

10361 Abstracts Collection and Executive Summary

Theory of Evolutionary Algorithms

— Dagstuhl Seminar —

Anne Auger¹, Jonathan L. Shapiro², L. Darrell Whitley³ and Carsten Witt⁴

¹ INRIA Saclay-Ile-de-France, Orsay, FR

`anne.auger@inria.fr`

² University of Manchester, GB

`jls@cs.man.ac.uk`

³ Colorado State University, US

`whitley@cs.colostate.edu`

⁴ Technical University of Denmark, DK

`cawi@imm.dtu.dk`

Abstract. From September 5 to 10, the Dagstuhl Seminar 10361 “Theory of Evolutionary Algorithms ” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general.

Keywords. Evolutionary algorithms, bio-inspired search heuristics, theoretical analysis, optimization time

1 Motivation and Goals

Evolutionary algorithms (EAs) are stochastic optimization methods that are based on principles derived from natural evolution. Mutation, recombination, and selection are iterated with the goal of driving a population of candidate solutions toward better and better regions of the search space. From a more general perspective, EAs are just one instance of bio-inspired search heuristics. Since the underlying ideas of bio-inspired search are easy to grasp and easy to apply, EAs and different bio-inspired search heuristics are widely used in many practical disciplines, mainly in computer science and engineering.

It is a central goal of theoretical investigations of search heuristics to assist practitioners with the tasks of selecting and designing good strategy variants and operators. Due to the rapid pace at which new strategy variants and operators are being proposed, theoretical foundations of EAs and other bio-inspired search heuristics still lag behind practice. However, EA theory has gained much momentum over the last few years and has made numerous valuable contributions to the field of evolutionary computation.

The theory of EAs today consists of a wide range of different approaches. Run-time analysis, schema theory, analyses of the dynamics of EAs, and systematic empirical analysis all consider different aspects of EA behavior. Moreover, they employ different methods and tools for attaining their goals, such as Markov chains, infinite population models, or ideas based on statistical mechanics or population dynamics. In the most recent seminar, more recent types of bio-inspired search heuristics were discussed. Results regarding the runtime have been generalized from EAs to ACO and PSO. Although the latter heuristics follow a different design principle than EAs, the theoretical analyses reveal surprising similarities in terms of the underlying stochastic process. New analytic approaches were also surveyed.

Theoretical studies of EAs in continuous domain have recently evoked interest of people working in the field of classical numerical optimization. Although stochastic and deterministic optimization algorithms address optimization of different types of problems—mainly convex and smooth for deterministic algorithms and noisy, multimodal, irregular for stochastic algorithms—the focuses of both fields became closer and closer: on the one hand many hybridizations of stochastic search and gradient-based algorithms have been proposed, on the other hand, derivative-free optimization is now a major part of the research in the classical optimization community.

The goals of the 2010 Dagstuhl seminar were twofold. The first goal was to discuss the potential and limitations of a unified theory of all types of bio-inspired search heuristics with focus on how to analyze the runtime of estimation-of-distribution algorithms, which themselves can be considered as a generalized model of ACO, PSO and EAs. The second goal of the seminar was to bridge the gap with the classical optimization community.

2 Participants and Results

The seminar brought together 45 participants from 12 countries, and from across the spectrum of EA theory. Talks have been arranged into seven sessions grouped around common themes such as, runtime analysis, multi-objective optimization, landscape analysis, black-box complexity, continuous optimization. The length of the different slots were between 30-45 minutes.

In addition, shorter talks called *spotlight talks* were introduced. These gave PhD students the opportunity to present in 5 minutes their main research interest together with questions they are tackling. This allowed PhD students who did not feel ready to give full talks the opportunity to introduce themselves and feel part of the seminar. This addition was very well received.

Thomas Jansen (University College Cork)

Run Time Analysis of EAs: Limitations of Counting Function Evaluations

Anton Eremeev (Sobolev Institute of Mathematics - Omsk)

Evolutionary Algorithms and Dynamic Programming

Leslie Valiant (Harvard University)

The Scope and Limits of Darwinian Evolution

Tobias Friedrich (MPI Saarbruecken)
Approximation Quality of the Hypervolume Indicator
Dimo Brockhoff (INRIA Saclay, France)
Recent Results on Optimal μ -Distributions for the Hypervolume Indicator
Spotlight: Madeleine Theile (TU Berlin)
How crossover speeds up EAs for MO shortest path problems
Guenter Rudolph (TU Dortmund)
Convergence Rates of MO-EAs
Walter J. Gutjahr (Universitaet Wien)
Runtime Analysis of an EA for Stochastic MO Combinatorial Optimization

