
Editorial Introduction

Kenneth De Jong

Computer Science Department
George Mason University
Fairfax, Virginia 22030 USA
kdejong@gmu.edu

Riccardo Poli

Department of Computer Science
University of Essex, UK
rpoli@essex.ac.uk

Jonathan Rowe

School of Computer Science
University of Birmingham, UK
j.e.rowe@cs.bham.ac.uk

1 Background

The seventh workshop on the foundations of genetic algorithms, FOGA-7, was held in Torremolinos (Malaga), Spain from September 2-4, 2002. This series of workshops was initiated in 1990 to encourage further research in the theoretical aspects of genetic algorithms, and have been held bi-annually ever since. The papers presented at these workshops are revised, edited and published in book form during the year following each workshop. This series of (now seven) books provides the best source of reference for the theoretical work in this field.

At the same time this series of books provides a clear picture of how the workshops have grown and matured along with the field to encompass many evolutionary computation paradigms including evolution strategies (ES), evolutionary programming (EP), and genetic programming (GP), as well as the continuing growth in interactions with other fields such as mathematics, physics, and biology.

A tradition of these workshops is organize them in such a way as to encourage lots of interaction and discussion by restrict the number of papers presented and the number of attendees, and by holding the workshop in a relaxed and informal setting. This year's

workshop was no exception. Forty two researchers met for three days to present and discuss twenty two papers in a delightful setting at a resort hotel in Torremolinos, Spain arranged by Carlos Cotta from the University of Malaga.

2 Overview of the Papers

To date there is no single unifying theory or analysis tool that is sufficient to characterize the behavior of Evolutionary Algorithms (EAs) in any general sense. Rather, a variety of analysis tools and techniques have emerged over the last 30 years that provide insights into particular properties of particular EAs under particular conditions.

This should not be construed as a negative statement about the field of Evolutionary Computation (EC). Rather, it is a reflection of the complexity of the algorithms under study and their application areas. It does, however, make it difficult to understand and make connections between a collection of papers such as the ones in this volume that represent state-of-the-art progress across the entire range of analysis tools and techniques. In this editorial introduction we attempt to do so.

2.1 Schema Theory

A traditional way of analyzing standard GAs is via "schema analysis" in which GAs are viewed as distributing new samples among competing subspaces (schemata) of increasingly finer-grained partitions of the search space. With sufficient knowledge about the fitness landscape, exact models of simple GAs can be constructed. The mathematical tractability of such models depends heavily on the epistatic interactions of the variables of the fitness landscape. Often, by choosing landscapes with simple interactions, the analysis can be made tractable enough to gain insights into GA performance. The paper by Furutani, "Schema Analysis of OneMax Problem - Evolution Equation for First Order Schemata", is an example of this approach. Using the well-studied OneMax Problem, schema evolution equations can be derived and insights into the interplay between crossover and mutation are obtained.

Real GAs, of course, generally do not have sufficient *a priori* knowledge of fitness landscapes to make strong convergence claims. Rather, they depend on sample estimates to guide the search process. It is intuitively clear that the reliability of such estimates depends heavily on the amount of epistasis in the fitness landscape. The paper by Heckendorn, "Partitioning, Epistasis and Uncertainty", begins to formalize this intuition by using Walsh analysis to express the uncertainty introduced by epistasis, and by showing how it affects the ability to accurately estimate the average fitness of competing schemata.

Because of the complex, non-linear interactions among the components of an EA, progress is often made by studying the effects of individual components in isolation. A well-known example of this from the biological literature is Geiringer's theorem that characterizes the steady state distribution of genotypes in a population undergoing ONLY crossover. The original theorem assumed fixed-length genotypes, while many of our EAs today have variable-length genotypes. The paper by Poli, Stephens, Wright, and Rowe, "A Schema-theory-based Extension of Geiringer's Theorem for Linear GP and Variable-length GAs

under Homologous Crossover”, uses a schema theory approach to extend this result to EAs that use variable-length linear genotypes and a particular form of ”homologous” crossover.

