

Theory of Evolutionary Algorithms

13.01. - 18.01.2002

organized by

H.-G. Beyer (Dortmund)

K. De Jong (Fairfax)

C. Reeves (Coventry)

I. Wegener (Dortmund)

The previous Dagstuhl workshop on the “Theory of Evolutionary Algorithms” held in February of 2000 had a great influence on the development of this field and provided a unique opportunity for the people working in this area to interact with each other. Therefore, we had many people who were interested in the new workshop and we could not invite all who asked us.

The idea was to discuss the different approaches to a theory of evolutionary algorithms. The participants were researchers with quite different scientific background. People influenced by computer science, mathematics, physics, biology, or engineering came together which led to vivid and fruitful discussions.

The organizers are happy to report that 40 researchers accepted an invitation to Dagstuhl. They came from Germany (13), USA (9), England (6), Belgium (2), Austria (2), India (1), Japan (1), Mexico (1), Netherlands (1), Romania (1), Russia (1), Spain (1), and Switzerland (1). The 32 talks captured all the aspects of a theory of evolutionary algorithms, among them EA-dynamics, non-static fitness and robustness, algorithmic aspects of EAs, recombination, fitness landscapes, global performance of EAs, and schema approaches. The schedule included an evening session showing “evolution strategies in action”.

Besides the official schedule the participants used unscheduled time for many discussions and some informal sessions with short talks, all inspired by the special Dagstuhl atmosphere.

Seminar Program

Monday, January 14th, 2002

- 9.15 – 9.50 **Jonathan E. Rowe**
Dynamics of Evolutionary Algorithms on Infinite Search Spaces
- 9.50 – 10.25 **Christopher Stevens**
Coarse Grained Formulations of Genetic Dynamics: An Overview
- 10.50 – 11.25 **Alden Wright**
A Fixed Point Analysis of a Gene Pool Genetic Algorithm
- 11.25 – 12.00 **Adam Prügel-Bennett**
When a Genetic Algorithm Outperforms a Hill-Climber
- 15.30 – 16.05 **Dirk V. Arnold**
Local Performance of Evolution Strategies
in the Presence of Noise
- 16.05 – 16.40 **Jürgen Branke**
Possible Selection Bias when Searching for Robust Solutions
- 16.50 – 17.25 **Stefan Droste**
On the Analysis of the (1+1) EA for a Dynamically
Changing ONEMAX
- 17.25 – 18.00 **Karsten Weicker**
On a Local Operator in a Dynamic Environment

Tuesday, January 15th, 2002

- 9.00 – 9.45 **Annie S. Wu**
The Proportional Genetic Algorithm (Keynote Talk)
- 9.45 – 10.20 **Michael Kolonko**
Intensifying Evolutionary Algorithms by Simulated Annealing
- 10.50 – 11.25 **Juan Julio Merelo**
Evolutionary Computation Software
- 11.25 – 12.00 **William B. Langdon**
Structure of the Genetic Programming Search Space
- 15.30 – 16.05 **Ingo Rechenberg**
Analysis of Evolution Strategies on Ridge Functions
- 16.05 – 16.40 **Clarissa Van Hoyweghen**
Modelling a Recombinative Hill-Climber on the HIFF Problem
- 16.50 – 17.25 **Dirk Thierens**
Adaptive Mutation Rate Parameter Control in
Genetic Algorithms
- 17.25 – 18.00 **Ingo Wegener**
Real Royal Road Functions
- 19.30 – 20.30 **Ingo Rechenberg**
Evolution Strategies in Action (Evening Session)

Wednesday, January 16th, 2002

- 9.00 – 9.45 **Peter F. Stadler**
The Structure of Fitness Landscapes (Keynote Talk)
- 9.45 – 10.20 **Colin Reeves**
Attractor Distributions in Fitness Landscapes
- 10.50 – 11.25 **Franz Rothlauf**
Redundant Representations for Genetic and Evolutionary Algorithms
- 11.25 – 12.00 **Marc Toussaint**
Neutrality and Self-Adaptation

Thursday, January 17th, 2002

- 9.00 – 9.35 **Anton Eremeev**
Performance Estimation of Some Mutation-Based Evolutionary Algorithms
- 9.35 – 10.10 **Jeffrey Horn**
Fundamental Limitations of Multispecies Artificial Evolution
- 10.45 – 11.20 **Kalyanmoy Deb**
Test Problems for Multi-objective Optimization
- 11.20 – 11.55 **Lothar Thiele**
Complexity Results for Multiobjective Evolutionary Algorithms
- 15.20 – 16.05 **Michael D. Vose**
Schema Analysis (Keynote Talk)
- 16.05 - 16.40 **Heinz Mühlenbein**
Schemata and Marginal Distributions
- 16.50 – 17.25 **Riccardo Poli**
Exact Schema Theorems and Markov Chain Models for Genetic Programming and Variable Length Genetic Algorithms
- 17.25 – 18.00 **James A. Foster**
Effective Expressed Fitness in Sorting Networks