Benjamin Doerr (MPI, Saarbruecken)
Drift, Drift, Drift
Per Kristian Lehre (Technical University of Denmark)
A new Drift Theorem for Runtime Analysis of Population-based EAs
Spotlight: Timo Koetzing (MPI, Saarbruecken)
Computing Runtime Bounds of ACO, EA
Dirk Sudholt (University of Birmingham)
Runtime and Speedups in Parallel Evolutionary Algorithms
Verena Heidrich-Meisner (Ruhr-Universität Bochum)
Racing Algorithms for Selection in Noisy Environments

Christian Mueller (ETH Zurich)
Exploring the common concepts of adaptive MCMC and Covariance Matrix Adaptation schemes
Alexander Melkozerov (Tomsk State University)
Implicit Filtering Algorithm Analysis
Marcus Gallagher (The University of Queensland)
Nonparametric Models and EDAs
Steffen Finck (Fachhochschule Vorarlberg)
Theoretical Strategy Comparison

Andrew M. Sutton (Colorado State University)
Local Statistics of Bounded Pseudo-Boolean Functions
Francisco Chicano (Universidad de Malaga)
Elementary Landscape Decomposition of Combinatorial Optimization Problems
Bill Langdon (University College London)
An analysis of GP fitness landscapes
Adam Pruegel-Bennett (University of Southampton)
MAXSAT: The Anatomy of the Landscape of a Hard Optimisation Problem

Jonathan E. Rowe (University of Birmingham)
Convergence of preference functions
Alden Wright (University of Montana - Missoula)

Modeling Drift in a Steady State GA

Gabriela Ochoa (University of Nottingham)

Mapping the topology of combinatorial landscapes

Sebastien Verel (INRIA, Lille)

Towards a Definition of Population based Fitness Landscapes

Christine Zarges (TU Dortmund, Germany)

Artificial Immune Systems as Randomized Search Heuristics

Carola Winzen (MPI, Saarbruecken)

Comparison-based black-box complexities

Olivier Teytaud (INRIA, Saclay, France)

Decidability and complexity in partially observable antagonist coevolution

Frank Neumann: (MPI, Saarbruecken)

Computational Complexity Analysis of Simple Genetic Programming On Two

Modeling Isolated Program Semantics

Another innovation in the 2010 seminar was *working parties*, whose goal was to bring small groups of researchers together to discuss new important results and map out future research directions. Four working parties took place:

Working parties:

- Natural Gradient
 - Fitness Landscapes
 - Future of Runtime Analysis
 - Genetic Programming on General Purpose Graphics Processing Units
-

Many of the presentations and discussions were dedicated to identifying the limitations of the various approaches, shedding light on their complementarity and arriving at a wider consent with regard to advantages and disadvantages of different techniques. In particular limitations and the complementarity of different approaches for discrete runtime analysis were discussed. This field appears to be reaching a certain maturity. While a great deal of progress has been made on mutation-based search that moves from one point to another, convergence results on more general population based algorithms have been more difficult to obtain. New generalized forms of drift analysis have been developed which now makes it easier to apply this theory to analyze the convergence behavior of randomized search algorithms on a wider range of problems. The method works by aggregating state information in complex population-based evolutionary algorithms. Talks were of high quality and showed remarkable progress in different areas. New results on fitness landscapes analysis were presented which show that certain metrics of search landscapes that usually require exponential time to compute (such as exact statistical moments, and measures such as auto-correlation) can be computed in polynomial time for many NP-Hard prob-

lems such as MAX-SAT, QAP and Graph-Coloring. The approach is based on a combination of Walsh Analysis and Elementary Landscapes.

Several new themes that have been largely absent from previous seminars have emerged as prominent research themes from the current seminar. The first theme is the theory for multi-objective optimization algorithms: first results on the linear convergence of a specific MO algorithm, overview of recent results on hypervolume based search algorithms have been presented. The second theme concerns black-box complexity, which provides bounds on the runtime of very general classes of randomized search heuristics. The last theme is the theory of Natural Evolutionary Algorithms: a recent major result presenting the state-of-the-art continuous optimizer CMA-ES as a natural gradient algorithms where a gradient descent is performed on the set of parameter of the search distribution has been discussed during Dagstuhl. This result is making a link between CMA and classical optimization algorithms but more generally the idea offers a very principled design technique for search algorithms that sample from a parametrized distribution.