2.2 Dynamical System Models

Another approach to understanding EAs is to model them as a dynamical system and analyze their trajectories, fixed points, etc. The paper by Wright, Rowe, Poli, and Stephens, ”Bistability in a Gene Pool GA with Mutation”, is a classic example of this approach. They are able to show that for even simple single-peaked fitness landscapes, certain EAs can have two stable fixed points.

In general, without considerable knowledge about a fitness landscape, it is difficult to analytically derive the fixed points of an EA. Some progress has been made in the case of simple GAs. For example, it is easy to see that homogeneous populations are fixed points for crossover-only GAs. It is more difficult, however, to determine stability and basins of attraction. This is often done empirically by repeatedly executing a model, or by running an actual GA. The paper by Reeves, ”The Crossover Landscape and the Hamming Landscape for Binary Search Spaces”, shows that there is a fairly tight coupling between local optima in Hamming space and the attractors of a crossover-only GA, suggesting that GA attractors might be more efficiently found using a local neighborhood search in Hamming space.

An ever-present issue in EA theory is the use of infinite population models. In many contexts, it is a key assumption that provides the mathematical tractability needed to prove theorems. At the same time it leaves open the question as to whether such analysis has any bearing on the finite population models used by practitioners on a day-to-day basis. On the other hand, finite population models frequently suffer from ”state space explosion” problems as they scale up to realistic population sizes. The paper by Prugel-Bennett, ”Modelling Finite Populations”, examines two approaches to modeling the fitness distributions of finite populations: 1) by explicitly modeling the distribution itself, and 2) by using the cumulant description techniques from statistical mechanics. The latter is shown to be much more precise for realistic population sizes, and the reason for this advantage is the fact that cumulants are ”self-averaging” quantities.

2.3 Stochastic Search

Another direction of analysis is to view EAs as instances of stochastic search algorithms that maintain implicit, dynamically changing probability distributions used to direct future search with the goal, of course, of having significant amounts of probability mass accumulate at the (unknown) global optima. Taking this idea one step further results in algorithms like PBIL and UMDA that replace the population with an explicit probability distribution and replace reproductive sampling with a distribution update (learning) rule. An interesting question is how difficult is it to represent and ”learn” such distributions particularly when there are higher order interactions among the variables.

The paper by Shapiro, ”The sensitivity of PBIL to Its Learning Rate, and How Detailed Balance Can Remove It”, shows how one can improve the robustness of a particular learning regime by focusing on its detailed balance equations. The paper by Muehlenbein and Mahnig, ”Evolutionary Algorithms and the Boltzmann Distribution”, suggests that

significant improvements can be made by assuming the underlying distributions are Boltzmann in character and thus reducing the problem to estimating the necessary Boltzmann distribution parameters.

2.4 Representation

An intriguing observation from biology is that many species use some form of diploid genotype, and much discussion has ensued regarding the advantages of this kind of genetic redundancy. Of continuing interest to the EA community is whether diploid representations increase problem-solving capabilities for certain classes of problems and certain types of EAs. The paper by Liekens, ten Eikelder, and Hilbers, "Modelling and Simulating Diploid Simple Genetic Algorithms", develops a Markov model of simple diploid GAs, and present a somewhat surprising result that diploid GAs can be simulated by haploid GAs, suggesting that for static landscapes there is nothing of significance to be gained by switching to a diploid representation.

Many of our standard EAs make no sharp distinction between the internal genotype space being manipulated and the phenotypic space on which fitness is defined. In many cases the two spaces are the same or have simple maps between them that are one-to-one and onto. Yet there appear to be clear advantages to representations that are redundant, include self-adaptation parameters, etc. The difficulty is formally analyzing and designing such systems. The paper by Marc Toussaint, "On the Evolution of Phenotypic Exploration Distributions", proposes a formalism for doing so in which neutral sets represent alternative exploration strategies, and then characterizes the evolution of the distribution over strategy space in terms of minimizing difference between it and a particular exponential fitness function.

