between the highest fitness value for the best rank and the lowest fitness value for the worst rank. In the baseline algorithm, linear interpolation is used and fitness is varied between the population size (highest) and unity (lowest). The ratio of these two fitnesses is a definition of the *selective pressure* of the assignment mechanism. Solutions of the same rank then have their fitnesses adjusted to the average of the original assignments for that rank.

Part of this study is concerned with the effect of diversity-preserving mechanisms. Therefore no manipulation of the above fitnesses through sharing is undertaken.

*Stochastic universal sampling* has been chosen as the selection mechanism [Baker, 1987]. This method achieves maximum spread with minimal bias, but is non-parallelisable. In total, 100 selections are required since the chosen reinsertion strategy is that all offspring replace all parents (no generational gap) and since for the chosen genetic operators two parents are required to produce two offspring.

Since five of the test problems feature real number decision variables, it is logical to use a real number representation for these problems. Hence, a candidate solution is described by a concatenation of phenotypic decision variables. The other test problem, ZDT-5, explicitly uses binary variables, thus a binary representation is natural for this problem.

Different representations require different search operators. For the binary chromosome case, the familiar single-point two-parent crossover and bit-flipping mutation operators are employed. Good results are known to be achievable using this simple approach [Zitzler *et al*, 2000]. For real representations, the so-called *naïve crossover* is used in conjunction with a *Gaussian mutation* operator. The former of these search tools is a very simple two-parent single-point crossover operator, where the crossover sites are limited to points between decision variables. This offers quite a low-power search, since it cannot generate any values for decision variables that were not present in the original population. However, when coupled with a complementary high-power search tool, the resulting search capabilities are considerable[1]. Gaussian mutation is one such operator. Its main benefit is that it provides tuneable search power in the form of the standard deviation. This can be exploited to provide on-line adaptation that avoids the generation of infeasible solutions and controls convergence speed by varying the search from near global early on to very local towards the end. *Sigmoidal* variation, as a function of the percentage of generations completed, of the standard deviation is useful because it allows concentrated periods of high- and low-power search [Purshouse and Fleming, 2001].

## 3.2 PERFORMANCE

Attainment surfaces illustrating the performance of the baseline algorithm are shown in Figure 1. Particularly good results were achieved for ZDT-1, ZDT-2, and ZDT-3 (Figures 1a, 1b, and 1c respectively) in terms of both closeness to the global Pareto front and diversity across the front. The tight envelopes of attainment indicate the high level of consistency achieved in these cases. The MOGA struggled to achieve good coverage of the surface as $f_1$ approaches zero on ZDT-2. Note that this is a region where there is little trade-off between the objectives.

As shown in Figure 1d, the wider envelopes of attainment produced for the multi-fronted ZDT-4 signify entrapment in a locally non-dominated front. On no occasions did the MOGA converge to the global trade-off surface although coverage across the identified fronts was good.

The baseline MOGA achieved reasonable closeness to the global front on ZDT-5. Performance on this deceptive test function is depicted in Figure 1e. Note that on no occasions was the algorithm able to identify the extreme right-hand section of the discrete trade-off surface.

Rather poor performance was observed on the non-uniform ZDT-6, as shown in Figure 1f. Coverage was especially poor on the less dense area of the front. This, together with the missing section of the ZDT-5 front, is

[1] Coincidentally, the incorporation of naïve crossover largely prevents the convergence failures encountered by Ikeda *et al* [2001], thus showing that MOEA failure cannot be solely blamed on the use of Pareto ranking in these cases.

