Editorial Introduction

Welcome to the fourth and final issue of the first volume of *Evolutionary Computation*! Not unlike parenthood, this first year with an infant journal has been both hectic and rewarding. I would like to thank the associate editors, the editorial board, and the staff at MIT Press for all the time and effort they have put into making this a successful first year.

As we move into the second year, we need to continue to expand the journal in several directions. I would like to see the inclusion of short technical notes, book reviews, and summaries of special workshops. With the emergence of worldwide information networks, we need to explore electronic modes of abstracting and distributing information. I'm sure you have many innovative ideas, and I encourage you to send me an e-mail note outlining them.

One of the most encouraging trends I have seen at conferences this past year is the dramatic increase in the number and variety of applications of evolutionary algorithms. The first paper in this issue is representative of the diversity and maturity of these applications. Lienig and Thulasiraman present an innovative and successful application of genetic algorithms to difficult VLSI layout problems. Their results either match or exceed the best published results for channel routing problems, and their architecture serves as a model for attacking other routing problems with GAs.

Traveling salesman problems (TSPs) continue to serve as benchmarks for testing combinatorial optimization algorithms, and they have been approached in a variety of ways by the EC community. TSPs are my favorite example for illustrating the importance of choosing a good representation and effective operators when applying GAs to hard problems. To date, there have been many adequate, but no really outstanding applications of GAs to TSP problems that are capable of extending the state-of-the-art.

The second paper, by Valenzuela and Jones, appears to have the potential to do so. They have developed an interesting hybrid approach in which a GA evolves increasingly better "divide and conquer" partitions of the Euclidean plane, rather than the tours themselves. At the leaf nodes of the partition, a local optimizer finds optimal subtours, and a repair mechanism joins the subtours into a global tour as one ascends the partition tree. The algorithm has good ($n \log n$) scaleup properties and is shown to be effective on problems involving thousands of cities.

One of the continuing discussions (controversies) in both evolutionary biology and evolutionary computation is the extent to which it is useful to view evolutionary forces as an optimization process (as opposed to emphasizing adaptation, complex co-evolution, etc.). The use of evolutionary algorithms to solve difficult optimization problems has been viewed as a "natural" application from the very beginnings of the field of evolutionary computation. However, when we build such applications, we invariably add features of our own design to "improve" performance. The net effect of all this is that we still do not have strong predictive theories about convergence and rates of convergence of our optimization-oriented evolutionary algorithms, except for the simplest of problems.

The third paper in this issue, by Mühlenbein and Schlierkamp-Voosen, presents new evidence that the science of breeding rather than the process of natural evolution may be a better model on which to base optimization-oriented applications. They present a strong predictive theory for their "breeder" GA based on the principles of "response to selection" and "heritability," which are well studied in the breeding literature. The theory is developed initially for the simpler case of additive gene effects, and then extended to cover interacting genes.

A frequently asked question is: What's the relationship between evolutionary algorithms and simulated annealing (SA) algorithms? Both seem to be useful for solving large optimization problems. Can we say anything about their relative strengths and weaknesses, and identify important similarities and differences? The final paper in this issue, by Rudolph, directly addresses these questions by presenting a formal framework that captures both approaches as instances of a common "Markovian optimization" model. The first part of the paper discusses the relationships between SA and (1 + 1)-ES models and $(1, p)$-ES models, which are easier to analyze but not used much in practice. The latter sections extend these ideas to spatially distributed evolutionary algorithms with population sizes $\gg 1$ and lead to an interesting hybrid SA/EA algorithm with significant potential for solving difficult multimodal optimization problems.

Volume 2 of the journal is already taking shape. It will contain our first special issue, the focus of which is classifier systems, edited by Rob Smith and Manuel Valenzuela-Rendón. Also in the pipeline are papers involving genetic programming, a variety of interesting EC applications, and self-adaptive evolutionary systems, to mention a few.

Enjoy!

Kenneth De Jong
George Mason University
Fairfax, VA 22030 USA
E-mail: ecj@cs.gmu.edu