3 Conclusion

Fruitful and stimulating discussions among varying groups of participants occurred throughout the week of the Dagstuhl seminar on “Theory of Evolutionary Algorithms”. Besides the presentations, the unconventional format of the working parties was successful in provoking discussions and stimulating the exchange of new ideas. The spotlight talks provided a great opportunity for new PhD students to introduce themselves to the community. The Dagstuhl seminars are firmly established in the community as biannual event, and we hope to be able to build on this success and continue to promote discussions between researchers in different areas of EA theory at further workshops in the future.

Anne Auger, Jonathan Shapiro, Darrell Whitley, Carsten Witt

4 Collection of Abstracts

Recent Results on Optimal μ -Distributions for the Hypervolume Indicator

Dimo Brockhoff (INRIA Saclay - Orsay, FR)

Hypervolume-based multiobjective evolutionary algorithms are the most recent type of stochastic algorithms for tackling multiobjective optimization problems—casting the multiobjective problem into a single objective (set-) problem: find the set of μ points among all solution sets of size μ that maximize the underlying hypervolume indicator. Such an optimal solution set is also known under the term optimal μ -distribution. From both the algorithm design and analysis point-of-view, the knowledge about optimal μ -distributions is necessary in terms of performance assessment: not only do we want to know where the optima of certain test functions lie in the search space but also how a hypervolume-based algorithm approaches them in order to finally establish algorithms that quickly and reliably converge to a set with the maximally achievable hypervolume value.

In this talk, I presented some of our most recent results on theoretically investigating optimal μ -distributions for the hypervolume indicator which have not presented elsewhere before. In particular, I presented new results on the existence of optimal μ -distributions and the exhaustive investigation of problems with linear bi-objective fronts.

Keywords: Evolutionary multiobjective optimization, hypervolume indicator, optimal μ -distributions

Elementary Landscape Decomposition of Combinatorial Optimization Problems

Francisco Chicano (Universidad de Málaga, ES)

We revise some important facts of the landscape theory that allow us to better understand the behaviour of local search methods in combinatorial optimization problems. In particular, we characterize what an elementary landscape is and show how a given landscape can be decomposed as the sum of several elementary landscapes. The decomposition of any landscape into elementary ones can be used as a base for the design of new search strategies or operators. Furthermore, this decomposition opens the door to the exact computation of the autocorrelation coefficient, an exact formal metric that models the performance of local search methods on the problem.

Keywords: Fitness landscapes, landscapes' theory, combinatorial optimization

Drift, Drift, Drift

Benjamin Doerr (MPI für Informatik - Saarbrücken, DE)

Drift analysis, brought into our field by He and Yao, quickly became one of the most powerful tools in the analysis of randomized search heuristics. The last 12 months saw a number of interesting new ideas concerning drift analysis. I'll report them and discuss some open problems arising from them. To be precise, all I'll talk about concerns drift analysis used to prove upper bounds on run-times.

Keywords: Drift analysis

Evolutionary Algorithms and Dynamic Programming

Anton Eremeev (Sobolev Institute of Mathematics - Omsk, RU)

Recently, it has been proven that evolutionary algorithms produce good results for a wide range of combinatorial optimization problems. Some of the considered problems are tackled by evolutionary algorithms that use a representation, which enables them to construct solutions in a dynamic programming fashion. We take a general approach and relate the construction of such algorithms to the development of algorithms using dynamic programming techniques. Thereby, we give general guidelines on how to develop evolutionary algorithms that have the additional ability of carrying out dynamic programming steps. Finally, we show that for a wide class of the so-called DP-benevolent problems (which are known to admit FPTAS) there exists a fully polynomial-time randomized approximation scheme (FPRAS) based on an evolutionary algorithm.