Friday, January 18th, 2002

- 9.00 – 9.35 **Jonathan L. Shapiro**
Learning Mechanisms in PBIL
- 9.35 – 10.10 **Lothar M. Schmitt**
Theoretical Aspects of Genetic Algorithms
- 10.45 – 11.20 **Thomas Jansen**
Introduction and Application of a New Framework for the Valuation of Algorithms for Black-Box Optimization
- 11.20 – 11.55 **Kenneth A. De Jong**
Understanding CoEvolution

Contents

Dirk V. Arnold: <i>Local Performance of Evolution Strategies in the Presence of Noise</i>	6
Jürgen Branke: <i>Possible Selection Bias when Searching for Robust Solutions</i>	7
Kenneth A. De Jong: <i>Understanding CoEvolution</i>	7
Kalyanmoy Deb: <i>Test Problems for Multi-Objective Optimization</i>	8
Stefan Droste: <i>On the Analysis of the (1+1) EA for a Dynamically Changing ONEMAX</i>	8
Anton Eremeev: <i>Performance Estimation of Some Mutation-Based Evolutionary Algorithms</i>	9
James A. Foster: <i>Effective Expressed Fitness in Sorting Networks</i>	9
Jeffrey Horn: <i>Fundamental Limitations of Multispecies Artificial Evolution</i>	9
Thomas Jansen: <i>Introduction and Application of a New Framework for the Valuation of Algorithms for Black-Box Optimization</i>	11
Michael Kolonko: <i>Intensifying Evolutionary Algorithms by Simulated Annealing</i>	11
William B. Langdon: <i>Structure of the Genetic Programming Search Space</i>	12
Juan Julio Merelo: <i>Evolutionary Computation Software</i>	13
Heinz Mühlenbein: <i>Schemata and Marginal Distributions</i>	13
Riccardo Poli: <i>Exact Schema Theorems and Markov Chain Models for Genetic Programming and Variable Length Genetic Algorithms</i>	14
Adam Prügel-Bennett: <i>When a Genetic Algorithm Outperforms a Hill-Climber</i>	15
Ingo Rechenberg: <i>Analysis of Evolution Strategies on Ridge Functions</i> . .	16
Colin Reeves: <i>Attractor Distributions in Fitness Landscapes</i>	17
Franz Rothlauf: <i>Redundant Representations for Genetic and Evolutionary Algorithms</i>	17
Jonathan E. Rowe: <i>Dynamics of Evolutionary Algorithms on Infinite Search Spaces</i>	18
Lothar M. Schmitt: <i>Theoretical Aspects of Genetic Algorithms</i>	19
Jonathan L. Shapiro: <i>Learning Mechanisms in PBIL</i>	19
Peter F. Stadler: <i>The Structure of Fitness Landscapes</i>	19
Christopher Stephens: <i>Coarse Grained Formulations of Genetic Dynamics: An Overview</i>	20
Lothar Thiele: <i>Complexity Results for Multiobjective Evolutionary Algorithms</i>	21

Dirk Thierens: <i>Adaptive Mutation Rate Parameter Control in Genetic Algorithms</i>	22
Marc Toussaint: <i>Neutrality and Self-adaption</i>	22
Clarissa Van Hoyweghen: <i>Modelling a Recombinative Hill-Climber on the HIFF Problem</i>	23
Michael D. Vose: <i>Schemata Analysis</i>	24
Ingo Wegener: <i>Real Royal Road Functions</i>	24
Karsten Weicker: <i>On a Local Operator in a Dynamic Environment</i>	25
Alden Wright: <i>A Fixed Point Analysis of a Gene Pool Genetic Algorithm</i>	25
Annie S. Wu: <i>The Proportional Genetic Algorithm</i>	25

Local Performance of Evolution Strategies in the Presence of Noise

Dirk V. Arnold

LS Informatik 11, University of Dortmund, Germany

Dirk.Arnold@udo.edu

Noise is a common factor in most real-world optimization problems. Sources of noise include, to name but a few, physical measurement limitations, stochastic simulation models, incomplete sampling of large spaces, and human-computer interaction. Evolutionary algorithms are general, nature-inspired heuristics for numerical search and optimization that are frequently observed to be particularly robust with regard to the effects of noise. However, their performance depends on a multitude of parameters, and in combination with fitness environments they form stochastic dynamical systems that are not easily analyzed and understood. With this talk, we attempt to contribute to the understanding of evolutionary optimization in the presence of noise by investigating the performance of evolution strategies, a type of evolutionary algorithm frequently employed for solving real-valued optimization problems. By considering simple noisy environments, it is possible to analytically obtain results that describe how the performance of the strategies scales with both parameters of the problem and of the strategies considered. Such scaling laws allow for comparisons of different variants of the strategies, for tuning evolution strategies for maximum performance, and they offer insights and an understanding of the behavior of the strategies that goes beyond what can be learned from mere experimentation.

In particular, by considering the $(\mu/\mu, \lambda)$ -ES and comparing its performance on the sphere model with that of evolution strategies that do not employ recombination, the potential of genetic repair for optimization in the presence of noise is investigated. Then, the performance of the cumulative mutation strength adaptation scheme on the noisy sphere is addressed. Finally, the performance of the algorithm is compared with that of other common optimization strategies, confirming the frequently claimed relative robustness of evolutionary optimization in the presence of noise.