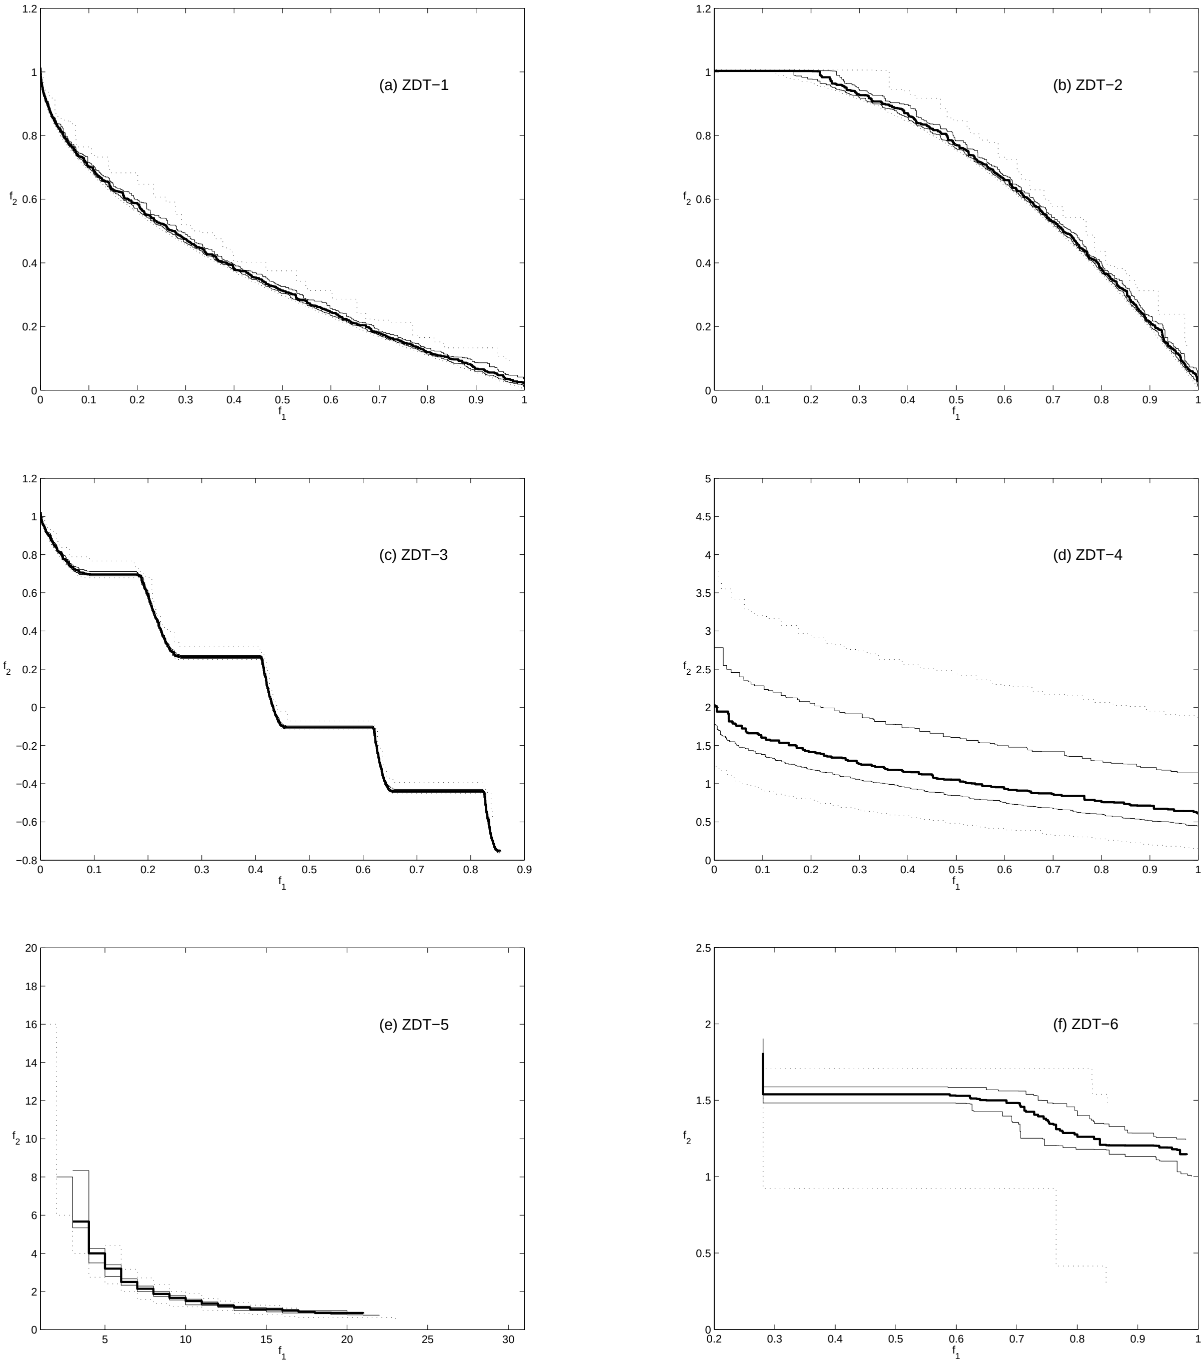


Figure 1: Attainment surfaces – baseline MOGA solving the ZDT problems

the strongest indication that density-based sharing would be beneficial. Closeness to the true Pareto front is also not good. Only the 0%-attainment surface lies on the global front, where coverage is particularly poor. Furthermore, the position of this front with respect to the median and quartiles suggests that this result is something of an outlier.

# 4 ELITIST STRATEGY

*Elitism* is the process of preserving previous high-performance solutions from one generation to the next. This is conventionally achieved by simply copying the solutions directly into the new generation. Elitism has long been considered an effective method for improving the efficiency of an EA [De Jong, 1975]. Various recent studies in the EMO community have indicated that the inclusion of an elitist element can considerably improve the performance of an MOEA [Zitzler *et al*, 2000; Deb *et al*, 2000]. The two main issues are (1) how to manage the size of the elite sub-population, and (2) how to use elitism to drive the search effectively.

The elitist strategy adopted in this study is a variant on the approach developed by Zitzler [1999] and is illustrated by the schematic in Figure 2. The key difference is that the archive size is allowed to vary within pre-defined limits, whilst the number of newly generated candidate solutions is varied such that the total population size (elites plus new solutions) is held constant.