Keywords: Combinatorial optimization, dynamic programming, evolutionary algorithms, FPTAS

Joint work of: Doerr, Benjamin; Eremeev, Anton; Horoba, Christian; Neumann, Frank; Theile, Madeleine

Theoretical Strategy Comparison

Steffen Finck (Fachhochschule Vorarlberg, AT)

Comparing different optimization strategies on a theoretical level yields a deeper understanding of the strategies. However, available theoretical results are often derived under specific assumptions and the form of the results depends on the analysis method applied. This makes a direct comparison of different strategies cumbersome or even impossible. By using the same theoretical analysis approach for the strategies of interest one can avoid these obstacles. In this work an analysis method based on dynamic systems theory is reviewed and used to

analyze Simultaneous Perturbation Stochastic Approximation (SPSA) on the sphere model. This approach was chosen because it allows to determine the dynamic behavior and the influence of the strategy specific parameters. The latter can be used to derive guidelines for setting these parameters.

The results obtained are later compared with results for various Evolution Strategies.

While for the noise-free sphere model SPSA exhibits a similar behavior than Evolutionary Gradient Search, the results differ for the sphere with constant normalized noise. There the efficiency of SPSA varies w.r.t. certain parameters. In the case of constant noise, SPSA can achieve small residual location errors, however, increasing the noise level yields a quadratic increase in the residual location error. For the ES variants considered the increase is linear.

Approximation Quality of the Hypervolume Indicator

Tobias Friedrich (MPI für Informatik - Saarbrücken, DE)

In order to allow a comparison of (otherwise incomparable) sets, many evolutionary multiobjective optimizers use indicator functions to guide the search and to evaluate the performance of search algorithms. The most widely used indicator is the hypervolume indicator. It measures the volume of the dominated portion of the objective space. Though the hypervolume indicator is very popular, it has so far not been shown that maximizing the hypervolume indicator is indeed equivalent to the overall objective of finding a good approximation of the Pareto front. In this talk, I will present several recent theoretical results on the additive and multiplicative approximation factor achieved by two-dimensional sets maximizing the hypervolume indicator.

Keywords: Multiobjective optimization, hypervolume indicator, approximation

Nonparametric Models and EDAs

Marcus Gallagher (The University of Queensland, AU)

Nonparametric models have seen only limited application in EDAs. This talk will discuss possibilities for employing kernel density estimators in EDAs. There are several open issues regarding: selection of the kernel bandwidth, adaptive estimators, using weighted estimators and relationships with other EAs and theoretical frameworks.

Runtime Analysis of an EA for Stochastic Multi-Objective Combinatorial Optimization

Walter J. Gutjahr (Universität Wien, AT)

For stochastic multi-objective combinatorial optimization (SMOCO) problems, the Adaptive Pareto Sampling (APS) framework has been proposed. APS is based on sampling and on solution of deterministic multi-objective sub-problems. We show that when plugging in the well-known SEMO algorithm as a sub-procedure into APS, a modification has to be performed to achieve fast convergence to the Pareto front. A general theorem is presented indicating how runtime complexity results for APS can be derived from corresponding results for SEMO. This may be a starting point for the runtime analysis of evolutionary SMOCO algorithms.

Racing Algorithms for Selection in Noisy Environments

Verena Heidrich-Meisner (Ruhr-Universität Bochum, DE)

Selection is a fundamental step in evolutionary search and optimization algorithms (EAs). In most current EAs, this step relies on ranking candidate solutions according to their quality. Because of noise and uncertainty, in many applications the quality of a search point is a random variable. We propose a statistically sound method for robust selection that allows for efficient search and optimization in the presence of such stochastic effects. We adapt racing algorithms based on Hoeffding and empirical Bernstein concentration inequalities for this task. Only assuming that the observed quality values are almost surely within known bounds, we derive confidence intervals for the quality of solutions and for the correctness of a selection decision based on repeated evaluations.

The proposed racing selection algorithm automatically adjusts the total number of evaluations per iteration of the search algorithm and, more importantly, how these evaluations are distributed among the candidate solutions. The estimation of the individual quality values is kept just accurate enough for a sufficiently good ranking, which in turn suffices for the EA to find better solutions. This saves objective function evaluations and addresses the trade-off between speed and accuracy.

Keywords: Racing algorithm, selection, uncertainty handling, covariance matrix adaptation evolution strategy

Joint work of: Heidrich-Meisner, Verena; Igel, Christian

Run Time Analysis of Evolutionary Algorithms: Limitations of Counting Function Evaluations

Thomas Jansen (University College Cork, IE)

(joint work with Christine Zarges (TU Dortmund))

When analyzing the optimization time of evolutionary algorithms most often the number of function evaluations is used to measure the actual computation time.