Possible Selection Bias when Searching for Robust Solutions

Jürgen Branke

Institute AIFB, University of Karlsruhe, Germany
jbr@aifb.uni-karlsruhe.de

In the presented talk, the problem of finding robust solutions was discussed and compared with the problem of finding solutions when the fitness function is noisy. It was demonstrated that when searching for robust solutions, one has to deal with skewed and uneven noise. The effect of this noise on some common genetic operators was examined on a simple model with two types of individuals with equal expected fitness, but unequal noise distributions. It was shown that common operators in evolutionary computing such as elitism, steady-state reproduction, or linear ranking selection are strongly influenced by the noise distributions and exhibit a strong bias towards one solution or the other.

Understanding CoEvolution

Kenneth A. De Jong (joint work with Paul Wiegand and William Liles)
Department of Computer Science, George Mason University, USA
kdejong@gmu.edu

There continues to be a growing interest in the design and implementation of co-evolutionary algorithms for the purpose of solving difficult computational problems. Experience, however, suggests that coevolutionary algorithms are generally more difficult to design than traditional evolutionary algorithms, primarily because of the increased complexity in the dynamics. In this presentation we explore the usefulness of evolutionary game theory as an analysis tool that can provide a better understanding of coevolutionary systems. In particular, we illustrate how symmetric multi-population evolutionary game theory models can be used to better understand a class of existing cooperative coevolutionary algorithms.

Test Problems for Multi-Objective Optimization

Kalyanmoy Deb
Indian Institute of Technology, Kanpur, India
`deb@iitk.ac.in`

The existing problems used for testing multi-objective optimization are not scalable and tunable to any desired degree. In this talk, systematic construction of test problems is discussed. There are two tasks a multi-objective EA (MOEA) must achieve: (i) Converge to the Pareto-optimal front and (ii) maintain a diverse set of solutions. By using three functionals, two objective test problems are designed to test an MOEA's ability to overcome hurdles of convergence and diversity preservation. Of them, non-convexity, non-uniformity, and discreteness of Pareto-optimal front can be easily achieved in test problems. Extending this idea, three different approaches – multiple function approach, bottom-up approach and constraint surface approach – are proposed and a number of interesting multi-objective optimization test problems are illustrated. Such problems would be helpful in comparing and performing scalability studies of MOEAs.

On the Analysis of the (1+1) EA for a Dynamically Changing ONEMAX

Stefan Droste
LS Informatik 2, University of Dortmund, Germany
`Stefan.Droste@udo.edu`

Although evolutionary algorithms (EAs) are often successfully used for dynamically changing optimization problems, there are only a few theoretical results about EAs for dynamic objective functions. Here, the runtime of the (1+1) EA is theoretically analyzed for a dynamic variant of ONEMAX. The main focus lies on determining the degree of change of the fitness function, where the expected runtime of the (1+1) EA rises from polynomial to super-polynomial. The proofs presented show methods how to rigorously analyse EAs on dynamically changing objective functions, as the error introduced by the simplifications made is bounded asymptotically exact.

Performance Estimation of Some Mutation-Based Evolutionary Algorithms

Anton Eremeev

(joint work with Pavel Borisovsky)

Sobolev Institute of Mathematics, Omsk State University, Russia

eremeev@iitam.omsk.net.ru

In this talk we discuss the upper and lower bounds on probability to generate the solutions of certain quality in the (1+1)-ES, the (1, λ)-ES evolution strategies and in a simplified mutation-selection genetic algorithm. The bounds are obtained in terms of the so-called monotone bounds on transition probability of the mutation operator. Some recent results on comparison of the (1+1)-ES to other mutation-based evolutionary algorithms and to the local search are presented. Also we consider some applications of computational complexity theory to evolutionary algorithms, showing certain limitations on capabilities of these algorithms when they are used to solve NP-hard problems.

Effective Expressed Fitness in Sorting Networks

James A. Foster

Department of Computer Science, University of Idaho, USA

foster@cs.uidaho.edu

Things break.

Most evolved systems – including living things, evolved hardware and software, ecologies, and societies – exist in a dangerous world. Fortunately, they usually work well even with faults in their constituent parts. We hypothesize that this is a general phenomenon, even when there is a strong separation between evolving genotypes and vulnerable phenotypes. We present empirical evidence that sorting networks, evolved with selection on one representation (compare-exchange gates) but with a completely different fault model (digital logic gates) are tolerant to realistic hardware faults (stuck-on and short circuit faults) than either traditionally designed or randomly selected sorting networks. We introduce the concept of *effective expressed fitness (EEF)*, the average quality of individuals in the space of phenotypes with distance determined by the fault model. Finally, we argue that the linkage between effective genotypic fitness and EEF should be considered when analysing real-world evolutionary processes.