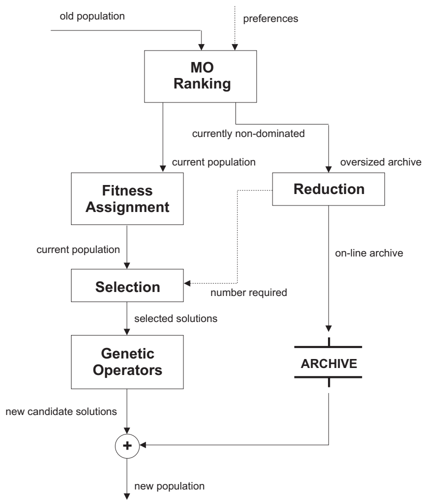


Figure 2: Elitist strategy

The on-line archive is initialised to the empty set, whilst the initial population is initialised to a random set of candidate solutions (possibly seeded with information provided by the decision-maker). The populations at subsequent iterations of the algorithm are the combination of new solutions and current elite solutions. The currently non-dominated solutions in the population are identified and are stored as the new, potentially over-sized, archive. Over-represented solutions are then eliminated from the archive, if necessary, using the *SPEA-2* truncation procedure [Zitzler *et al*, 2001]. This is an effective reduction technique for two-criterion problems.

When the new elite set has been finalised, the size of this set is known, and thus the number of new candidate solutions required to fill the population can be calculated. These solutions are created through the selection and genetic manipulation of members of the current population. The new solutions are then combined with the elite set to form the total population, which completely replaces the old population.

This elitist strategy has been integrated within the baseline MOGA and has been applied to the test problems. Randomisation test results between the elitist model and the baseline are shown in Figure 3. Observed differences to the left of the randomisation distribution offer evidence in favour of the elitist version outperforming the baseline case.