Moreover, results are most often stated in asymptotic form concentrating only on the most dominant terms in the optimization time ignoring finer details. However, there are studies where more detailed and precise results about the optimization time are presented. Considering one case study it is demonstrated that such results can be completely misleading with respect to the actual computation time. On this level of detail the number of function evaluations is no longer an appropriate measure and implementation details need to be taken into account.

Inspired by the field of algorithm engineering it is shown how theoretical and experimental analysis can be combined so that results that are closer to the actual performance of evolutionary algorithms are obtained.

Computing Runtime Bounds of ACO/EA

Timo Koetzing (MPI für Informatik - Saarbrücken, DE)

In analyzing evolutionary algorithms or ant colony optimization algorithms, one frequently needs to lower bound the chance of getting a significant improvement over the best so far. As many of the the analyzed settings, a new sample is in expectation near the best so far. Hence, we want to get lower bounds on the tails of certain distributions, but unfortunately, few of such tailbounds are known.

Keywords: Lower tail bounds, ACO, Evolutionary Algorithms, Theory

2-bit Flip Mutation Elementary Fitness Landscapes

William Langdon (University College London, GB)

Genetic Programming parity with only XOR is not elementary. GP parity can be represented as the sum of $k/2 + 1$ elementary landscapes. Statistics, including fitness distance correlation (FDC), of Parity's fitness landscape are calculated. Using Walsh analysis the eigen values and eigenvectors of the Laplacian of the two bit flip fitness landscape are given.

Tests support λ/μ as a measure of the ruggedness of elementary landscapes for predicting problem difficulty. An elementary needle in a haystack (NIH) landscape is given.

Keywords: Genetic Algorithms, Genetic Programming, search, optimisation, graph theory, Laplacian, Hamming cube

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2814>

A new Drift Theorem for Runtime Analysis of Population-based EAs

Per Kristian Lehre (Technical University of Denmark, DK)

Drift analysis is one of the primary mathematical techniques for runtime analysis of evolutionary algorithms (EAs). This technique has mostly been applied to analyse single-individual based EAs, e.g. the (1+1) EA. One barrier in applying classical drift analysis to population-based EAs is the state aggregation problem: How to represent the state of the population by a single real value.

We introduce a new drift theorem aimed at analysing population-based EAs. The effects of the selection mechanism and the variation operator are decoupled in the conditions of the theorem. The drift conditions are with respect to a random walk of a single individual, and not the population as a whole, thus alleviating the state aggregation problem.

Keywords: Runtime analysis

Full Paper:

<http://www2.imm.dtu.dk/~pkle/papers/selpres2009.pdf>

Implicit Filtering Algorithm Analysis

Alexander Melkozerov (Tomsk State University, RU)

In the talk, we consider the analysis of the implicit filtering (IF) algorithm, a direct search and optimization method which uses gradient approximation. The aim of the work is to apply methods from the theory of evolution strategies to the IF analysis and thus show that these methods can be used to analyze not only ES, but other optimization algorithms as well.

The first step of this work was the analysis of a simplified IF algorithm on the noisy sphere model. A theoretical stationary state fitness function formulae have been obtained for IF with forward and central difference gradient approximations and compared with experiments. The satisfactory agreement between theoretical predictions and experimental results for a wide range of search space dimensionalities confirmed the validity of the proposed approach.

Keywords: Implicit filtering, evolution strategies, analysis, theory

Exploring the common concepts of adaptive MCMC and Covariance Matrix Adaptation schemes

Christian Mueller (ETH Zürich, CH)

In the past decade, adaptive Markov Chain Monte Carlo (MCMC) methods gained considerable attention in the statistics community. In many modeling situations, one is confronted with the task of drawing samples from a target distribution that is only known up to a normalizing constant. Common examples arise in Bayesian statistics and statistical physics. Adaptive MCMC methods try to learn an optimal proposal distribution from previously accepted samples in order to efficiently explore the target distribution. I try to present both theoretical and practical advancements that have been achieved in the past years. Special attention is given to the commonalities and differences of these methods and Evolution Strategies with Covariance Matrix Adaptation. I also present how both types of methods can be unified within the Gaussian Adaptation framework.

