

Assessing the Convergence of Rank-Based Multiobjective Genetic Algorithms

Rajeev Kumar & Peter Rockett

Department of Electronic & Electrical Engineering, University of Sheffield, Mappin St, Sheffield S1 3JD, UK

Abstract Many problems in engineering and related areas require the simultaneous optimisation of multiple objectives and to this end, rank-based genetic algorithms have proved very successful. The key issue of convergence of vector optimisations, however, has not hitherto been explicitly addressed. In this paper we introduce rank histograms to both assess convergence of a given single genetic optimisation and to combine results from multiple runs to test for the adequacy of the individual optimisations. Results are presented on two analytic benchmark multiobjective problems where the optimal solution set is known *a priori*, and on a problem in partitioning a pattern recognition task.

1 - Introduction

Many problems in engineering and related areas require the *simultaneous* genetic optimisation of a number of, possibly competing, objectives. One method which has been employed in the past to tackle this problem has been to combine the multiple objectives into a single scalar by some linear combination. The combining coefficients, however, are usually based on heuristics or guesswork and can exert an unknown influence on the outcome of the optimisation. A more satisfactory approach is to use the notion of Pareto optimality first proposed by Goldberg [1] in which an optimal *set* of solutions prescribes some surface - the Pareto front - in the vector space of the objectives. For a solution on the Pareto front no objective can be improved without simultaneously degrading at least one other. Key to Pareto optimality is the concept of *domination* where one vector solution in m -elements is defined as dominating another:

$$(\mathbf{X}_1 \prec \mathbf{X}_2) \triangleq (\forall_m)(x_{m1} \leq x_{m2}) \wedge (\exists_m)(x_{m1} < x_{m2})$$

Although a number of important results have been reported for multiple objective GAs [2..4], one issue that has been largely ignored has been that of convergence. For a single objective GA gauging convergence is almost trivial. For multiple objectives, however, this is no simple matter. Horn *et al* [3] observed that they were in no position to appreciate whether their obtained solution set was indeed a true optimal solution. For objective spaces of dimensionalities equal to or less than three some visualisation of the solution space could possibly be employed but for higher dimensionality this is not feasible. The approaches adopted in the past have been either to run the GA for some fixed number of iterations, or to terminate the optimisation when some fraction of the population has become non-dominated. The first of these is unsatisfactory since either a large amount of CPU time could be wasted producing further generations for an optimisation which has already converged; alternatively there is no way of knowing that a particularly stubborn problem is still far from convergence. The second option is ill-conceived since solutions are non-dominated *relative to the population sample* not its universe; just because a solution dominates all others in the current population does not imply that it lies on the (desired) Pareto front. In the course of a multiobjective optimisation, it is completely

normal for solutions which are non-dominated at some stage in the computation to become dominated by a superior solution at some later stage.

In this paper we present a straightforward method for gauging the convergence of a rank-based multiobjective GA which makes principled comparison of the population states at points in the computation. In the following sections the rationale behind the method is set-out and results presented on both well-known benchmark analytic problems and a real optimisation.

2 - Rationale

If we consider a multiobjective GA with a set of N -objectives, the Pareto front is some general surface in the N -dimensional space of the objectives. For the general case of real-valued objectives the Pareto front comprises an infinite set of points over which the GA performs a search with a finite population. Due to the finite size of the population, from any initialisation of the population there is a finite set of genetic material to be permuted and combined. Towards the end of the GA run most of the chromosomes will become rather similar and so crossover becomes a weak driver for population advancement; most of the gains will effectively be made by random walk due to mutation, albeit 'walking' from a set of promising solutions. At some stage, the rate at which the population improves slows and few further gains of significance are achieved. Hopefully, here the GA will have converged to the Pareto front but conceivably it has got 'stuck' at some sub-optimal point - either way there is little point in continuing the optimisation and it should be terminated. Detecting the corresponding point in a single objective GA is fairly obvious and a number of convergence criteria can be formulated - for example, the scalar score of the best performer in the population being unchanged over some number of iterations. For a multiobjective GA, we need to formulate some comparison over the sets of vector solutions at two stages in the optimisation and to this end we introduce the use of rank histograms.