2.5 Behavior Spaces

One of the most fascinating and complex areas of EC is that of evolving behaviors. Typically, one evolves executable objects (finite state machines, rules, Lisp code, etc.) which, when executed, provide behavioral traces from which fitness is calculated. Except for the fact that most executable representations have a many-to-one map onto behavioral spaces, little is known in general about the properties of behavior spaces. The paper by Langdon, "How many Good Programs are there? How Long are they?", provides considerable insight into the case where the executable objects are random, linear, non-recursive programs and the behavioral repertoire is that of boolean functions. In this case Langdon is able to show that, in the limit, as the length of such programs increases, most behaviors correspond to constant functions or extremely simple boolean functions. In addition, bounds are provided on how long programs need to be on average before this limiting distribution is closely approximated.

2.6 Coevolution

Coevolutionary EAs have captured the imagination of many EA researchers and practitioners, but have often resulted in disappointment and frustration. One reason for this is that the dynamics of a coevolutionary system is considerably more complex than a standard EA, and the observed behaviors are often counterintuitive to expectations drawn from

standard EA experiences. One promising analysis tool is evolutionary game theory, but requires significant extensions since in its standard form contains only replicator dynamics and no reproductive variation. The paper by Wiegand, Liles, and De Jong, "Modeling Variation in Cooperative Coevolution Using Evolutionary Game Theory", takes a step in that direction by showing how, for cooperative coevolutionary EAs, the standard Vose-like mixing matrices can be easily incorporated into evolutionary game theory models. This extended model is then used to study the fixed points and basins of attraction on simple fitness landscapes.

A frequent goal of implementers of competitive coevolutionary systems is the emergence of an "arms race" in which the overall performance of the competitors continues to rise. In practice, arms races are difficult to obtain and seldom seen. The paper by Bucci and Pollack, "A Mathematical Framework for the Study of Coevolution", introduces a framework for studying such systems based on the mathematics of ordered sets, and shows how the order structure of the coevolving populations can be a key source of difficulty, and in special cases closely related to notions of Pareto dominance.

One of the difficulties with competitive coevolutionary systems is that fitness is based on internal evaluations involving other members of the current population. Hence, it is possible for evolving individuals to have high internal fitness, but low fitness from an external, objective point of view. The paper by Luke and Wiegand, "Guaranteeing Coevolutionary Objective Measures", focuses on identifying the conditions under which an external objective fitness measure exists such that the dynamics of a standard EA using this measure are identical to a coevolutionary EA with an internal fitness measure.

2.7 Optimization-oriented Analysis

One of the difficulties in comparing the performance of competing optimization algorithms is in selecting unbiased evaluation criteria. One approach is to define test problem generators for a class of problems (such as boolean satisfiability). Algorithms are then compared in terms of their performance on a set of randomly drawn problem instances. The algorithmic difficulty of such a scenario can be further increased by having the problem instances represented as "black boxes", i.e., the only information available to an algorithm about a test problem instance is obtained by sampling, i.e., by injecting input sequences into the black box and observing the output.

The paper by Droste, Jansen, Tinnefeld, and Wegener, "A New Framework for the Valuation of Algorithms for Black-box Optimization", shows how such black-box scenarios can be used as the basis for a formal analysis of optimization algorithms. By using a sample probe (a function evaluation) as the basic unit of computational cost, the complexity of algorithms and problems are characterized in terms of the expected number of evaluations before the optimum is first encountered. The usefulness of this framework is illustrated by establishing lower bounds on the complexity of a variety of black-box optimization problems, and then evaluating the effectiveness of particular optimization algorithms in terms of their ability to approach these lower bounds.