Fundamental Limitations of Multispecies Artificial Evolution

Jeffrey Horn

Department of Mathematics and Computer Science, Northern Michigan
University, USA
jhorn@nmu.edu

As soon as the number of interacting species reaches three, complex dynamics can emerge, including perpetual oscillations. For five or more species, we can find apparently chaotic behavior. This can happen even under the simplest multi-species regimes, using only proportionate selection and the natural method of resource sharing. Up to now, we have known that in expectation, multiple species can be maintained at stable equilibria when there is little or no competition for resources between different species. And for two species, we have previously shown that increasing overlap of the two resource niches eventually leads to pure competition and expected takeover by the superior species. However, once we had three or more species with mutually overlapping niches, we found oscillations before reaching stable equilibrium. Now, inspired by recent discoveries by Huisman and Weissing in theoretical ecology models of plankton species, we have altered our resource sharing method slightly, and induced perpetual oscillations (i.e., no stable equilibria) in the expected species populations. For five species mutually competing, the dynamics appear to become chaotic, as Huisman and Weissing found. These discoveries imply that even in the simplest of EAs for speciation, high-order interactions can make equilibria and hence expected behavior difficult to predict or control. However, by discovering and addressing the underlying causes of these dynamics, we have some hope of controlling such processes for practical use. For example, by altering resource sharing from a direct division of resources to an indirect measure of "niche count", as in fitness sharing (Goldberg and Segrest, 1987), we can achieve much clearer and stronger multi-species solutions to layout and packing problems.

Introduction and Application of a New Framework for the Valuation of Algorithms for Black-Box Optimization

Thomas Jansen

(joint work with Stefan Droste, Karsten Tinnefeld, and Ingo Wegener)

Department of Computer Science, George Mason University, USA

tjansen@gmu.edu

Black-box optimization algorithms cannot use the specific parameters of the problem instance, i. e., of the fitness function. Their run time is measured as the number of function evaluations. This implies that the usual algorithmic complexity of a problem cannot be used in the black-box scenario. Therefore, a new framework for the valuation of algorithms for black-box optimization is presented allowing the notion of the black-box complexity of a problem. It can be concluded that randomized search heuristics whose (worst-case) expected optimization time for some problem is close to the black-box complexity of the problem are provably efficient (in the black-box scenario). It is proved that black-box complexity and algorithmic complexity are uncomparable measure. The new approach is applied to some problems based on typical example functions as well as to the class of unimodal functions.

Intensifying Evolutionary Algorithms by Simulated Annealing

Michael Kolonko

Institut für Mathematik, TU Clausthal, Germany

kolonko@math.tu-clausthal.de

The paper presents a new flexible way to control the temperature parameter in the simulated annealing optimization algorithm. The temperature controls the probability to accept uphill moves (in minimization problems). In the usual set-up, the sequence of temperatures is given in advance making the process of solutions an inhomogeneous Markoff chain.

We replace this fixed sequence by a stochastic process. The temperature in the n -th step may depend on the whole history up to that time. If the temperature process converges stochastically, then the solutions will converge against the stationary distribution belonging to the homogeneous Markov chain with the limiting temperature as fixed value. For small limiting values this will concentrate around the optimal solutions. Stochastically converging temperatures that

are suitable for simulated annealing can be constructed with a recursive scheme. Here, temperature is decreased if new solutions are accepted and is raised if a new solution is rejected.

This is an intuitively appealing temperature control that leads to a solution process which is convergent in the above sense. Applications to the job shop scheduling problem show that it also has an excellent empirical performance.

Within a genetic algorithm run, simulated annealing can be applied to a single 'good' individual. This individual will then tend to be ahead of the rest population with respect to fitness, it forms a kind of isolated avantgarde. Via crossover and selection, the population will be able to share the progress of the single individual very rapidly. Iterating this procedure – sending one individual ahead by intensive simulated annealing and then catching it up by genetic operators – speeds up the progress of the whole optimization. Some practical results with this kind of procedure in a multicriterial railway time-table optimization are presented.

Structure of the Genetic Programming Search Space

William B. Langdon

Department of Computer Science, University College, London, UK

W.Langdon@cs.ucl.ac.uk

It is known that the fitness distribution of programs tends to a limit as they get bigger. We use Markov chain convergence theorems to give both upper and lower bounds on program size needed for convergence.

Results are presented for four types of computer models. The bounds on the length of random programs needed to converge depends upon the size of memory N . Bounds are exponential in N , $N \log N$ and smaller, depending on the model.

Evolutionary Computation Software

Juan Julio Merelo

Department of Computer Architecture and Technology, University of Granada,
Spain

`jmerelo@geneura.ugr.es`

Excellent (or at least good enough) libraries in a wide range of different programming languages exist for implementing evolutionary algorithms. Most of them follow the object-oriented programming paradigm, which wraps data structures and methods to manipulate them in a single, easy to use, package called a class. Using these OOP design principles, most packages create classes for individuals, fitness functions, genetic and population level operators, and, finally, algorithm classes, which, once instantiated, can be applied to a population to yield an evolved one.

New and emerging technologies can also be, and sometimes are, used in these EA packages: among them, XML (eXtensible Markup Language) is the one that holds the most promise. XML can be used as a language to express evolutionary algorithms, with which you can build documents that, once parsed, can be used to generate programs that run the algorithms in any language. These documents can also be used as an interchange format for different EC packages, should they exist.