3 - Rank Histograms

Supposing at two points, $(t-1)$ and t in the optimisation we have two populations, Pop_{t-1} and Pop_t . Solely for the purpose of gauging the state of convergence, we combine the two populations to form $Pop_{t-1} \cup Pop_t$ and rank the resulting union. Taking each rank in turn, we can

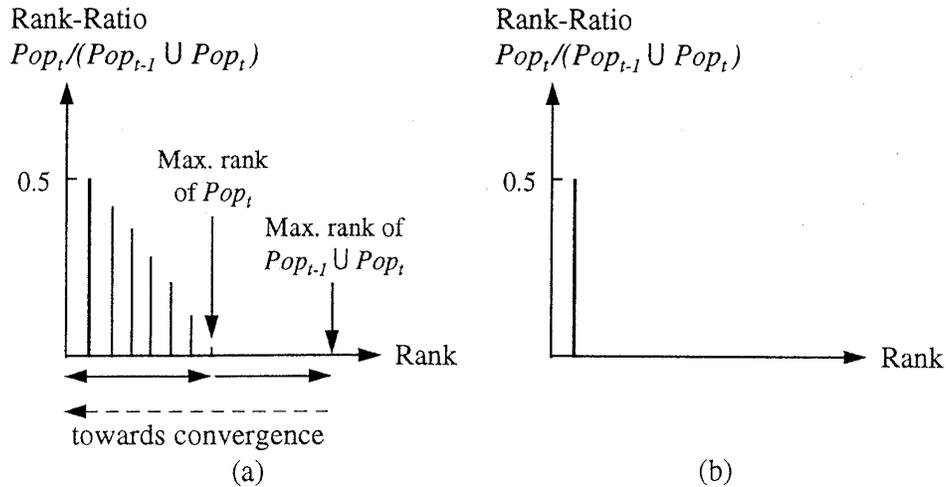


Figure 1: Rank-ratio histogram of a population. (a) The population consists of both dominated and non-dominated individuals and is in an unconverged state. (b) Rank ratio histogram for a converged population state.

generate a histogram of the fraction of the members from Pop_t in $Pop_{t-1} \cup Pop_t$. If the optimisation has progressed to perfect convergence, this *rank ratio histogram* will have a single non-zero entry of 0.5 (Figure 1(b)) in the bin corresponding to unity rank indicating that no solutions superior to those in Pop_{t-1} have been generated in evolving the later generation, Pop_t . Clearly the computation can be stopped at this point. On the other hand, if the optimisation is still far from its converged state then the rank ratio histogram will possess a significant tail of non-zero entries for ranks higher than unity (Figure 1(a)). Thus the emergence of a rank ratio histogram of the required form dictates the earliest useful stopping point for the computation. (In practice the stopping point is judged on the size of the tail above unity rank.)

Continuing the optimisation beyond the point when the rank ratio histogram contains a single non-zero entry results - in our experience - of a redistribution of non-dominated solutions along the Pareto front, usually

leading to a more even sampling [5].

We emphasise that the above stopping criterion does not necessarily imply that the solutions are on the Pareto front. Conceivably the GA has got 'stuck' due to the finite nature of its initial genetic material and the stopping criterion merely denotes the point beyond which further appreciable gain is rather unlikely. We can acquire more substantial evidence as to whether or not the Pareto front has been located by re-running the GA with other randomly initialised states, again stopping using the rank ratio histogram. This leads to a series of populations: $Pop_1, Pop_2, \dots, Pop_m$. If we consider only the non-dominated solutions from $Pop_1, Pop_2, \dots, Pop_m$ and rank them we can again generate a frequency distribution of ranks (Figure 2(a)); if all the component GAs have been run to proper convergence then this *inter-population rank histogram* should contain only one non-zero entry at unity rank meaning that all the non-dominated solutions are equivalent (Figure 2(b)). Although not proof-positive this greatly strengthens the

