
On The State of Evolutionary Computation

Kenneth De Jong
George Mason University
Fairfax, VA 22030
kdejong@aic.gmu.edu

William Spears
Naval Research Laboratory
Washington, DC 20375
spears@aic.nrl.navy.mil

Abstract

In the past few years the evolutionary computation landscape has been rapidly changing as a result of increased levels of interaction between various research groups and the injection of new ideas which challenge old tenets. The effect has been simultaneously exciting, invigorating, annoying, and bewildering to the old-timers as well as the new-comers to the field. Emerging out of all of this activity are the beginnings of some structure, some common themes, and some agreement on important open issues. We attempt to summarize these emergent properties in this paper.

1 INTRODUCTION

The field of evolutionary computation (EC) is in a stage of tremendous growth as witnessed by the increasing number of conferences, workshops, and papers in the area as well as the emergence of a central journal for the field. Until recently, the field could be characterized as a fairly amorphous collection of independent research groups each representing a particular approach to evolutionary computation with little interaction in the form of migration and cross-fertilization of ideas.

However, in the past few years the evolutionary computation landscape has been rapidly changing. There has been a significant increase in the levels of interaction between various research groups. New arrivals to the field have resulted in the injection of new ideas which challenge old tenets. The effect has been exciting and invigorating, annoying at times, and in some cases bewildering to the old-timers as well as the new-comers to the field.

Emerging out of all of this activity are the beginnings of some structure, some common themes, and some agreement on important open issues. We attempt to summarize these emergent properties in the remainder of the paper.

2 A HISTORICAL VIEW OF THE FIELD

The most obvious and straightforward way to characterize the field is in terms of its historical evolution. There have been three well-defined paradigms which have served as the basis for much of the activity in the field: genetic algorithms (GAs), evolution strategies (ESs), and evolutionary programming (EP). Each of these paradigms has acquired an admittedly over simplified characterization of its basic tenets which, in spite of its shortcomings, serves as a useful starting point for the subsequent discussions of issues and open questions.

GAs owe their name to an early emphasis on representing and manipulating individuals in terms of their genetic makeup rather than using a phenotypic representation. Much of the early work used a universal internal representation involving fixed-length binary strings with "genetic" operators like mutation and crossover defined to operate in a domain-independent fashion at this level without any knowledge of the phenotypic interpretation of the strings (Holland, 1975; De Jong, 1975). This universality was also reflected in a strong emphasis on the design of robust adaptive systems with a broad range of applications. Equally important was the early emphasis on theoretical analysis resulting in "the schema theorem" and characterizations of the role and importance of crossover.

By contrast, ESs were developed with a strong focus on building systems capable of solving difficult real-valued parameter optimization problems (Rechenberg, 1973; Schwefel, 1981). The "natural" representation was a vector of real-valued "genes" which were manipulated primarily by mutation operators designed to perturb the real-valued parameters in useful ways. Analysis played a strong role here as well with initial theorems on convergence to global optima, rates of convergence, and other ES properties such as the "1/5" rule.

Universality was also a central theme of the early work on EP. The direction this took was the idea of representing individuals phenotypically as finite state machines capable of responding to environmental stimuli, and developing operators (primarily mutation) for effecting structural and behavioral change over time (Fogel et al, 1966). These ideas were then applied to a

broad range of problems including prediction problems, optimization, and machine learning.

These early characterizations, however, are no longer all that useful in describing the enormous variety of current activities on the field. GA practitioners are seldom constrained to universal fixed-length binary implementations. ES practitioners have incorporated recombination operators into their systems. EP is used for more than just the evolution of finite state machines. Provocative new terms and ideas like "messy GAs" (Goldberg, 1991) and "genetic programming" (de Garis, 1990; Koza, 1992) have emerged on the landscape.

This suggests that we need to rethink how we think about evolutionary computation. Categorizing things according to historical roots appears to have outgrown its usefulness. Rather, we need to identify important dimensions along which the broader field of EC can be expanded and evaluated. In the remainder of the paper, we present a first cut at identifying such properties and also suggest potentially useful dimensions which are currently unexplored.