Keywords: Adaptive MCMC, Markov Chains, CMA-ES, Gaussian Adaptation, Metropolis criterion

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2010/2813>

See also: <http://www.mosaic.ethz.ch/research/pubs/docs/Muller2010a.pdf>

Computational Complexity Analysis of Simple Genetic Programming On Two Problems Modeling Isolated Program Semantics

Frank Neumann (MPI für Informatik - Saarbrücken, DE)

Analyzing the computational complexity of evolutionary algorithms for binary search spaces has significantly increased their theoretical understanding. With this paper, we start the computational complexity analysis of genetic programming. We set up several simplified genetic programming algorithms and analyze them on two separable model problems, ORDER and MAJORITY, each of which captures an important facet of typical genetic programming problems. Both analyses give first rigorous insights on aspects of genetic programming design, highlighting in particular the impact of accepting or rejecting neutral moves and the importance of a local mutation operator.

Joint work of: Durrett, Greg; Neumann, Frank; O'Reilly, Una-May

Mapping the topology of combinatorial landscapes

Gabriela Ochoa (University of Nottingham, GB)

We have recently introduced a model of combinatorial landscapes: Local Optima networks (LON) .

This model allows the use of complex network analysis techniques in connection with the study of fitness landscapes and problem difficulty in combinatorial optimization. The model, inspired by work in the physical sciences on energy surfaces, is based on the idea of compressing the information given by the whole problem configuration space into a smaller mathematical object which is the graph having as vertices the optima configurations of the problem and as edges the possible transitions between these optima.

Our initial work considered binary search spaces and the NK family of abstract landscapes. Recently, we have turned our attention to more realistic combinatorial spaces (permutation spaces) and have analyzed the structure of two classes of instances of the Quadratic Assignment problem (QAP). For these landscape classes, we have found interesting results concerning the clustering or community structure of the local optima. We have additionally explored alternative definition of edges that produce more manageable graphs, and would allow a more direct connection with the dynamics of heuristic search algorithms. This talk will summarize our recent and potential future developments, with particular emphasis on the possibilities for visualizing the structure of landscapes, and analyzing the dynamics of heuristic search algorithms.

Keywords: Fitness Landscape Analysis, Local Optima, Complex Networks, Combinatorial Spaces, QAP, Visualization

Joint work of: Ochoa, Gabriela, Verel, Sebastian, Tomassini, Marco, Daolio, Fabio

Fixed Parameter Evolutionary Algorithms and Maximum Leaf Spanning Trees

Pietro Oliveto (University of Birmingham, GB)

We investigate the NP-hard problem of computing a spanning tree that has a maximal number of leaves by evolutionary algorithms in the context of

fixed parameter tractability (FPT) where the maximum number of leaves is the parameter under consideration. Investigating two common mutation operators, we show that an operator related to spanning tree problems leads to an FPT running time in contrast to a general mutation operator that does not have this property.

Keywords: Evolutionary Algorithms, Parameterized Complexity, Maximum Leaf Spanning trees, Runtime Analysis

Joint work of: Oliveto, Pietro; Kratsch, Stefan; Lehre, Per Kristian; Neumann, Frank

MAXSAT: The Anatomy of the Landscape of a Hard Optimisation Problem

Adam Pruegel-Bennett (University of Southampton, GB)

The landscape of random instances of Max-3-Sat is analysed empirically. It is shown why this problem becomes difficult for local search algorithms as the problem size becomes large. It is shown that there exists very distinctive long-range correlations in the landscape. A new population-based algorithm (landscape guided search—LGS) is presented which exploits the long-range correlation. We show that this algorithm can significantly out-perform the current state-of-the-art algorithms for this problem.

Keywords: Combinatorial optimisation, landscape, MAXSAT

Convergence of preference functions

Jonathan E. Rowe (University of Birmingham, GB)

A preference function is a function which selects a subset of objects based on (partial) information. As information increases, different objects may be selected. We examine conditions under which the selection of objects converges to the choice that would be made if full information were available, making use of tools from domain theory. The work is motivated by previous research (by Ficici, Bucci, Popovici) on co-evolutionary algorithms in which an evolving population of agents interact with each other and, it is hoped, produce better and better quality behaviour. The formalisation of how quality can be measured in this context has introduced the concept of a convex (or, as Ficici calls it, "monotonic") preference function (or solution concept). We simplify and extend the scope of this previous work, examining the relationship between convexity and convergence properties.

Keywords: Preference functions, convergence, continuity, convex function, co-evolutionary algorithms

Joint work of: Rowe, Jonathan E.; Jung, Achim

Convergence Rates of Multiobjective Evolutionary Algorithms

Guenter Rudolph (TU Dortmund, DE)

We prove that certain versions of multiobjective evolutionary algorithms (MOEAs) are algorithmically equivalent to certain versions of singleobjective evolutionary algorithms (EAs) under specific scalarizations of the vector-valued objective function.

