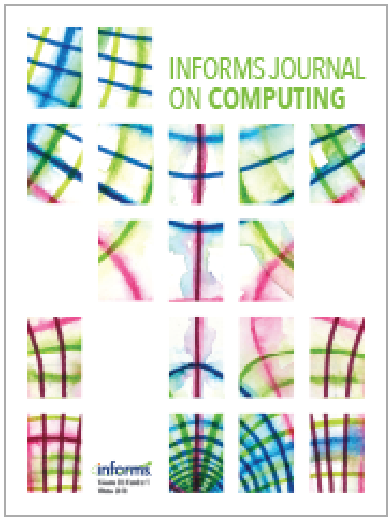

INFORMS Journal on Computing

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Rejoinder—Genetic Algorithms: No Panacea, but a Valuable Tool for the Operations Researcher

Colin R. Reeves,

INFORMS Journal on Computing
Vol. 9, No 3, Summer 1997

0899-1499/97/0903-0263 $05.00

REJOINDER

Genetic Algorithms: No Panacea, but a Valuable Tool for the Operations Researcher

COLIN R. REEVES / *School of Mathematical and Information Sciences, Coventry University, Priory St., Coventry CV1 5FB, United Kingdom, Email: CRReeves@coventry.ac.uk*

When I was invited to contribute this feature article, I did not realize quite how much work I was taking on! Thus, I am naturally pleased that the commentators all found something useful and interesting as a result of my labors, and I would like to thank them for their comments. There is a clear common thread running through the remarks of all of them—the question of how to decide what genetic algorithms (GAs) can do, or alternatively, the related question of how to make them do it better. Focusing on this question was one of my aims in writing the feature article, so the way they have picked this up and added their own insights is very helpful.

Ahuja and Orlin present a very pragmatic approach to the implementation of GAs for operations research—an approach with which I feel very comfortable. From a practical standpoint, it is obviously better to implement techniques that work, rather than to adhere rigidly to a particular conceptual framework. Their appreciation of the place of GAs in the broader family of heuristic methods is also, in my view, an important point. They argue that the interaction between representation and the crossover operator is often of crucial importance. Although this is a frequent empirical observation, and one that seems obvious in principle, it is somewhat harder to pin down exactly how these factors interact in any given case, and how selection also enters into the equation. In [10, 11], we analyzed rather simple bitwise-defined functions in some detail, and found that it is possible for selection to mitigate some of the bad effects of an inappropriate choice of representation and operator. How far such conclusions can be extended to more general combinatorial optimization problems is a matter for conjecture.

Methods based on perturbations of local optima, or on combining information obtained from multiple restarts of hill-climbing procedures, have recently received renewed attention. One explanation for their apparent success in solving combinatorial optimization problems is the "big-valley" hypothesis,[3] an idea that would also apply to the path relinking technique that Ahuja and Orlin discuss. The hypothesis is that the local optima induced by commonly used neighborhood search methods often occur in clusters, and that these clusters are relatively close to the global optimum. Furthermore, it often appears that the better the local optimum is, the closer it is to the global optimum. In [3], this is investigated experimentally for the traveling salesman and graph bisection problems. With some modification to the idea of "distance" between solutions, it also seems to hold for flowshop sequencing.[15]

If true in general, this would also explain some recent successful applications of GAs of the type Ahuja and Orlin describe. These applications are characterized by the replacement of traditional crossover by a more extensive search in a reduced solution space. The fundamental idea of introducing neighborhood search into crossover is not new—it had been proposed earlier in [16] and [14]. However, Ahuja and Orlin's suggestion of taking it to the extent of finding the *optimal* child is a novel idea that clearly works extremely well in the case of the independent set problem.[1] Without going quite that far, finding locally optimal children also seems to produce excellent results in several other cases,[9, 18] especially when, as Ahuja and Orlin suggest, the local optimization steps can be implemented efficiently.

It is interesting that Ahuja and Orlin find initializing GA populations with high-quality solutions is beneficial, as this conflicts with the comments of Levine! Kershenbaum, too, regards it as important that the GA should be initialized with "truly random" solutions. My personal experience is more in line with that of Ahuja and Orlin; perhaps the type of problems we have tried to solve, or the ways in which we have tried to adapt the GA, are different from those of Kershenbaum and Levine. Certainly, with respect to the arguments of Kershenbaum and Levine, if the big valley hypothesis is valid, good local optima are probably not uniformly spread throughout the landscape, and it would make sense to concentrate the search in a likely area from the start, rather than to wait for the GA to discover it. In any case, in many problems the solution space is so large that it is somewhat doubtful that a GA could explore the entire space in any really thorough sense.