3 CENTRAL EC DIMENSIONS

In this section we discuss the most obvious candidates for important dimensions of the field. The more speculative dimensions are deferred to next section.

3.1 CHARACTERISTICS OF FITNESS LANDSCAPES

The majority of the EC work to date has been with problem domains in which the fitness landscape is time-invariant and the fitness of individuals can be computed independently from other members of the current population. This is a direct result of the pervasiveness of optimization problems and the usefulness of evolutionary algorithms (EAs) in solving them. This has led to considerable insight into the behavior of EAs on such surfaces including such notions as "GA-easy", "GA-hard", and "deception".

However, most of these insights evaporate if we attack problem classes for which the fitness landscapes violate one or more of our traditional assumptions. There are at least three important problem classes for which results are badly needed: autonomously changing landscapes, the evolution of cooperative behavior, and ecological problems.

Problems involving autonomously changing landscapes frequently arise when fitness is defined in terms of one or more autonomous agents in the environment whose behavior can change independently of any of the search activity of an EA. Typical examples are mechanical devices which age, breakdown, etc, or changes in weather patterns which dramatically change the "fitness" of a particular ship on the open sea. If we apply typical optimization-oriented EAs to such problems, the strong

pressures to converge generally result in a loss of the population diversity needed to respond to such changes. We currently have very little insight regarding how to design EAs for such problems.

Holland's classifier systems (1986) and iterated prisoner's dilemma problems are examples of problems in which fitness is a function of how well an individual complements other individuals in the population. Rather than searching for a single optimal individual, the goal is to evolve groups of individuals (generalists, specialists, etc.) which collectively solve a particular problem.

If we apply typical optimization-oriented EAs to such problems, the strong pressures towards homogeneity in the population make it difficult to maintain different but cooperative individuals. Additional mechanisms for rewarding groups of individuals seem to be required (e.g., bucket brigades, profit sharing), but we have little in the way of theory to guide us.

Ecology-oriented problems present a third and perhaps most difficult class of landscapes in which the shape of the fitness landscape is directly affected by the evolutionary process itself. Perhaps a better way to think of this is in co-evolutionary terms in which multiple interacting evolutionary processes are at work modeling the availability of resources, prey-predator relationships, host-parasite interactions, etc. Very few of our insights from the optimization world appear to carry over here.

The interest in using EAs to solve problems like these which violate traditional assumptions continues to grow. We already have examples of EAs which are powerful function optimizers, but which are completely ineffective for evolving cooperative behavior or tracking a changing landscape. Modified EAs are now being developed for these new problem classes, but are also much less useful as traditional optimizers. These developments have created both the need and the opportunity to gain a deeper understanding of the behavior of EAs.

3.2 INTERNAL REPRESENTATIONS

One of the most critical decisions made in applying evolutionary techniques to a particular class of problems is the specification of the space to be explored by an EA. This is accomplished by defining a mapping between points in the problem space and points in an internal representation space.

The EC community differs widely on opinions and strategies for selecting appropriate representations, ranging from universal binary encodings to problem-specific encodings for TSP problems and real-valued parameter optimization problems. The tradeoffs are fairly obvious in that universal encodings have a much broader range of applicability, but are frequently outperformed by problem-specific representations which require extra effort to implement but exploit additional knowledge about a particular problem class.

Although there are strong historical associations between GAs and binary string representations, between ESs and vectors of real numbers, and between EP and finite state machines, we can clearly imagine (and, in fact, have many examples of) using representations other than the traditional ones. Claiming one EA approach is better than another on a particular class of problems is not meaningful any more without motivating and specifying (among other things) the representations chosen.

What is needed, but has been difficult to obtain, are theoretical results on representation theory. Holland's schema analysis (1975) and Radcliffe's (1991) generalization to formae are examples of how theory can help guide representation choices. The recent work on "fitness correlation" (Manderick, et al, 1991) suggests another approach. Clearly, much more work is required along this important dimension if effective representations are to be easily selectable.