As a consequence, we can transcribe results proven for singleobjective EAs to MOEAs with little effort.

The results presented at the seminar are published in the following proceedings:

N. Beume, M. Laumanns, and G. Rudolph: Convergence Rates of (1+1) Evolutionary Multiobjective Algorithms, pp. 597-606 in: R. Schaefer et al. (eds.): Proceedings of 11th Int'l Conf. on Parallel Problem Solving from Nature (PPSN XI), Part I, Springer: Berlin 2010.

N. Beume, M. Laumanns, and G. Rudolph: Convergence Rates of SMS-EMOA on Continuous Bi-Objective Problem Classes, in: H.-G. Beyer and W. Langdon (eds.): Proceedings of 11th Int'l Conf. on Foundations of Genetic Algorithms (FOGA XI), ACM Press, 2011 (at press).

Keywords: Multiobjective optimization, SMS-EMOA, convergence rate

Joint work of: Rudolph, Guenter; Beume, Nicola; Laumanns, Marco

Runtime and Speedups in Parallel Evolutionary Algorithms

Dirk Sudholt (University of Birmingham, GB)

Parallelization is a popular way for speeding up optimization tasks. Parallel evolutionary algorithms often distribute the evolution on subpopulations, each evolving on a distinct processor or "island." Good individuals are exchanged between the islands with respect to a given communication topology connecting the islands. I will review recent theoretical advances on the runtime of parallel evolutionary algorithms, with a focus on possible speedups in terms of the parallel computation time.

Keywords: Parallel evolutionary algorithms, island model, migration

Local Statistics of Bounded Pseudo-Boolean Functions

Andrew M. Sutton (Colorado State University, US)

The short-term dynamics of local search and mutation-based evolutionary algorithms are strongly influenced by the statistical structure of the fitness landscape region that is near the current search point. We show that if a fitness function is epistatically bounded, then low moments of the function over local Hamming regions can be constructed in polynomial time. The method takes advantage of the fact that k -bounded pseudo-Boolean functions have a sparse representation in the Walsh basis expansion and each Walsh polynomial is an eigenfunction of a generalized adjacency structure. We then show that, when the cardinality of the function range is asymptotically bounded, the distribution of fitness values over

any Hamming region can be approximated by a solution of an underdetermined linear system involving the moments. Finally, we show how the analysis can be used in practice to direct local hill-climbing search algorithms in MAX-SAT across plateau regions.

Keywords: Pseudo-Boolean functions, Walsh analysis, MAX-SAT, fitness landscape analysis

Joint work of: Sutton, Andrew M.; Whitley, L. Darrell; Howe, Adele E.

Decidability and complexity in partially observable antagonist coevolution

Olivier Teytaud (INRIA Saclay - Orsay, FR)

The usual formalization of decidability in partially observable problems has the drawback that a problem might be decidable, without the optimal move being actually computable. We propose an alternate formulation, show that settings which are decidable for the usual definition are not for this one.

We conjecture that the complexity in decidable subcases (bounded horizon) is also much higher and show upper complexity bounds by analyzing coevolutionary algorithms.

Keywords: Coevolution, partial observability

How Crossover Speeds Up Evolutionary Algorithms for the Multi-Objective All-Pairs-Shortest-Path Problem

Madeleine Theile (TU Berlin, DE)

Understanding the impact of crossover in evolutionary algorithms is one of the major challenges in the theoretical analysis of these stochastic search algorithms.

Recently, it has been shown that crossover provably helps to speed up evolutionary algorithms for the classical all-pairs-shortest path (APSP) problem. This approach can be extended to the NP-hard multi-criteria APSP problem. Based on rigorous runtime analyses it is possible to show that crossover leads to better worst case bounds than previous known results. This is the first time that rigorous runtime analyses have shown the usefulness of crossover for an NP-hard multi-criteria optimization problem.

In this talk I gave a short overview of my research interests with a particular focus on the most recent results on crossover and the multi-objective APSP problem.

Keywords: Crossover, multi-objective problems, APSP problem

The Scope and Limits of Darwinian Evolution

Leslie Valiant (Harvard University, US)

Living organisms function according to complex mechanisms that operate in different ways depending on conditions. Darwin's theory of evolution suggests that such mechanisms evolved through random variation guided by natural selection. However, there has existed no theory that would explain quantitatively which mechanisms can so evolve in realistic population sizes within realistic time periods, and which are too complex.