The paper by Borisovsky and Ereemev, "A study on the Performance of the (1+1)-Evolutionary Algorithm", presents an alternative approach to estimating the time (number of samples) required by an EA to find the optimum of a particular problem. The basic idea

is to focus on the dynamically changing probability mass of the upper tail of the fitness distribution and argue that "optimal" algorithms will never reduce this probability mass, i.e., they satisfy a "monotonicity" property. Using this approach the black-box complexity of OneMax can be rederived, and the performance of a (1+1)-EA can be shown to be optimal on a variety of problem classes.

One of the useful properties of EAs when solving optimization problems is their relative insensitivity to noisy objective functions. In such situations most optimization algorithms use some sort of sampling technique in order to obtain accurate estimates of particular function values and/or their derivatives. The paper by Hartley, "The Long-term Behaviour of Genetic Algorithms with Stochastic Evaluation", shows how such sampling techniques are in a sense "built into" certain population-based methods in the sense that by choosing an adequate population size, there is no need for an EA to do more than one evaluation when assessing the fitness of individuals.

Similarly, motivated by noisy industrial optimization applications, the paper by Beyer, Olhofer, and Sendhoff, "On the Behaviour of $(\mu/\mu_I, \lambda)$ -ES Optimizing Functions Disturbed by Generalized Noise", extends the existing simple (but tractable) models of noise to more general models that better reflect the situations encountered in real-world optimization problems such as aerodynamic design. Although these more general noise models are intractable in general, the paper shows how the results from the simpler noise models can be used to provide both qualitative and quantitative results regarding ES performance for the more complex cases.

2.8 EA Design

One of the issues facing an EA practitioner is that of selecting the best (or at least an appropriate) EA to solve a particular problem. In addition to the fact that there are now a wide variety of EAs to choose from, most EAs are themselves parameterized, requiring an additional parameter tuning stage.

One approach to improving this situation is to design EAs that are self-tuning. Historically, the most widely studied mechanisms of this type have been self-tuning mutation operators. While intuitively pleasing, the design of a robust feedback control mechanism for such a self-tuning process is difficult. One approach is the so-called "self-adaptive" approach in which the parameters to be tuned are embedded in the genome and are adapted using the basic evolutionary mechanisms of selection and variation. The paper by Smith, "Parameter Perturbation Mechanisms in Binary Coded GAs with Self-adaptive Mutation", analyzes two such mechanisms, one endogenous and one exogenous, using a dynamical systems approach and characterizes their robustness on several difficult fitness landscapes.

A more complex method for self-adapting mutation is described by Corne, Oates, and Kell in their paper "Fitness Gains and Mutation Patterns: Deriving Mutation Rates by Exploiting Landscape Data". They define several possible criteria for dynamically selecting among competing mutation operators, each of which requires knowledge of the fitness landscape to be effective. They then consider the extent to which such knowledge can be inferred during the evolutionary process, focusing on situations in which there are large populations with parallel evaluation capabilities, and test these ideas experimentally on several landscapes.

An alternative approach is to develop an effective design methodology that would serve as a guide to practitioners, allowing them to match EAs to problems in a more principled way. The paper by Weicker and Weicker, "Toward Qualitative Models of Interactions in Evolutionary Algorithms", presents a strategy for doing so by capturing EA design knowledge in the form of a hypergraph that supports the ability to do qualitative reasoning about the effects of design decisions.

It seems intuitively clear that EA performance on difficult problems could be significantly improved by taking advantage of multi-population models. The difficulty, generally, is in designing an effective way for multiple populations to interact. The paper by Schaefer and Kolodziej, "Genetic Search Reinforced by the Population Hierarchy", presents an approach in which the population interactions are structured hierarchically, with lower levels representing more fine-grained search, and present results showing its effectiveness relative to single-population EAs and traditional multi-population island models.

3 Concluding Remarks

So there you have it - a collection of twenty two papers representing the state-of-the-art research in evolutionary computation theory. We hope you find them as interesting and provocative as we have, and that these papers serve as a catalyst for further progress to be reported at the next FOGA workshop in 2004.