In the future, XML and other technologies could be integrated in visual programming environments, as well as in production environments such as web services and operating systems.

Schemata and Marginal Distributions

Heinz Mühlenbein

Fraunhofer-Institut für Autonome Intelligente Systeme, Sankt Augustin,
Germany

`heinz.muehlenbein@gmd.de`

Holland (1975) introduced the concept of schemata in order to facilitate the understanding of genetic algorithms. We show that schemata are nothing else than definitions of marginal distributions derived from the distribution of the population. Let ξ be a schema. Then Holland claimed that genetic algorithms behave approximately like

$$\frac{dP(\xi, t)}{t} = P(\xi, t)(\hat{\mu}_\xi(t) - \hat{\mu}(t)).$$

$P(\xi, t)$ is the probability of schema ξ in the population at generation t , $\hat{\mu}_\xi(t)$ is the average fitness of the schema (Holland, p. 88, 1975). We show that the above equation is fulfilled if P is derived from the Boltzmann distribution

$$p(x, t) = \frac{e^{t\rho(x)}}{\sum_y e^{t\rho(y)}}.$$

In the main part of the talk I speak about the Boltzmann distribution algorithm. A major theorem describes the factorization of the Boltzmann distribution

$$p(x, t) = p(c_1, t) \prod_{i=2}^m p(x_{b_i} | x_{c_i}, t).$$

I describe the connection of this theory to the theory of graphical models.

Exact Schema Theorems and Markov Chain Models for Genetic Programming and Variable Length Genetic Algorithms

Riccardo Poli

Department of Computer Science, University of Essex, UK

`rpoli@essex.ac.uk`

In my talk I have given an overview of the recent advances in the schema theory for genetic programming (GP) and variable length genetic algorithms (GAs). I have also indicated how GP theory is a generalization of the corresponding theory for GAs operating on fixed length strings. In the talk I have also shown how one can extend the Nix and Vose Markov chain model for GAs to GP and variable length GAs. Finally I have briefly indicated some applications of the schema theory, including extensions of Geiringer's theorem to variable length strings under homologous and subtree crossover.

When a Genetic Algorithm Outperforms a Hill-Climber

Adam Prügel-Bennett

Department of Electronics and Computer Science, University of Southampton,
UK

apb@ecs.soton.ac.uk

In this talk, we introduced a new toy-problem with many local optima. The problem consisted of a binary string of length L . The fitness is a function of the Hamming distance from a global distance. The function consists of a periodic function imposed on a gradient. As a consequence there is a series of local optima separated by a fixed Hamming distance. The difficulty of the model was controlled by a parameter n characterising the Hamming distance between local optima. The characteristic first-passage time to the global optimum was considered for a number of search algorithms. For hill-climbers with independent mutations at each site the mean first-passage time scales as L^n . For a stochastic hill-climber allowed only single mutations the first-passage time scales as L^{n+1} . For a genetic algorithm (GA) using crossover (and for n sufficiently small) the first-passage time was empirically found to scale as $L \log(L)$ —the dependence on n only appearing in the prefactor. The superiority of the GA is down to the crossover operator which concentrates the search to parts of the string where solutions disagree—this is particularly advantageous for this type of problem. For a GA without crossover the first-passage time scaled similarly to that for hill-climbing, although, interestingly, the GA was found to take a small number of function evaluations to find the global optimum.

The talk concluded by a discussion on modelling a GA without crossover acting on this problem. A model using the probability of a member of the population being in a Hamming states was examined. However, this model gave a very poor description of the dynamics. The reason for this could be traced to taking the ensemble average. For probability distribution the ensemble average is not a representative distribution. This contrasts with the statistical mechanics approach which uses cumulants of the probability distribution. The averages of these quantities were good representative of the ensemble and thus give a very good approximation to the ensemble dynamics.

Analysis of Evolution Strategies on Ridge Functions

Ingo Rechenberg

FG Bionik und Evolutionstechnik, TU Berlin, Germany

rechenberg@fb10.tu-berlin.de

Main effort in nonlinear theory of the Evolution Strategy is made to calculate the rate of progress f . This convergence quantity denotes the distance moved uphill divided by the number of generations. At the ends of a scale of possible hill formations stand the circular summit and the parabolic ridge. The circular summit (sphere model in n dimensions) is the simplest isotropic quadratic function, which should describe the situation if you are near to the target. You approach the optimum centrally. However, if the optimum is far away you have to move forward into one direction. For this situation the ridge function should be the suitable model, where the target lies at infinity.

For the case of high dimensionality ($n \rightarrow \infty$) an asymptotic theory will be developed for both hill models. Universal variables are introduced to obtain a manageable solution. For the sphere model the progress rate, originally dependent on six parameters (including quality noise), results to be - in the universal notation - a function of the universal mutation step size only. For the ridge function with a given power m the elegant analytic result found for the sphere model could be repeated. The universal rate of progress turns out to be a function of the universal mutation step size again and - additionally - of the heritability h^2 , commonly used in theoretical population biology to describe environmental noise. The mathematical solutions allow calculating analytically the first derivations. This leads to explicit expressions for the maximum speed of progress and the corresponding mutation step size. The formulas give the possibility to check the effectiveness of algorithmic rules for mutation step-size adaptation in the field of the Evolution Strategies.