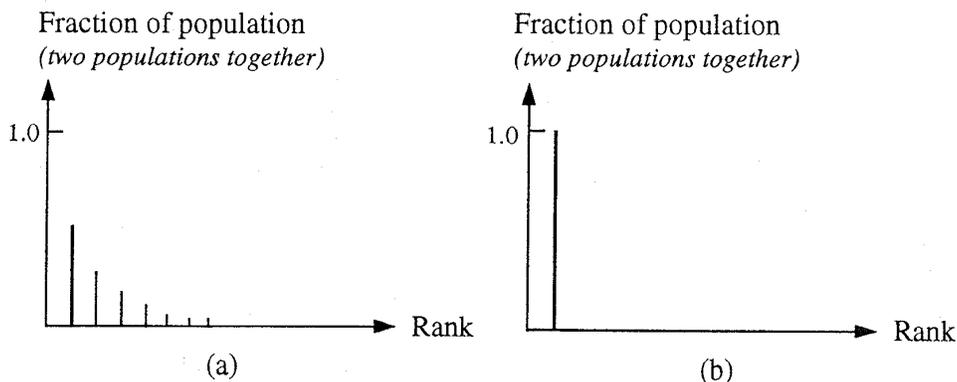


Figure 2: Inter-population rank histograms: (a) Some non-dominated solutions are demoted to dominated, and (b) none of the non-dominated solution is demoted thus the population appears to be in a converged state.

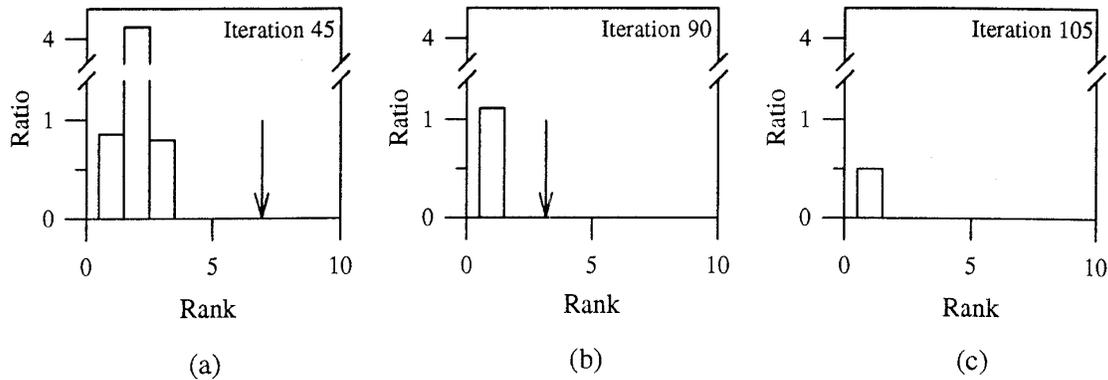


Figure 3: Progression of the rank-ratio histogram for function $F1$ at different stages of population evolution. In (a) and (b) the population is in an unconverged state. At iteration 105, the population has 'converged'. (The down arrow indicates the highest rank of the combined population.)

hypothesis that the true Pareto front has been located. (Again, in practice we seek an inter-population rank histogram with an insignificant tail; any solutions which are not of unity rank are clearly dominated and can be discarded.)

4 - Results

We have investigated the application of rank ratio and inter-population rank histograms to a number of multiobjective genetic optimisations. Firstly, we report the results on two, much-studied analytic multiobjective functions which have the advantage that the location of the Pareto front is known *a priori*. We also give the results from a 'real' optimisation involving an NP-complete partitioning problem in pattern recognition.

Throughout we have used the Pareto Converging Genetic Algorithm (PCGA) [5] although since the present results deal only with convergence, they are applicable to any search procedure which eventually converges. The encodings and other parameters used are identical to those previously used in the literature. For a population size of N , we have calculated the rank ratio histograms at successive points ($N/2$) individual breeding cycles apart.