There is considerable evidence, clearly shown by the results in Figure 3, that the elitist algorithm produces results closer to the true front than the baseline for ZDT-1, 2, 3, 4, and 6. Superior performance in terms of diversity is strongly suggested for ZDT-1, 2, 4, 5, and 6.

Elitism increases the convergence speed of the algorithm. The danger of sub-optimal convergence is somewhat reconciled by the distributed nature of the elite set. High-power search operators, such as the Gaussian mutation operator used in this work, can also reduce the risk of premature convergence. Hence, the increased convergence exhibited in this study is expected.

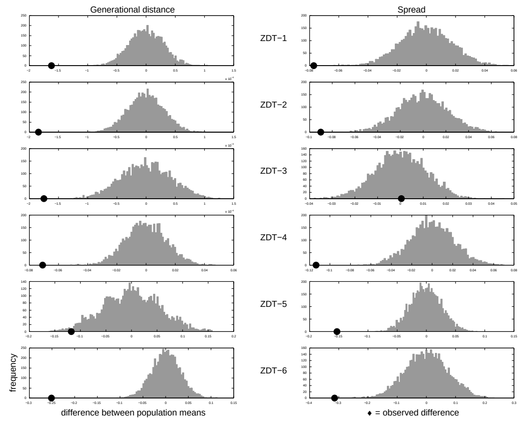


Figure 3: Elite versus baseline

The elitist scheme also maintains the characteristics of the currently identified trade-off surface within the on-line population. Thus, diversity of non-dominated solutions in the population is maintained and encouraged (through the thinning of similar criterion vectors) by the truncation mechanism. This helps to explain the improvement in diversity seen in the results. However, the truncation process only represents the current distribution: it does not, directly (though fitness), drive the search towards a superior distribution. Despite this fact, the inclusion of elitism did lead to improved diversity on the non-uniformly distributed ZDT-6. Modifications to the fitness, such as those arising through sharing, may assist further in improving diversity across the trade-off surface.

# 5 SHARING STRATEGY

## 5.1 INTRODUCTION

One of the aims of a multi-objective evolutionary algorithm is to obtain a suitable *distribution* of candidate solutions in regions of interest to the decision-maker. In an evolutionary algorithm, this can be achieved through the formation of sub-population clusters – known as *niches* – within the global population. *Fitness sharing* is the most popular method for fostering this niching process [Goldberg and Richardson, 1987]. In this approach, the raw fitness value of a candidate solution is reduced by a factor dependent on the local population density. This measure should be made in the domain over which a good distribution is of interest: usually criterion-space.

## 5.2 PARAMETER-BASED METHODS

Fitness sharing has been shown to combat the problem of *genetic drift* (population convergence to a single point due to stochastic selection errors), thus helping to attenuate the possibility of sub-optimal convergence and to enhance coverage of trade-off surfaces. However, the power law equation on which the technique is based requires a definition of *closeness* in order to calculate the population densities. This can be difficult to estimate in practice. Furthermore, the method is sensitive to choice of this *niche size* parameter. Several methods have been proposed in order to estimate the niche size, for example Deb and Goldberg [1989] and Fonseca and Fleming [1993], of which the dynamic approach of Fonseca and Fleming [1995] is particularly interesting.

Fonseca and Fleming [1995] noted the similarity between the power law sharing function and the *Epanechnikov* kernel density estimator used by statisticians. The kernel smoothing parameter used in the estimator was found to be directly analogous to the fitness sharing niche size parameter. The key benefit of this is that statisticians have developed successful techniques for estimating the value of this parameter [Silverman, 1986]. Furthermore, the approach is amenable to update at each generation of the EA population. This approach can be regarded as parameter-based sharing with automatic tuning.

Epanechnikov sharing has been added to the baseline MOGA and has been applied to the benchmark problems. Sharing is performed using the Euclidean distance metric in the criterion domain. Results of a randomisation comparison with the baseline algorithm are shown in Figure 4. Observed values that favour the sharing scheme will lie to the left of the randomisation distribution.