3.3 MODELING THE DYNAMICS OF POPULATION EVOLUTION

At a high level of abstraction we think of evolutionary processes in terms of the ability of more fit individuals to have a stronger influence on the future makeup of the population by surviving longer and by producing more offspring which continue to assert influence after the parents have disappeared. How these notions are turned into computational models varies quite dramatically within the EC community. This variance hinges on several important design decisions discussed briefly in the following subsections.

3.3.1 Population Size Issues

Most current EAs assume a constant population size N which is specified as a user-controlled input parameter. So called "steady state" EAs rigidly enforce this limit in the sense that each time an offspring is produced resulting in $N+1$ individuals, a selection process is invoked to reduce the population size back to N . By contrast, "generational" EAs permit more elasticity in the population size by allowing $K \gg 1$ offspring to be produced before a selection process is invoked to delete K individuals.

Although we understand that the size of an EA's population can affect its ability to solve problems, we have only the beginnings of a theory strong enough to provide *a priori* guidance in choosing an appropriate fixed size (e.g., (Goldberg, 1989a)), not much theory regarding appropriate levels of elasticity (K), and even less understanding as to the merits of dynamically adjusting the population size.

3.3.2 Deletion Strategies

The processes used to delete individuals varies significantly from one EA to another and includes strategies such as uniform random deletion, deletion of the K worst, and inverse fitness-proportional deletion. It is clear that "elitist" deletion strategies which are too strongly biased towards removing the worst lead can lead to premature loss of diversity and suboptimal solutions. It is equally clear that too little fitness bias results in unfocused and meandering search. Finding a proper balance is important but difficult to determine *a priori* with current theory.

3.3.3 Parent Selection

Similar issues arise with respect to choosing which parents will produce offspring. Biasing the selection too strongly towards the best individuals results in too narrow a search focus, while too little bias produces a lack of needed focus. Current methods include uniform random selection, rank-proportional selection, and fitness-proportional selection.

Clearly, the parent selection and individual deletion strategies must complement each other in terms of the overall effect they have on the exploration/exploitation balance. We have some theory here for particular cases such as Holland's "optimal allocation of trials" characterization of traditional GAs (Holland, 1975), and the "1/5" rule for ESs (Schwefel, 1981), but much stronger results are needed.

3.3.4 The Mechanisms of Reproduction

In addition to these selection processes, the mechanisms used for reproduction also affect the balance between exploration and exploitation. At one extreme one can imagine a system in which offspring are exact replicas of parents (asexual reproduction with no mutation) resulting in rapid growth in the proportions of the best individuals in the population, but with no exploration beyond the initial population members. At the other extreme, one can imagine a system in which the offspring have little resemblance to their parents, maximizing exploration at the expense of inheriting useful parental characteristics.

The EC community has focused primarily on two reproductive mechanisms which fall in between these two extremes: 1-parent reproduction with mutation and 2-parent reproduction with recombination and mutation. Historically, the EP and ES communities have emphasized the former while the GA community has emphasized the latter.

However, these traditional views are breaking down rapidly. The ES community has found recombination to be useful, particularly in evolving adaptive mutation rates (Bäck et al, 1991). Various members of the GA community have reported improved results by not using

recombination (de Garis, 1990), by not using mutation (Koza, 1992), or by adding new and more powerful mutation operators (Eshelman and Schaffer, 1991).

As before, we have the tantalizing beginnings of a theory to help understand and guide the use and further development of reproductive mechanisms. Beginning with Holland's initial work (1975), the GA community has analyzed in considerable detail the role of crossover and mutation (see, for example, (Goldberg, 1989b), (Vose and Liepins, 1991), (Booker, 1992), (De Jong and Spears, 1992), or (Spears, 1992a)). The ES community has developed theoretical models for optimal mutation rates with respect to convergence and convergence rates in the context of function optimization (Schwefel, 1981).

However, the rapid growth of the field is pressing these theories hard with "anomalous results" (Forrest and Mitchell, 1993) and new directions not covered by current theory. One of the important issues not well understood is the benefit of adaptive reproductive operators. There are now a variety of empirical studies which show the effectiveness of adaptive mutation rates (e.g., (Fogarty, 1989), (Bäck et al, 1991), or (Fogel and Atmar, 1992)) as well as adaptive recombination mechanisms (e.g., (Schaffer and Morishima, 1987), (Davis, 1989), or (Spears, 1992b)).