An important point for comparing the number of iterations required for convergence is that PCGA is a

steady-state not a generational GA.

4.1 - Vincent & Grantham's $F1$ Function [6]

This well-known two objective function in a single variable is fairly straightforward and has been studied previously [3,7]. The progression of the rank ratio histograms for a population size of 30 is shown in Figure 3 in which the downward arrow represents the highest tied rank in the combined population. At iteration 45 there is a significant range of ranks. By iteration 90 (at which point comparison is made with the population at iteration 75) there are only non-dominated individuals present although, as indicated by the down-arrow in Figure 3(b), the union comparison set contains members up to a rank of three. By iteration 105, no new non-dominated solutions have been generated relative to iteration 90 and so the optimisation can be assumed to have reached a logical stopping point. Detailed examination of the populations at each iteration stage confirmed the expected behaviour and that by iteration 105 all the solutions lay on the (known) Pareto front. Continuing the GA run up to 450 iterations produced no change in the rank ratio histogram and served only to redistribute points over the Pareto front producing a more uniformly spaced sampling.

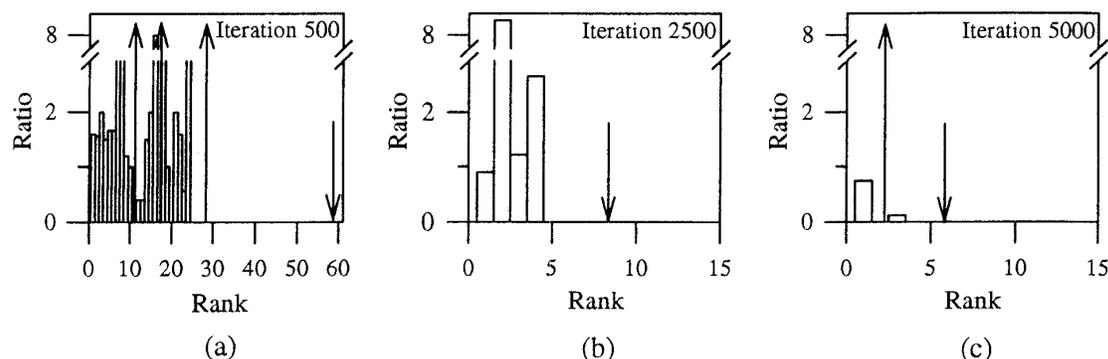


Figure 4: The rank-ratio histograms for function $F2$ at various stages in the GA. The population has not quite converged after 5000 iterations. (Up arrow indicates a division by null entry.)

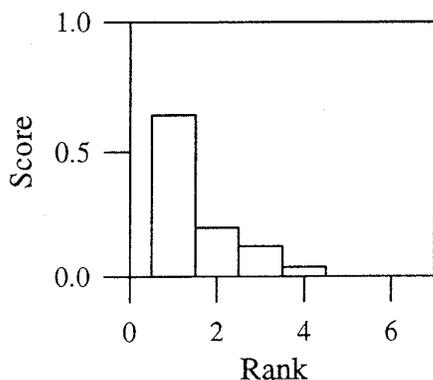


Figure 5: Inter-population rank histogram for function $F2$ formed by combining the non-dominated solutions of two independent runs. The histogram shows that some of the non-dominated individuals are demoted to dominated as indicated by histogram tail.

4.2 - $F2$ Function [8]

The $F2$ function is another analytic n -variable minimisation problem in two symmetrical objectives [8] and is a much harder task than $F1$. For comparison with Fonseca & Fleming [8] we have taken $n = 8$; identical parameters have been used except our PCGA method does not use mating restrictions or sharing. The rank ratio histograms are shown in Figure 4 at iterations: 500, 2500 and 5000. Initially there is a large spread of ranks (Figure 4(a)) but by the time 2500 new individuals (~ 50 new generations) have been evolved (Figure 4(b)) the range of ranks has been greatly reduced. At iteration 5000 (~ 100 new generations) the population is still not quite converged as can be seen from Figure 4(c) and also more directly from viewing the disposition of points in objective space. We have selected this stopping point for direct comparison with the results from Fonseca & Fleming's MOGA technique [8] showing that this optimisation has not quite run to completion; ranks up to six are still present.