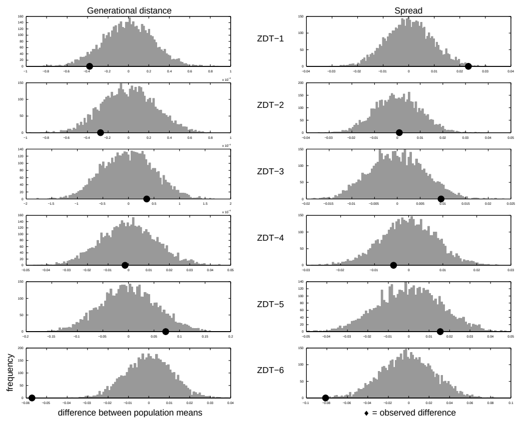


Figure 4: Epanechnikov versus baseline

The inclusion of Epanechnikov sharing has improved both aspects of performance on the non-uniform ZDT-6. Note in particular that a method designed to improve diversity has also helped to improve convergence, thus suggesting the strong interaction between the two performance criteria. However, no improvements in either diversity or closeness have been achieved for any other test function. Indeed there is some evidence to suggest deterioration in diversity on ZDT-1. The lack of improvement to diversity is of particular concern, since the elitist results in Section 4 have indicated that diversity *can* be greatly improved on these problems. A possible explanation for the lack of success is that the automatic parameter selection is providing poor estimates.

## 5.3 PARAMETER-LESS METHODS

The difficulty and inconvenience involved in determining the niche size value has led many researchers to investigate parameter-less methods for achieving niching. A new approach is presented here that increases the resolution of the Fonseca and Fleming [1993] Pareto-based ranking procedure through the inclusion of population density information. An *intra-ranking* is performed on candidate solutions of identical Pareto-based rank, discriminating on the basis of population density at that rank. Solutions in less dense areas receive a superior intra-ranking to their counterparts in denser regions. This approach requires a definition of *distance* (Euclidean nearest neighbour is used herein) but does not require a definition of closeness. In practice, the distance metric is likely to be problem dependent and could conceivably include decision-maker preference information. Following the new fine-grained ranking process, the fitness assignment procedure remains unchanged.

Using this scheme, if one candidate solution is preferred to (dominates) another, then the former is guaranteed to have a superior fitness value. Also, when all solutions are non-dominated, discrimination is based purely on density. If, in addition, the density is globally uniform then all fitnesses are identical.

With any type of ranking scheme, information content is lost. Ranking indicates that one solution lies in a more densely packed region than another solution but the actual difference in density between the two is lost. This limits the amount of information available to the search procedure but protects against premature convergence to locally *superfit* solutions and removes the requirement for a niche size setting.

The results for this new sharing scheme, compared to the non-sharing baseline model, are shown in Figure 5. The central aim of sharing is to improve the distribution of solutions in criterion-space and this should be primarily evident in the spread results. There is strong evidence to suggest that the new method improved spread on ZDT-3 and ZDT-4. The use of the Epanechnikov kernel, by contrast, did not improve results on these problems. In no case was the absence of a sharing mechanism shown to be preferable. However, there is little evidence to suggest that the use of sharing made any difference to the results for ZDT-6. This is particularly disappointing since this problem has a non-uniform distribution across its trade-off surface: a situation in which sharing is considered a highly appropriate strategy.