4 NEW AND POTENTIALLY IMPORTANT DIMENSIONS FOR EC RESEARCH

In the previous section, we summarized the current state of the art with respect to central EC issues and indicated where additional research on these issues is required. In this section, we discuss some more speculative areas which are likely to play a more central role in the near future.

4.1 REPRESENTATION AND MORPHOGENESIS

In the earlier section on representation issues we discussed the tradeoffs between problem-independent and problem-specific representations. Closely related to this is the biological distinction between the more universal genotypic descriptions of individuals in the form of plans for generating them and the phenotypic descriptions of the actual generated structures.

Historically, much of the EA work has involved the evolution of fairly simple structures could be represented in phenotypic form or be easily mapped onto simple genotypic representations. However, as we attempt to evolve increasingly more complex structures (e.g., Lisp code (Koza, 1992) or neural networks (de Garis, 1990)), it becomes increasingly difficult to define forms of mutation and recombination which are capable of producing structurally sound and interesting new individuals. If we look to nature for inspiration, we don't

see many evolutionary operators at the phenotype level (e.g., swapping arms and legs!). Rather, changes occur at the genotype level and the effects of those changes instantiated via growth and maturation. If we hope to evolve such complexity, we may need to adopt more universal encodings coupled with a process of morphogenesis (e.g., see Dawkins (1987) or Harp, et al (1989)).

4.2 INCLUSION OF LAMARCKIAN PROPERTIES

Although EAs may be inspired by biological systems, many interesting properties arise when we include features not available to those systems. One common example is the inclusion of Lamarckian operators, which allow the inheritance of characteristics acquired during the lifetime of an individual.

In the EC world this is beginning to show up in the form of hybrid systems in which individuals themselves go through a learning and/or adaptation phase as part of their fitness evaluation, and the results of that adaptation are passed on to their offspring (e.g., see Grefenstette (1991)).

Although initial empirical results are encouraging, we presently have no good way of analyzing such systems at a more abstract level.

4.3 NON-RANDOM MATING AND SPECIATION

Currently, most EAs incorporate a random mating scheme in which the species or sex of an individual is not relevant. One problem with this, as with real biological systems, is that the offspring of parents from two species are often not viable. As we move to more complex systems which attempt to evolve cooperating behavior and which may have more than one evolutionary process active simultaneously, the roles of non-random mating and speciation will become an important issue.

Some solutions to these problems have been suggested, such as crowding (De Jong, 1975), sharing (Goldberg, 1987), and tagging (Booker, 1982). Unfortunately, these solutions tend to make fairly strong assumptions, such as the number of species and/or the distribution of niches in the environment. For some problems these assumptions are reasonable. However, in many cases such properties are not known *a priori* and must evolve as well.

4.4 DECENTRALIZED, HIGHLY PARALLEL MODELS

Because of the inherent natural parallelism within an EA, much recent work has concentrated on the implementation of EAs on both fine and coarse grained parallel machines. Clearly, such implementations hold promise of significant decreases in the execution time of EAs.

More interestingly, though, for the topic of this paper, are the evolutionary effects that can be naturally implemented with parallel machines, namely, speciation, nicheing, and punctuated equilibria. For example, non-random mating may be easily implemented by enforcing parents to be neighbors with respect to the topology of the parallel architecture. Species emerge as local neighborhoods within that topology. Subpopulations in equilibrium are "punctuated" by easily implemented migration patterns from neighboring subpopulations.

However, each such change to an EA significantly changes its semantics and the resulting behavior. Our admittedly weak theory about traditional EAs is currently of little help in understanding these parallel implementations.

4.5 SELF-ADAPTING SYSTEMS

Another theme that has been arising with increasing frequency is the inclusion of self-adapting mechanisms with EAs to control parameters involving the internal representation, mutation, recombination, and population size. This trend is due in part to the absence of strong predictive theories which specify such things *a priori*. It is also a reflection of the fact that EAs are being applied to more complex and time-varying fitness landscapes.