Kershenbaum focuses most of his commentary on the important question of representation. His five principles are worth further consideration by anyone who wishes to implement a GA for a particular problem. However, as he argues, it is hard to choose a good encoding a priori. Several attempts can be found in the GA literature to explore the possibility of adapting the encoding dynamically (see [13], for example), but no firm conclusions could be drawn. Some very recent work by Whitley[17] has provided theoretical grounds for believing that Gray coding is, in fact, better for most practical problems than standard binary coding, in the case of bit-string representations of integer or real values. However, Kershenbaum's remarks are still very relevant to discrete optimization.

Kershenbaum also has an interesting point regarding the utility of GAs as a kind of benchmark for other heuristics. Given the relative ease of implementing a GA, this idea makes a lot of sense.

Although the feature article refers briefly to the potential of implementing GAs in parallel, it is an area of which I have no personal experience. I am glad, therefore, that David Levine, who has extensive experience with parallel GAs, has been able to provide some insights into what will surely become an increasingly important area of research and application. I am also happy to concur with the recommendations he makes in his Section 5, in giving some "rules of thumb" for GA implementations. Apart from the fact that I tend to use ranking-based selection, and often experiment with nonstandard crossover operators, all of his prescriptions are ones that I generally adopt for GA applications. For any GA novices who read these commentaries, his recommendations are ones that are worth following.

Finally, I come to the comments of Peter Ross. His examples of successful applications are a valuable addition to those listed in the feature article. He also stresses (as have some of the other commentators) the fact that the genetic algorithm is really a family of approaches. The two examples of "alternative GAs" that he cites—the Breeder GA and CHC—are certainly among the more important members of the family. As Ross argues, crossover seems to be the essence of any method that can properly be termed a genetic algorithm.

However, if we recognize the dual function of crossover that has been argued for in the main article (i.e., that crossover first reduces the search space and then applies an operator), it is perhaps not so far from concepts used by other heuristic techniques. Tabu search, for example, effectively reduces the search space at each iteration by incorporating memory. Multistart strategies, such as those proposed in [3], set out to reduce the search space explicitly, while the idea of *thermostatistical persistency* in simulated annealing[4] tries to focus the search in yet another way. Perhaps we should say then that the distinctive feature of most genetic algorithms is that they usually accomplish this focusing of the search by pairwise "mating" of two solutions. (Some research has been carried out into "multiparent" strategies, but using just two parents is far more common.)

The operator aspect of crossover is somewhat less distinctive, because it depends to a greater extent on what form of crossover is used in a given implementation. However, assuming that traditional one-point crossover is relatively the most popular, the implication is that Culberson's "complementary crossover"[5] is typically the operator used to define neighborhoods in the reduced search space. The fact that this has interesting links to Gray-coded representations and their apparently superior properties, may explain why this often seems to work well for problems using binary-coded continuous variables.

The circumstances under which this approach is appropriate may not be as wide as some early proponents of GAs had hoped, as is discussed in the paper that Ross references.[7] Certainly, the case of combinatorial problems, such as those discussed in the feature article, shows that GAs cannot be applied blindly if we want to be successful. Ross also has some sensible things to say about the No Free Lunch theorem, which are clearly related to the question of when particular GA strategies should be used. Unfortunately, general answers to this question cannot yet be given, but by continued care in implementing GAs, and by thoughtful evaluation of the results obtained, we may hope to start to find some answers.

Finally, Ross raises an interesting general point about the relationship of GAs to Darwinian evolution. One sometimes hears this metaphor being put forward almost as an axiomatic justification for treating GAs as if they really were "the only game in town," to the exclusion of other techniques. The argument is that Darwinian evolution has taken us from the primordial soup to space travel, so GAs must be the way to go. Hardly anyone doubts that species do indeed adapt to their environments, but as Ross points out, some scientists (for example, Goodwin[8] or Kauffman[12]) are saying that the neo-Darwinian scheme of random variation and natural selection alone (as exemplified, for example, by the writings of Dawkins[6]) is inadequate as a general explanation for the complexity and order of the world around us. Behe[2] has recently argued even more provocatively that many molecular biological systems are so complex that no neo-Darwinian "explanation" is even conceivable.