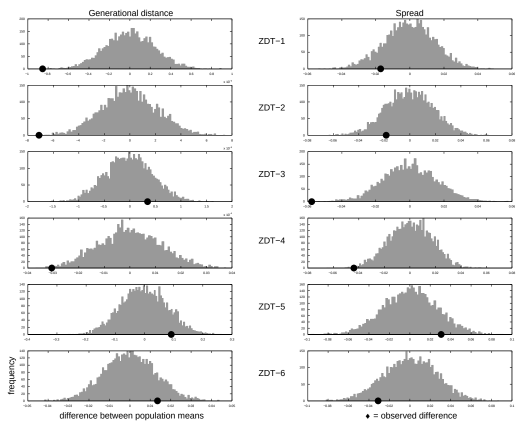


Figure 5: New sharing versus baseline

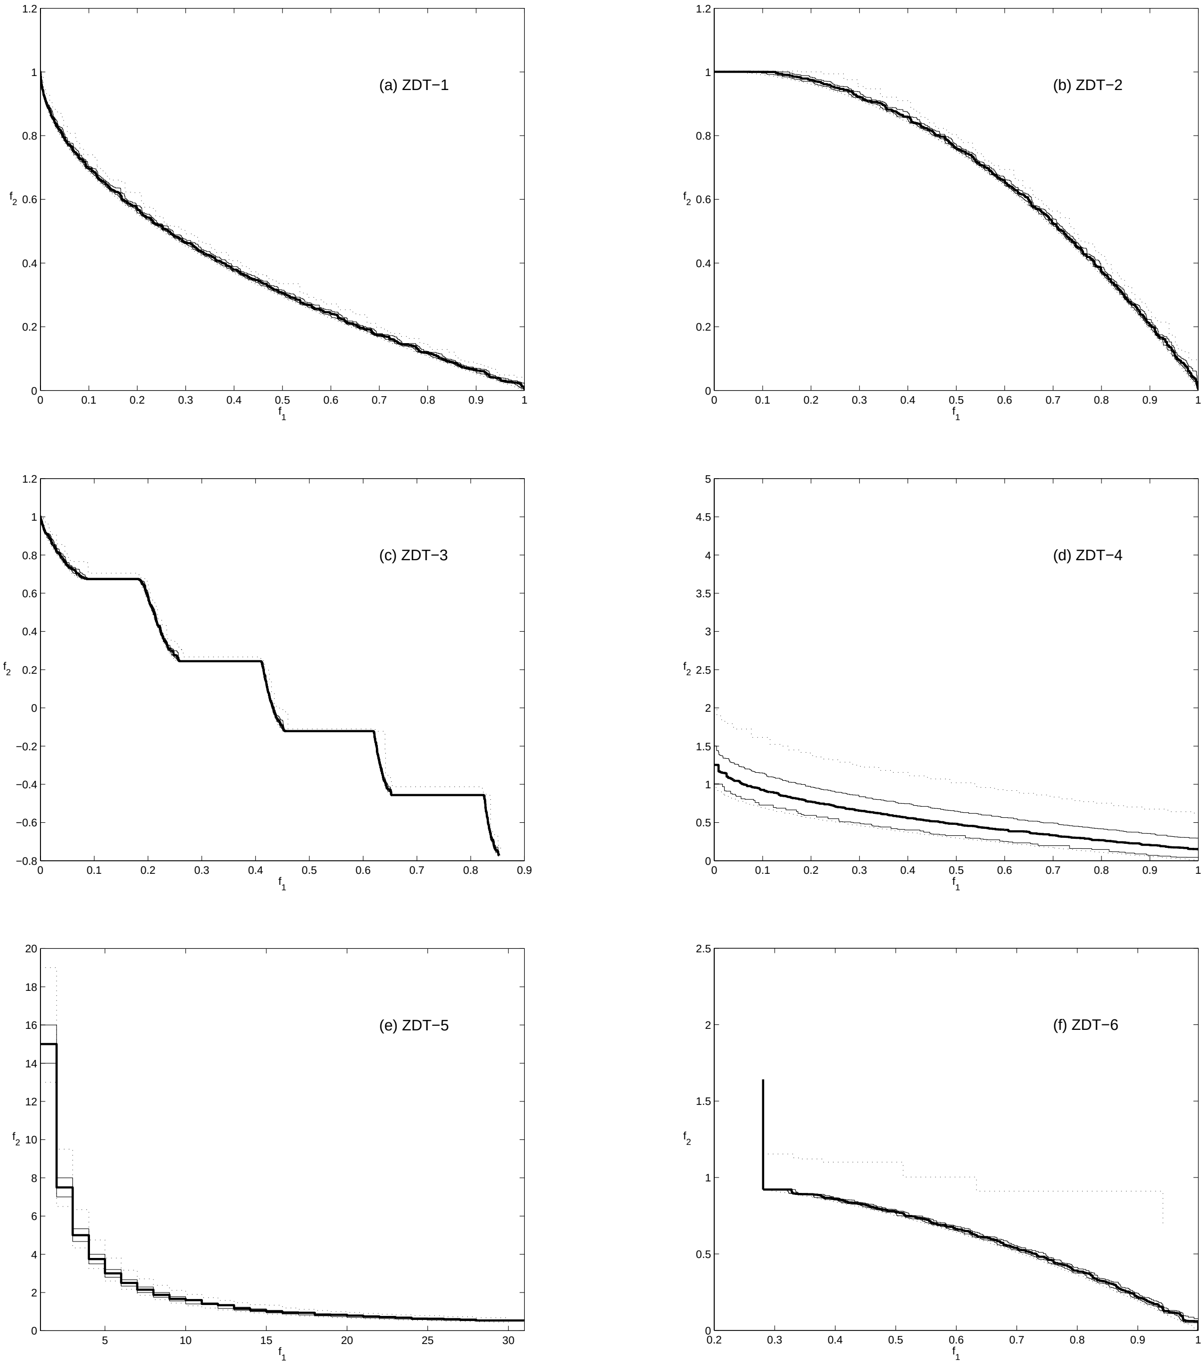


Figure 6: Attainment surfaces – elitist, parameter-less sharing MOGA solving the ZDT problems

# 6 HIGH-PERFORMANCE MOGA

The use of an elitist strategy or a parameter-less sharing strategy in isolation has been shown to offer improved performance. It is instructive to also consider the effect of these schemes in combination. Attainment surfaces for such an algorithm are shown in Figure 6. The envelopes of attainment are generally very tight, indicating good consistency. As evident from Figure 6d, closeness has been greatly improved on ZDT-4: indeed the 25%-attainment surface lies very close to the global front of this difficult test problem. Complete coverage of the right-hand portion of the trade-off surface has been achieved for ZDT-5, as shown in Figure 6e. Finally, closeness and diversity have been much improved on ZDT-6 (Figure 6f).

Comparisons with the baseline MOGA are made using randomisation testing in Figure 7. Observed differences that lie to the left of the randomisation distribution favour the new algorithm. Compelling evidence points to the algorithm substantially outperforming the baseline in terms of diversity across all six benchmark problems. The *combination* of elitism and