Some important issues that need to be solved involve the self-adaptation mechanism itself. For example, do we use an EA or some other mechanism? If we use an EA, how do we use fitness as a performance feedback for self-adaptation?

On a positive note, the EC community has already empirically illustrated the viability of self-adaptation of mutation and recombination as noted earlier, as well as adaptive representations like Argot (Shaefer, 1987), messy GAs (Goldberg, 1991), the dynamic parameter encoding (DPE) scheme of Schraudolph and Belew (1992), and the Delta coding of Whitley et al (1991).

4.6 CO-EVOLUTIONARY SYSTEMS

Hillis' recent work (1990) on the improvements achievable by co-evolving parasites along with the actual individuals of interest gives an exciting glimpse of the behavioral complexity and power of such techniques. Holland's Echo system (1992) reflects an even more complex ecological setting with renewable resources and predators.

Both systems suggest an important future role for co-evolution in EAs, but raise more questions than they answer concerning a principled method for designing such systems as well as the kinds of problems for which this additional level of complexity is both necessary and effective.

5 SUMMARY AND CONCLUSIONS

This is an exciting time for the EC field. The increased level of EC activity has resulted in an infusion of new ideas and applications which are challenging old tenets and invalidating historical categorizations. In spite of this rapidly changing EC landscape, some common structures and agreement on important issues is beginning to emerge.

We have presented a framework containing both basic and more speculative dimensions which are useful for both characterizing the current state of the field, and also point out important open issues which need further research. We believe that a view of this sort is an important and necessary part of the continued growth of the field.

References

- Bäck, T., Hoffmeister, F., & Schwefel, H.-P. (1991) A survey of evolution strategies. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 2-9. La Jolla, CA: Morgan Kaufmann.
- Belew, R. K., & Booker, L. B. (eds.) (1991) *Proceedings of the Fourth International Conference on Genetic Algorithms*. La Jolla, CA: Morgan Kaufmann.
- Booker, L. B. (1982) *Intelligent Behavior as an adaptation to the task environment*. Doctoral Thesis, Department of Computer and Communication Sciences. University of Michigan, Ann Arbor.
- Booker, L. B. (1992) Recombination distributions for genetic algorithms. *Proceedings of the Foundations of Genetic Algorithms Workshop*. Vail, CO: Morgan Kaufmann.
- Davis, L. (1989) Adapting operator probabilities in genetic algorithms. *Proceedings of the Third International Conference on Genetic Algorithms*, 60-69. La Jolla, CA: Morgan Kaufmann.
- Dawkins, R. (1987) *The Blind Watchmaker*, New York, W. W. Norton and Company.
- de Garis, H. (1990) Genetic programming: modular evolution for darwin machines. *Proceedings of the 1990 International Joint Conference on Neural Networks*, 194-197. Washington, DC: Lawrence Erlbaum.
- De Jong, K. A. (1975) *An analysis of the behavior of a class of genetic adaptive systems*. Doctoral Thesis, Department of Computer and Communication Sciences. University of Michigan, Ann Arbor.
- De Jong, K. & Spears, W. (1992) A formal analysis of the role of multi-point crossover in genetic algorithms", *Annals of Mathematics and Artificial Intelligence Journal*, Volume 5, #1, pgs. 1-26.
- Eshelman, L. J., & Schaffer, J. D. (1991) Preventing premature convergence in genetic algorithms by preventing incest. *Proceedings of the Fourth*