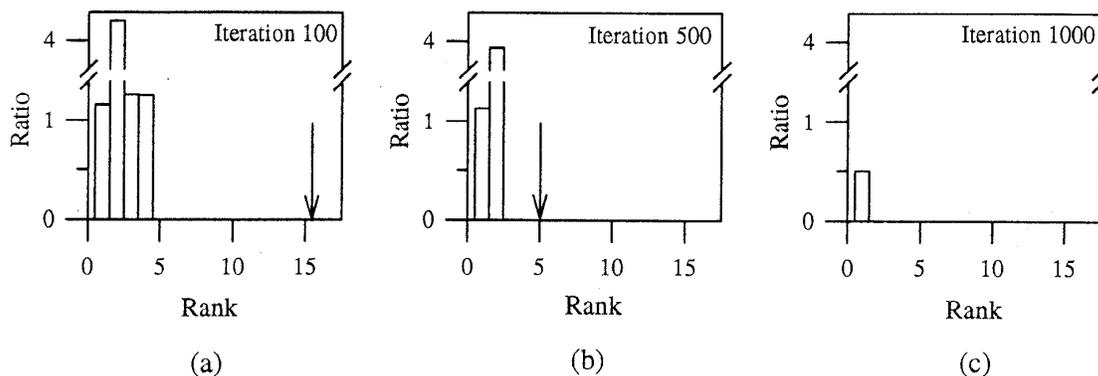


Figure 6: The rank-ratio histograms for an NP-complete pattern space partitioning problem. The population reaches convergence after 1000 iterations. Running for a further 1000 iterations produces the same number of alternative non-dominated solutions.

Further evidence that 5000 iterations are insufficient can be gleaned from the inter-population rank histograms for the $F2$ problem in Figure 5. By combining the non-dominated solutions from two independent GA runs and ranking these, it is apparent from the tail in this histogram that solutions that would have been adjudged non-dominated on the basis of a single GA run have been downgraded to dominated status. Thus inter-population rank histograms act as an additional check on the sufficiency of the stopping criteria.

4.3 - Pattern Space Partitioning Problem

The multiobjective problem examined here is an NP-complete partitioning of a pattern recognition space into an arbitrary number of hyperspheres subject to a number of criteria such as: (1) minimising the hypersphere overlap, (2) minimising the intrinsic dimensionality of each hypersphere, (3) maximising inclusion of data points, (4) maximising the classification rate within hyperspheres, (5) producing the most compact partitioning, and (6) minimising the maximum fraction of included patterns in a single hypersphere. Details of this work will be published elsewhere [9] but it constitutes a rather difficult optimisation. The rank ratio histograms for this problem are shown in Figure 6. For a population size of 200, the population converged in 1000 iterations equivalent to ~ 10 generations of a generational GA. Running the evolution for a further 1000 iterations produced an alternative but rank-equivalent set of non-dominated solutions indicating that the rank-ratio histogram has correctly identified the earliest reasonable stopping point. The inter-population rank histogram is shown in Figure 7 for the combination of the non-dominated individuals from two independent runs of the GA. A significant number of non-dominated individuals are demoted to being dominated during this combination as evidenced by the histogram tail stretching down to rank 27. This forcibly illustrates the point that the rank-ratio histogram test has found only a sensible *stopping point*, not the *whole* Pareto front itself.