new sharing was required in order to achieve this notable result: neither elitism nor sharing alone was shown to be sufficient. Improved closeness was observed for ZDT-1, 2, 4, and 6 (the result for ZDT-5 is not significant at the 1%-level).

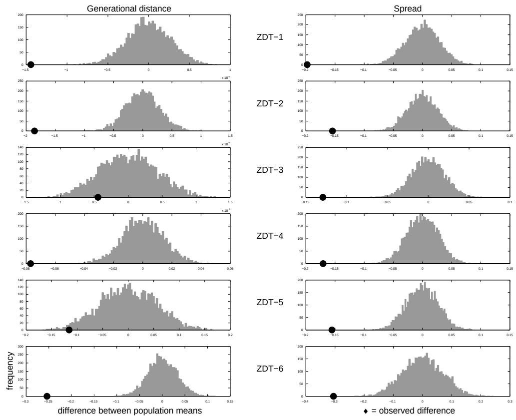


Figure 7: Elitist, sharing MOGA versus baseline

# 7 CONCLUSION

Using a progressive and tractable experimental approach, supported by appropriate statistical and visual analyses, this paper has shown that elitist and sharing strategies can significantly improve the performance of an evolutionary multi-criterion optimiser. Existing elitist heuristics are again shown to be beneficial, this time using a new analysis technique and in the context of MOGA. However, the shortcomings of a popular parameter-based sharing technique have been exposed, as have the dangers of relying too heavily on an automatic parameter-setting method. A new parameter-less method of sharing has been introduced and has been shown to be more reliable than the standard method. Impressive results were achieved when both elitism and sharing were used together. As a final word of caution, these results have been obtained for two-criterion problems: further research is required to ascertain the effectiveness of these methods as the dimension of the problem increases.