Attractor Distributions in Fitness Landscapes

Colin Reeves

School of Mathematical and Information Sciences, Coventry University, UK
c.reeves@coventry.ac.uk

One indicator that an instance of a problem may be difficult to solve is that the fitness landscape possesses many optima (most local optima).

We first review previous work on the estimation of ν – the number of optima (or, more generally, “attractors”) that exist in a fitness landscape – by formulating statistical models for measurable quantities that depend on ν . It is shown that fitting such parametric models can provide good estimates in the case that the landscape is isotropic, but that the estimates may be severely negatively biased otherwise.

We then describe *non-parametric* methods based on statistical principles of resampling – the jackknife, the bootstrap and a “plug-in” estimator. Empirically, we find that these methods provide rather good estimates of ν for some NK-landscapes and for the Low Autocorrelation Binary Sequence problem.

Redundant Representations for Genetic and Evolutionary Algorithms

Franz Rothlauf

Lehrstuhl für Wirtschaftsinformatik, University of Bayreuth, Germany
rothlauf@uni-bayreuth.de

When using representations for genetic algorithms (GAs) every optimization problem can be separated into a genotype-phenotype and a phenotype-fitness mapping. The genotype-phenotype mapping is the used representation and the phenotype-fitness mapping is the problem that should be solved.

A representation is denoted to be redundant if the number of genotypes is higher than the number of phenotypes. In redundant representations, a phenotype is represented on average by more than one genotype. Investigating redundant representations more closely, we recognize that - when neglecting the influence of crosscompetition - redundancy can be addressed as a matter of building block (BB)-supply. Therefore, representations that give more copies to high quality solutions in the initial population result in a higher performance of GEAs, whereas encodings where high quality solutions are underrepresented make a problem more difficult to solve. Uniform redundancy, however, has no influence on the performance of GEAs. A representation is denoted to be uniformly redundant if

every phenotype is represented on average by the same number of genotypes. Based on the Gambler’s ruin population sizing model we present a theoretical model describing the effects of redundant representations on the performance of GEAs. With r denoting the number of genotypic building blocks (BBs) that represent the best phenotypic BB, and k_r denoting the order of redundancy (number of genotypic bits that represent one phenotypic bit), the population size grows with $O(2^{k_r}/r)$ and the time to convergence goes with $O(const - r/2^{k_r k - 1})$. k denotes the size of the BBs in the phenotype. A comparison of the empirical results with the theoretical predictions for the one-max problem and the deceptive trap problem shows that the presented model accurately predicts GA performance.

Dynamics of Evolutionary Algorithms on Infinite Search Spaces

Jonathan E. Rowe

Department of Computer Science, University of Birmingham, UK

j.e.rowe@cs.bham.ac.uk

We assume a countably infinite search space. Probability distributions over this space are then sequences in the vector space of “absolutely convergent sequences”. The evolution equation for the simple genetic algorithm can be characterised in this space for proportional selection and mutation. The simplex is not compact in this vector space and so the Brouwer fixed-point theorem does not apply. It is therefore possible to construct a mutation operator and fitness function so that the iterates of the infinite population equation do not converge. This gives one potential source of “bloat” in which the average size of programs in a population increases without bound. A second source comes from the observation that “average length” is not a continuous function. It is thus possible that the population does converge, and yet the average length diverges. A method for constructing such a system is given, based on a generalisation of a theorem from the finite case, which guarantees that (under certain assumptions) any population can arise as the fixed-point of a GA.

Theoretical Aspects of Genetic Algorithms

Lothar M. Schmitt

Department of Computer Hardware, University of Aizu, Japan

lothar@u-aizu.ac.jp

We give an overview on results (and extensions) in *Theoretical Computer Science 259 (2001)*, 1 – 61. The model considered is the tensor-string model with finite-size populations. We discuss transition of our model to the multi-set model via “permutation-averaging”. The analysis is based upon inhomogenous Markov chains. Several mutation operators (single/multiple spot) are considered. Based on separate analysis of mutation M and crossover C we discuss spectral analysis of $M \cdot C = C \cdot M$ and applications. We discuss the mutation-flow equation

$$\|(1 - P_u)Mv\|_1 \leq (1 - \beta) + \beta\|(1 - P_u)v\|_1,$$

$\|v\|_1 = 1$, which together with a shrinking property of fitness selection implies convergence to uniform populations for small-mutation rate. Non-convergence is shown for the simple GA and other scaled GA. Convergence (to optima) is shown for scaled mutation and power-law scaled proportional fitness selection with logarithmic increase.