5 - Discussion

The notion of convergence of any iterative algorithm can be simply stated as: that at some point no further improvement is obtained. For optimising on a single objective this concept is trivial. This paper has presented a generalisation of this notion to vector optimisation across any arbitrary number of objectives; in this context it is quite natural to employ the concepts Pareto ranking. The use of histograms has a particular purpose in that a complex situation is summarised in a very simple manner. Moreover, the development of the histograms has a straightforward form - from a wide spread of ranks down to (hopefully) a single non-zero entry corresponding to the non-dominated individuals. Although we have not explored the possibility in the present work, it is conceivable that our histogramming approach could be extended to determining if a GA was moving towards convergence in a satisfactory manner on a particularly difficult problem. Insofar as successive histograms represent the rate of convergence, the present methodology could potentially be used to maximise the rate of the evolutionary progress on a given problem with respect to GA-internal parameters such as crossover and mutation rates.

6 - Conclusions

In this paper we have introduced rank histograms both for assessing the state of convergence of a rank-based multiobjective GA and combining evidence about the satisfactory convergence of a series of GA runs. Convergence of a given GA optimisation can be determined from a rank ratio histogram's progression to having a single non-zero entry at unity rank. The equivalence of this condition with locating the Pareto front has been demonstrated for two analytic multiobjective problems although, in general, this condition must be interpreted as the point beyond which little evolutionary gain is likely.

Inter-population rank histograms have been shown to combine the results from a number of GA runs and act as measure of adequacy of the individual contributing optimisations as well as providing additional evidence that the Pareto front for the problem has indeed been located successfully.

7 - Acknowledgements

One of us (RK) gratefully acknowledges support of a Commonwealth Scholarship Commission research studentship.

8 - References

- [1] D. E. Goldberg, *Genetic Algorithms in Search, Optimisation and Machine Learning*, Reading, MA: Addison-Wesley, 1989
- [2] C. M. Fonseca and P. J. Fleming, "An overview of evolutionary algorithms in multiobjective optimisation," *Evolutionary Computation*, 3(1), pp. 1-16, 1995
- [3] J. Horn, N. Nafpliotis, and D. E. Goldberg, "A niched Pareto genetic algorithm for multiobjective

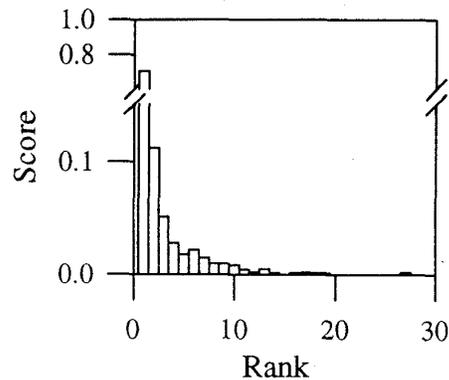


Figure 7: Inter-population rank histogram for the partitioning problem combining the non-dominated solutions of two GA runs to get an improved estimate of the Pareto-front.

optimisation," in *Proc. IEEE Int. Conf. Evolutionary Computation*, 1994, pp. 82-87.

[4] N. Srinivas and K. Deb, "Multiobjective optimisation using nondominated sorting in genetic algorithms," *Evolutionary Computation*, 2(3), pp. 221-248, 1994

[5] R. Kumar and P. I. Rockett, "Improved sampling of the Pareto front in multiobjective genetic optimisations by neo-stationary evolution: A Pareto converging genetic algorithm" Submitted to *IEEE Transactions of Evolutionary Computation*, 1997.

[6] T. L. Vincent and G. J. Grantham, *Optimality in Parametric Systems*, New York: John-Wiley, 1981

[7] J. D. Schaffer, "Multiple objective optimisation with vector evaluated genetic algorithms," in *Proc. 1st Int. Conf. Genetic Algorithms*, J. J. Grefenstette, ed., 1985, pp. 93-100

[8] C. M. Fonseca and P. J. Fleming, "Multiobjective genetic algorithms made easy: selection, sharing and mating restriction," in *Proc. IEE/IEEE Conf. Genetic Algorithms in Engineering Systems: Innovations & Applications (GALESIA 95)*, IEE Conference Publication No. 414, 1995, pp. 45-52.

[9] R. Kumar and P. I. Rockett - To be published.