The results described in this paper, together with an extended research report, are available for download from the following site:

`http://www.shef.ac.uk/~acse/research/students/r.c.purshouse/`

**Acknowledgments**

The first author gratefully acknowledges support from UK EPSRC. The authors would also like to thank the anonymous reviewers for their helpful comments and suggestions, and would like to express special thanks to Dr Nick Fieller from the Department of Probability and Statistics at the University of Sheffield for his valuable advice.

**References**

Baker, J. E., 1987, *Reducing bias and inefficiency in the selection algorithm*, Proceedings of the 2nd International Conference on Genetic Algorithms, 14-21.

Cleveland, W. S., 1993, *Visualizing Data*, Hobart Press, Summit, NJ.

Deb, K., Agrawal, S., Pratab, A., and Meyarivan, T., 2000, *A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II*, KanGAL report 200001.

Deb, K. and Goldberg, D. E., 1989, *An Investigation of Niche and Species Formation in Genetic Function Optimization*, Proceedings of the 3rd International Conference on Genetic Algorithms, 42-50.

De Jong, K. A., 1975, *An analysis of the behavior of a class of genetic adaptive systems*, PhD thesis.

Fonseca, C. M. and Fleming, P. J., 1993, *Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization*, Proceedings of the 5th International Conference on Genetic Algorithms, 416-423.

Fonseca, C. M. and Fleming, P. J., 1995, *Multiobjective Genetic Algorithms Made Easy: Selection, Sharing and Mating Restriction*, Proceedings of GALESIA '95, 45-52.

Fonseca, C. M. and Fleming, P. J., 1996, *On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers*, Parallel Problem Solving from Nature – PPSN IV, Lecture Notes in Computer Science, **1141**:584-593.

Goldberg, D. E. and Richardson, J., 1987, *Genetic Algorithms with Sharing for Multimodal Function Optimization*, Proceedings of the 2nd International Conference on Genetic Algorithms, 41-49.

Ikeda, K., Kita, H., and Kobayashi, S., 2001, *Failure of Pareto-based MOEAs: Does Non-dominated Really Mean Near to Optimal?*, CEC 2001, **2**:957-962.

Manly, B. F. J., 1991, *Randomization and Monte Carlo Methods in Biology*, Chapman and Hall, London.

Michalewicz, Z. and Fogel, D. B., 2000, *How to Solve It: Modern Heuristics*, Springer-Verlag, Berlin.

Purshouse, R. C. and Fleming, P. J., 2001, *The MultiObjective Genetic Algorithm Applied To Benchmark problems – An Analysis*, ACSE Research Report 796, `http://www.shef.ac.uk/~acse/research/students/r.c.purshouse/reports/report796.pdf`.

Silverman, B. W., 1986, *Density Estimation for Statistics and Data Analysis*, Monographs on Applied Statistics and Probability **26**, Chapman and Hall.

Veldhuizen, D. A. V., 1999, *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*, PhD thesis.

Zitzler, E., 1999, *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*, PhD thesis.

Zitzler, E., Deb, K., and Thiele, L., 2000, *Comparison of Multiobjective Evolutionary Algorithms: Empirical Results*, Evolutionary Computation **8**(2):173-195.

Zitzler, E., Laumanns, M., and Thiele, L., 2001, *SPEA2: Improving the Strength Pareto Evolutionary Algorithm*, TIK-Report 103, ETH Zürich.