In this paper we suggest such a theory. Evolution is treated as a form of computational learning from examples in which the course of learning is influenced only by the aggregate fitness of the current hypothesis on the examples, and not otherwise by specific examples. We formulate a notion of evolvability that distinguishes function classes that are evolvable with polynomially bounded resources from those that are not.

It is shown that in any one phase of evolution where selection is for one beneficial behavior, monotone Boolean conjunctions and disjunctions are demonstrably evolvable over the uniform distribution, while Boolean parity functions are demonstrably not. The framework also allows a wider range of issues in evolution to be quantified.

Less formally, we suggest that the overall mechanism that underlies biological evolution is evolvable target pursuit, which consists of a series of evolutionary stages, each one pursuing an evolvable target in the technical sense suggested above, each target being rendered evolvable by the serendipitous combination of the environment and the outcomes of previous evolutionary stages.

Keywords: Evolvability, learning, complexity

Towards a Definition of Population based Fitness Landscapes

Sebastien Verel (INRIA - Lille, FR)

In combinatorial optimization, fitness landscapes are usually defined on the solutions space with the neighborhood defined with the mutation operator. This approach helps to design local searches and to understand why they succeed or not. But it is not clear how the study of the solutions space is related to the dynamics of evolutionary algorithms which use a population, and neither how to study the crossover operator. This work tries to define a fitness landscapes where the points are a population of solutions.

Keywords: Fitness landscapes

Monotone Black-Box and Higher arity Black-Box Complexities

Carola Winzen (MPI für Informatik - Saarbrücken, DE)

Building on recent work of Lehre and Witt (GECCO 2010) we introduce an alternative black-box model, which we call *monotone black-box model*.

In contrast to the unbiased model, where the algorithm learns absolute fitness values, we assume in the monotone black-box model that the algorithm does only learn the *relative fitness* of the queried search points.

Surprisingly, the restriction of information does not increase the black-box complexity of the generalized OneMax functions. This picture changes completely if we consider the generalized BinaryValue functions. Whereas their (*-ary) unbiased black-box complexity is of order $O(\log n)$, their monotone black-box complexity is shown to be linear.

Joint work of: Winzen, C.; Doerr, B.; Johannsen, D.; Kötzing, T.; Lehre, P.; Wagner, M.

Keywords: Randomized search heuristics, black-box optimization, runtime analysis

Modeling Drift in a Steady State GA

Alden Wright (University of Montana - Missoula, US)

The general building block hypothesis says that crossover can combine partial solutions from two individuals into one individual. A genetic algorithm (GA) applied to the concatenated trap problem provides a prototypical example of the building block hypothesis. A GA using a panmictic population must use a large population to counteract drift.

How such a GA applied to this problem scales with problem size is an unsolved problem. The ONEMAX problem can be used as a surrogate for the concatenated trap problem if the initial frequency of ones (which correspond to trap optima) is less than 1/2. This paper gives two interrelated models of a steady state crossover-selection (no mutation) GA applied to the ONEMAX problem. Properties of these models are shown.

A crucial assumption of these models is linkage equilibrium, and by using a variation of the GA which uses a form of gene pool crossover, it is shown that the linkage-equilibrium based models are not very accurate.

Keywords: Drift, genetic algorithm, building block, crossover, Moran process, Markov chain

Artificial Immune Systems as Randomized Search Heuristics

Christine Zarges (TU Dortmund, DE)

Artificial Immune Systems (AIS) are an emerging new field of research in Computational Intelligence. They are applied in many different areas and ways, e.g., optimization, anomaly detection and classification. Like Evolutionary Algorithms (EAs) they are a special class of biologically inspired randomized search heuristics, in their case based on the immune system of vertebrates. Applied to optimization problems many similarities between EAs and AIS can be found. But in contrast to EAs, which emerged from a single main concept, AIS derive from various immunological theories, namely the clonal selection principle, negative selection, immune networks, and the danger theory.

Despite many empirical studies, the theoretical foundation of AIS is still in its infancy. The theoretical foundation of AIS is stated as a major challenge for the future of the field. In this talk, different AIS approaches for optimization, in particular concrete algorithms based on clonal selection and immune networks are presented. We discuss existing results and important aspects for future research.