Has this discussion strayed too far from genetic algorithms and operations research? Not really, for it suggests that an appeal to "evolution" as a justification for using GAs is ill-advised. The basic concepts of selection, recombination, and mutation may be useful in modeling (albeit in a rather imprecise and rudimentary fashion) the way systems adapt to their environments, but where the concern is optimization, these concepts alone are unlikely to be effective in more than a few cases. In fact, as Ross says, it is amazing that it works at all. Perhaps the claims for the universality of GAs as optimization tools are less common than they were, but Levine's observation, that a common complaint by the practitioners whom he questioned related to the dangers of inappropriate usage of GAs, also shows that a sober and realistic appraisal of GAs and their abilities is still needed. At the same time, if the pendulum *is* swinging away from Darwinian concepts, too close an identification with Darwinism could provide an excuse for GA-skeptics (and some do exist) to write them off. This would be equally misguided. We need to recognize that GAs do bring some advantages, and are indeed a valuable addition to our arsenal of techniques for optimization.

In my feature article I have tried to maintain this balance, and the remarks of all the commentators have been helpful in this regard. It is my hope that the readers of the *Journal on Computing* will find the entire package a worthwhile and useful aid in developing their own research and applications.

References

1. C.C. AGGARWAL, J.B. ORLIN, and R.P. TAI, 1997. Optimized Crossover for the Independent Set Problem, *Operations Research* 45, 226–234.
2. M. BEHE, 1996. *Darwin's Black Box: The Biochemical Challenge to Evolution*, Free Press, New York.

3. K.D. BOESE, A.B. KAHNG, and S. MUDDU, 1994. A New Adaptive Multi-start Technique for Combinatorial Global Optimizations, *Operations Research Letters 16*, 101–113.
4. P. CHARDAIRE, J.L. LUTTON, and A. SUTTER, 1995. Thermostatistical Persistency: A Powerful Improving Concept for Simulated Annealing Algorithms, *European Journal of Operational Research 86*, 565–579.
5. J.C. CULBERSON, 1995. Mutation-Crossover Isomorphisms and the Construction of Discriminating Functions, *Evolutionary Computation 2*, 279–311.
6. R. DAWKINS, 1986. *The Blind Watchmaker*, Longman, London.
7. L.J. ESHELMAN and J.D. SCHAFFER, 1995. Crossover's Niche, in *Proceedings of 5th International Conference on Genetic Algorithms*, S. Forrest (ed.), Morgan Kaufmann, San Mateo, CA, 9–14.
8. B. GOODWIN, 1995. *How the Leopard Changed Its Spots*, Weidenfeld and Nicholson, London.
9. C. HÖHN and C.R. REEVES, 1996. Graph Partitioning using Genetic Algorithms, in *Proceedings of the 2nd Conference on Massively Parallel Computing Systems*, G.R. Sechi (ed.), IEEE Computer Society, Los Alamitos, CA, 31–38.
10. C. HÖHN and C.R. REEVES, 1996. The Crossover Landscape for the *onemax* Problem, in *Proceedings of the 2nd Nordic Workshop on Genetic Algorithms and their Applications*, J. Alander (ed.), University of Vaasa Press, Vaasa, Finland, 27–43.
11. C. HÖHN and C.R. REEVES, 1996. Are Long Path Problems Hard for Genetic Algorithms? in *Parallel Problem-Solving from Nature—PPSN IV*, H-M. Voigt, W. Ebeling, I. Rechenberg, and H-P. Schwefel (eds.), Springer-Verlag, Berlin, 134–143.
12. S. KAUFFMAN, 1995. *At Home in the Universe*, Penguin, London.
13. K. MATHIAS and D. WHITLEY, 1993. Remapping Hyperspace during Genetic Search: Canonical Delta Folding, in *Foundations of Genetic Algorithms 2*, D. Whitley (ed.), Morgan Kaufmann, San Mateo, CA, 167–186.
14. V.J. RAYWARD-SMITH, 1995. A Unified Approach to Tabu Search, Simulated Annealing and Genetic Algorithms, in *Applications of Modern Heuristic Techniques*, V.J. Rayward-Smith (ed.), Alfred Waller, Henley-on-Thames, UK, 17–38.
15. C.R. REEVES, 1997. Landscapes, Operators and Heuristic Search, to appear in *Annals of Operations Research*.
16. C.R. REEVES, 1994. Genetic Algorithms and Neighbourhood Search, in *Evolutionary Computing: AISB Workshop, Leeds, UK, April 1994; Selected Papers*, T.C. Fogarty (ed.), Springer-Verlag, Berlin, 115–130.
17. D. WHITLEY, 1997. Personal communication.
18. T. YAMADA and R. NAKANO, 1996. Job-Shop Scheduling, in *Genetic Algorithms in Engineering Systems*, P.J. Fleming and A. Zalzala (eds.), Peter Peregrinus, London.