Learning Mechanisms in PBIL

Jonathan L. Shapiro

Department of Computer Science, University of Manchester, UK

jls@cs.man.ac.uk

Distribution estimation algorithms (DEA) are like evolutionary algorithms, except instead of evolving a population, one evolves the parameters of a parametrized distribution which generates a population. PBIL is the simplest example of a DEA, in which all sites are treated independently. Experimental comparisons between PBIL and a more complex DEA “Mimic” show that the selection mechanism has a larger effect on performance than the model complexity (Johnson and Shapiro, 2001). Here we show that this can be understood in terms of a fundamental sensitivity of PBIL to the learning rate. The algorithm converges to the optimum only if the learning rate is smaller than a problem-dependent function of the stringlength. This can be rectified by using rejection sampling to insure that when the fitness landscape is flat, the invariant fixed-point distribution is uniform.

The Structure of Fitness Landscapes

Peter F. Stadler

Institute for Theoretical Chemistry and Structural Biology, University of
Vienna, Austria
`studla@tbi.univie.ac.at`

Fitness landscapes have proven to be a valuable concept in evolutionary biology, combinatorial optimization, and the physics of disordered systems. A fitness landscape is a mapping from a configuration space into the real numbers. The underlying set may be finite, countably infinite, or a continuum, depending on the application. This set is equipped with some notion of adjacency, nearness, distance, or accessibility.

Landscape theory has emerged as an attempt to devise suitable mathematical structures for describing the “static” properties of landscapes as well as their influence on the dynamics of adaptation. Structural properties of landscapes may be discussed both in algebraic terms, e.g., using correlation measures, and in topological or geometric terms, focussing e.g. on critical points.

Coarse Grained Formulations of Genetic Dynamics: An Overview

Christopher Stephens

Institute for Nuclear Sciences, Universidad Nacional Autónoma de México,
Mexico
`stephens@nucleu.unam.mx`

The dynamics of “genetic” systems, such as genetic algorithms, evolving under the action of genetic operators, such as selection, mutation and recombination are exceedingly complicated. Recombination in particular, being bilinear in the string probabilities, poses special challenges, implying that the dynamical equations governing the evolution of binary strings consist of 2^N coupled non-linear difference equations. In this talk I review representations of the dynamics based on coarse-grained degrees of freedom – building-block schemata. Creation events of a particular string C_i for a given crossover mask m can be written in the form $P'(C_i(m), t)P'(C_i(\bar{m}), t)$ where $P'(C_i(m), t)$ represents the probability to select the building-block schema $C_i(m)$, inherited from the first parent, and $P'(C_i(\bar{m}), t)$ the building-block from the second parent. The equations have a natural recursive structure wherein building-block schemata of a given degree of fine-grainedness (schema order) are related to more coarse-grained building-block schemata of

lower order which in turn are related to even more coarse-grained blocks, etc., the ultimate building-blocks being 1-schemata. I show how this enhances the understanding of the Schema Theorem and Building-block hypothesis within the context of an exact formulation of the dynamics. I also show that the dynamical equations are form-invariant under a coarse-graining and that the set of coarse grainings have a semi-group structure. Finally, I give a very brief overview of concrete results that have been obtained directly from the formulation.

Complexity Results for Multiobjective Evolutionary Algorithms

Lothar Thiele

Department of Electrical Engineering, ETH Zürich, Switzerland

thiele@tik.ee.ethz.ch

Two particular results are presented in the talk:

1. An on-line algorithm for archiving pareto-optimal points with finite memory is derived. It is based on the concept of epsilon-dominance and guarantees convergence and diversity. Let $F = (f_1, f_2, f_3, \dots)$ be an infinite sequence of objective vectors, one by one passed to the ε -update algorithm, and F_t the union of the first t objective vectors of F . Then for any $t > 0$, the following holds: the archive A at time t contains an epsilon-Pareto set of F_t and has finite size which polynomially depends to the precision ε .
2. For the first time, a running time analysis of a multi-objective evolutionary algorithm for a discrete optimization problem is given. To this end, a simple pseudo-Boolean problem (Lotz: leading ones - trailing zeroes) is defined and a population-based optimization algorithm (FEMO). We show, that the algorithm performs a black box optimization in $O(n^2 \log n)$ function evaluations where n is the number of binary decision variables.

Adaptive Mutation Rate Parameter Control in Genetic Algorithms

Dirk Thierens

Institute of Information and Computing Sciences, University of Utrecht,
Netherlands

`dirk@cs.uu.nl`

Adaptation of control parameter values allows to use feedback information from the ongoing search process to optimize performance. In addition it frees the user from making non-trivial decisions about optimal fixed control parameter values. Unfortunately many existing schemes are rather complicated and therefore often not used by novice practitioners. Here we describe an extremely simple adaptive mutation rate parameter control scheme, and show its feasibility in comparison with more complicated self-adaptive schemes and deterministically decreasing dynamic schedules.

Neutrality and Self-adaption

Marc Toussaint

Institute for Neuroinformatics, University of Bochum, Germany

`mt@neuroinformatik.ruhr-uni-bochum.de`

How can an evolutionary exploration distribution adapt in such a way as to exploit the information gained about the problem at hand? In particular in the case of indirect encoding, how can the exploration distribution vary and what about “adaptive genotype-phenotype mappings”? I argue that a neutral set spans a space of different exploration distributions in the phenotype space – different genotypes may induce different topologies “around” the same phenotype (rather than being redundant). Enabling transitions between genotypes in a neutral set (by appropriate mutation operators) enables self-adaptation of the exploration distribution.