- International Conference on Genetic Algorithms*, 115-122. La Jolla, CA: Morgan Kaufmann.
- Fogarty, T. C. (1989) Varying the probability of mutation in the genetic algorithm *Proceedings of the Third International Conference on Genetic Algorithms*, 104-109. Fairfax, VA: Morgan Kaufmann.
- Fogel, L. J., Owens, A. J., & Walsh, M. J. (1966) *Artificial Intelligence Through Simulated Evolution*. New York: Wiley Publishing.
- Fogel, D. B., & Atmar, J. W. (eds.) (1992) *Proceedings of the First Annual Conference on Evolutionary Programming*. La Jolla, CA: Evolutionary Programming Society
- Forrest, S. & Mitchell, M. (1993) What makes a problem hard for a genetic algorithm? Some anomalous results and their explanation. To appear in *Machine Learning Journal*.
- Goldberg, D. E. & Richardson, J. (1987) Genetic algorithms with sharing for multimodal function optimization. *Proceedings of the Second International Conference on Genetic Algorithms*, 41-49. Cambridge, MA: Lawrence Erlbaum.
- Goldberg, D. E. (1989a) Sizing populations for serial and parallel genetic algorithms. *Proceedings of the Third International Conference on Genetic Algorithms*, 70-79. Fairfax, VA: Morgan Kaufmann.
- Goldberg, D. E. (1989b) *Genetic Algorithms in Search, Optimization & Machine Learning*. Reading, MA: Addison-Wesley.
- Goldberg, D. E., Deb, K., & Korb, B. (1991) Don't worry, be messy. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 24-30. La Jolla, CA: Morgan Kaufmann.
- Grefenstette, J. J. (1991) Lamarckian learning in multi-agent environments. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 303-310. La Jolla, CA: Morgan Kaufmann.
- Harp, S., Samad, T., & Guha, A. (1989) Towards the genetic synthesis of neural networks. *Proceedings of the Third International Conference on Genetic Algorithms*, 360-369. Fairfax, VA: Morgan Kaufmann.
- Hillis, D. W. (1990) Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D* 42, 228-234.
- Holland, J. H. (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor, Michigan: The University of Michigan Press.
- Holland, J. H. (1992) *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press. Second Edition.
- Holland, J. H. (1986) Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. Michalski, J. Carbonell, T. Mitchell (eds.), *Machine Learning: An Artificial Intelligence Approach*. Los Altos: Morgan Kaufmann.
- Koza, J. R. (1992) *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.
- Manderick, B., de Weger, M., & Spiessens, P. (1991) The genetic algorithm and the structure of the fitness landscape. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 222-229. La Jolla, CA: Morgan Kaufmann.
- Radcliffe, N. J. (1991) Forma analysis and random respectful recombination. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 222-229. La Jolla, CA: Morgan Kaufmann.
- Rechenberg, I. (1973) *Evolutionsstrategie: Optimierung Technischer Systeme nach Prinzipien der Biologischen Evolution*. Frommann-Holzboog, Stuttgart.
- Schaffer, J. D. & Morishima, A. (1987) An adaptive crossover distribution mechanisms for genetic algorithms. *Proceedings of the Second International Conference on Genetic Algorithms*, 36-40. Cambridge, MA: Lawrence Erlbaum.
- Schraudolph, N. N., & Belew, R. K. (1992) Dynamic parameter encoding for genetic algorithms. *Machine Learning Journal*, Volume 9, Number 1, 9-22.
- Schwefel, H.-P. (1981) *Numerical Optimization of Computer Models*. New York: John Wiley & Sons.
- Shaefer, C. G. (1987) The ARGOT strategy: adaptive representation genetic optimizer technique. *Proceedings of the Second International Conference on Genetic Algorithms*, 50-58. Cambridge, MA: Lawrence Erlbaum.
- Spears, W. M. (1992a) Crossover or mutation? *Proceedings of the Foundations of Genetic Algorithms Workshop*, Vail, Colorado: Morgan Kaufmann.
- Spears, W. M. (1992b) Adapting crossover in a genetic algorithm. *Naval Research Laboratory AI Center Report AIC-92-025*. Washington, DC 20375 USA.
- Vose, M. D., & Liepins, G. E. (1991) Schema disruption. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 237-242. La Jolla, CA: Morgan Kaufmann.
- Whitley, D., Mathias, K., & Fitzhorn, P. (1991) Delta coding: an iterative search strategy for genetic algorithms. *Proceedings of the Fourth International Conference on Genetic Algorithms*, 77-84. La Jolla, CA: Morgan Kaufmann.