Modelling a Recombinative Hill-Climber on the HIFF Problem

Clarissa Van Hoyweghen
(joint work with Bart Naudts)

Department of Mathematics and Computer Science, University of Antwerp,
Belgium

hoyweghe@ruca.ua.ac.be

The Hierarchically If-and-only-If (HIFF) problem was designed by (Watson, 1998) to be ideal for crossover and impossible for mutation. The problem is spin-flip symmetrical and has the property that at each of the $\log \ell$ levels (with ℓ the problem size), local improvements contradict global improvements with probability 0.5. A standard genetic algorithm running on this problem quickly gets stuck in a population that is not diverse enough for crossover to be able to combine blocks into larger ones. A recombinative hill-climber, on the other hand, is ideally equipped to optimize the problem.

In this seminar we presented a model and exact dynamics of a recombinative hill-climber equipped with one-point crossover. This model quantitatively explains to which extent the algorithm advances at different levels simultaneously. We also indicated that the probability model used for one-point crossover is insufficient for two-point crossover, and that strong approximations might be necessary to make a model with a polynomial number of different states.

We finally presented a efficient simulation algorithm that mimicks the dynamics of the hill-climber for both types of crossover operators. It is useful to build an intuition about the action of the hill-climber on the problem.

Schemata Analysis

Michael D. Vose

Department of Computer Science, University of Tennessee, USA

vose@cs.utk.edu

When modeling or coarse-graining stochastic search (via schemata for example) the question of accuracy arises. Relevant to this question is the concept of compatibility of the transition function of the system with respect to an equivalence relation; this concerns the agreement between the trajectory of the system and the trajectory of the model. Compatibility is introduced by way of a simple example, and subsequently defined within the context of Random Heuristic Search. Basic results concerning compatibility are indicated.

Real Royal Road Functions

Ingo Wegener

(joint work with Thomas Jansen)

LS Informatik 2, University of Dortmund, Germany

wegener@ls2.cs.uni-dortmund.de

Royal road functions have been defined by Forrest, Holland, and Mitchell to prove that crossover can decrease the expected optimization time significantly. However, the $(1 + 1)$ EA outperforms each GA on the royal road functions. We present for the first time functions f_n and g_n resp. where it can be proved that one-point crossover and uniform crossover resp. can decrease the expected optimization time from exponential to polynomial. Therefore, f_n and g_n are *real* royal road functions.

On a Local Operator in a Dynamic Environment

Karsten Weicker

Fakultät Informatik, University of Stuttgart, Germany

`weickerk@informatik.uni-stuttgart.de`

In the presented work, the limits of a local standard operator in time-dependent environments are examined empirically.

First, a simple 2-dimensional moving circle problem is considered where the circle moves into a constant direction. There, a self-adaptive, isotropic mutation using the Gaussian pdf is not able to track the circle with a stable accuracy when increasing the severity of the dynamics. This leads to a considerable fraction of runs not being able to track the optimum. A better self-adaptive prediction of the direction of the dynamics suffices to reach a stable accuracy.

Second, the tracking task is combined with an optimization task which is realized in a moving corridor where the direction of the dynamics and the direction of the fitness improvement are orthogonal to each other. The self-adaptation mechanism of evolution strategies is misled by the optimization task which results in an increase of the strategy variables and a loss of the ability to track the corridor. A two-fold selection strategy is introduced where the strategy variables are selected on the basis of a more conservative mechanism. This suffices to dampen the effects significantly.

Both examples demonstrate that new difficulties emerge solely from the introduction of dynamics to the problem. Standard techniques need to be re-evaluated for their suitability to the new demands.

A Fixed Point Analysis of a Gene Pool Genetic Algorithm

Alden Wright

Department of Computer Science, University of Montana, USA

`wright@cs.umt.edu`

Gene pool recombination in a genetic algorithm takes the population to linkage equilibrium in one step. Such a population can be described (in expectation) by l parameters where l is the string length. This allows for a considerable simplification of the infinite population dynamical systems model. This model was used to show that there can be bistability (multiple stable fixed points) in the model when applied to the needle-in-the-haystack fitness function.

The Proportional Genetic Algorithm

Annie S. Wu

School of Electrical Engineering and Computer Science, University of Central
Florida, USA

`aswu@cs.ucf.edu`

I introduce a genetic algorithm (GA) with a new representation method which we call the proportional GA (PGA). The PGA is a multi-character GA that relies on the existence or non-existence of genes to determine the information that is expressed. The information represented by a PGA individual depends only on the types and proportion of characters on the individual, and not on the order of the characters. As a result, the order of the encoded information is free to evolve in response to factors other than the value of the solution, for example, in response to the identification and formation of building blocks. The PGA is also able to dynamically evolve the resolution of encoded information. In this talk, I will describe our motivations for developing this representation and discuss its benefits and drawbacks. Initial comparisons of the PGA with a canonical GA gave promising results. I will present these results and discuss directions of future work based on these preliminary studies.