

NATURAL COMPUTING SERIES

L. Kallel · B. Naudts · A. Rogers (Eds.)

Theoretical Aspects of Evolutionary Computing



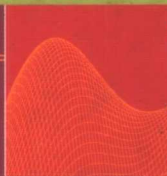
Springer

Quantum Computing

Neural Networks

Evolutionary Computing

DNA Computing



TP - 8832 - 72048

TP-399

Leila Kallel • Bart Naudts • Alex Rogers (Eds.)

Theoretical Aspects of Evolutionary Computing

With 129 Figures and 15 Tables



Springer

Editors

Dr. Leila Kallel
Ecole Polytechnique
CMAP
91128 Palaiseau Cedex, France
kallel@cmapx.polytechnique.fr
Dr. Bart Naudts
Universiteit Antwerpen (RUCA)
Departement Wiskunde & Informatica
Groenenborgerlaan 171
2020 Antwerpen, Belgium
bnaudts@ruca.ua.ac.be

Alex Rogers
University of Southampton
Department of Electronics and Computer Science
Highfield
SO17 1BJ Southampton, UK
ar97r@ecs.soton.ac.uk

Series Editors

G. Rozenberg (Managing Editor)
Th. Bäck, A.E. Eiben, J.N. Kok,
H.P. Spaiink
Leiden Center for Natural Computing
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands
rozenber@cs.leidenuniv.nl

Library of Congress Cataloging-in-Publication Data

Theoretical aspects of evolutionary computing/Leila Kallel, Bart Naudts, Alex Rogers (eds.).
p. cm. – (Natural computing series)

Based on the lectures and workshop contributions of the second EvoNet Summer School on Theoretical Aspects of Evolutionary Computing held at the Middelheim campus of the University of Antwerp, Belgium during the first week of September 1999.

Includes bibliographical references and index.

ISBN 3540673962 (alk. paper)

I. Evolutionary programming (Computer science) I. Kallel, Leila, 1972– II. Naudts, Bart, 1973– III. Rogers, Alex, 1970– IV. EvoNet Summer School on Theoretical Aspects of Evolutionary Computing (2nd: 1999: University of Antwerp, Belgium) V. Series.

QA76.618.T47 2000
006.3–dc21

00-061910

ACM Computing Classification (1998): F.2.2, I.2.8, I.2.6, G.3, G.1.6, J.3, J.2

ISBN 3-540-67396-2 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York,
a member of BertelsmannSpringer Science+Business Media GmbH
<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2001
Printed in Germany

The use of general descriptive names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover Design: KunkelLopka, Heidelberg
Typesetting: Camera ready by the editors
Printed on acid-free paper SPIN 10765204 45/3142SR – 5 4 3 2 1 0

Preface

During the first week of September 1999, the Second EvoNet Summer School on Theoretical Aspects of Evolutionary Computing was held at the Middelheim campus of the University of Antwerp, Belgium. Originally intended as a small get-together of PhD students interested in the theory of evolutionary computing, the summer school grew to become a successful combination of a four-day workshop with over twenty researchers in the field and a two-day lecture series open to a wider audience.

This book is based on the lectures and workshop contributions of this summer school. Its first part consists of tutorial papers which introduce the reader to a number of important directions in the theory of evolutionary computing. *The tutorials are at graduate level and assume only a basic background in mathematics and computer science. No prior knowledge of evolutionary computing or its theory is necessary.* The second part of the book consists of technical papers, selected from the workshop contributions. A number of them build on the material of the tutorials, exploring the theory to research level. Other technical papers may require a visit to the library.

The theory of evolution is at the crossroads of mathematics, computer science, statistics, and the theoretical sides of physics, chemistry, and biology. There is often too little interaction between these disciplines. One of the summer school's goals was to get researchers from many disciplines together. As a result, this book contains papers from researchers with a background in complexity theory, neural networks, probability theory, population genetics, statistical physics, and mathematics. Only a fraction of the authors grew up in the evolutionary computing community itself. This is reflected in the book, which presents a rich variety of approaches to the study of the behavior of evolutionary algorithms.

Overview of the Tutorials

The tutorial part of the book starts with a general introduction to evolutionary computing, from the point of view of an experienced engineering researcher, A. J. Keane.

who uses evolutionary algorithms as one of his main tools in practical optimization. This first tutorial has been written in cooperation with A. Rogers.

Before plunging into the theory, we present a second tutorial on evolutionary algorithms, written by A. E. Eiben. He discusses the use of evolutionary algorithms for constrained search problems.

A first course on the microscopic approach to the dynamics of the simple genetic algorithm, as started by M. D. Vose in the early 1990s, is given by J. E. Rowe. Filling in the actual numbers for a concrete example, but also commenting on the latest development, the tutorial is all the reader needs to get started with this exciting mathematical framework.

The statistical physics approach to the dynamics of evolutionary algorithms is a macroscopic one: instead of modeling the exact dynamics, only a few statistics of the dynamics are computed. A. Prügel-Bennett and J. L. Shapiro have written two complementary tutorials on this line of research.

Following the line of macroscopic descriptions is the tutorial on evolution strategies by H.-G. Beyer and D. V. Arnold. Evolution strategies operate on real spaces. Their dynamics on a number of fitness models have now been fully analyzed.

Apart from being successful in practical applications, optimization algorithms based on a model of the fitness distribution can be used to approximate genetic algorithms. H. Mühlenbein and T. Mahnig introduce this approach to studying the dynamics of genetic algorithms, also showing the link with population genetics and learning with Bayesian networks.

Closing the gap between theory and practice is the tutorial by L. Kallel, B. Naudts, and C. R. Reeves on properties of fitness landscapes and their impact on problem difficulty. The tutorial explores issues such as Walsh coefficients, epistasis, correlations in the fitness landscape, basins of attraction, and interaction structure.

Overview of the Technical Papers

In the first technical paper, A. Rogers and A. Prügel-Bennett build on the tutorials by Prügel-Bennett and Shapiro. They analyze the dynamics of a genetic algorithm with ranking selection, uniform crossover, and mutation on the ‘basin with a barrier’ model, which contains one feature commonly found in hard search problems. A relationship between the optimal mutation rate and population size is found, which is independent of the problem size.

M. Oates and coworkers also examine the relationship between population size and mutation rate. Their exhaustive experimental work on a number of search problems is presented in the second technical paper.

The third and last paper which deals with search in time-static problems is by D. V. Arnold, who presents the state of the art of evolution strategies on noisy fitness models. Beyer and Arnold’s tutorial on evolution strategies serves as reference material.

The quasispecies model from theoretical biology is studied by J. E. Rowe, and also by C. Ronnewinkel, C. O. Wilke, and T. Martinetz. Both use the microscopic model from Rowe's tutorial to study the quasispecies adaptability to fitness function changes by computing error thresholds and optimal mutation rates. However, their approach to time-dependence is different. The former identifies cyclic attractors, whereas the latter accumulate the time-dependency of the change cycle.

A. Berny demonstrates the link between statistical learning theory and combinatorial optimization. Rather than searching in the space of bit strings, he demonstrates how to search, using gradient techniques, in the space of distributions over bit strings.

The genetic drift caused by multi-parent scanning crossover is rigorously calculated by C. A. Schippers. He finds that occurrence-based scanning produces a stronger than usual, self-amplifying genetic drift.

Reference material for the next six papers is the tutorial by Kallel, Naudts, and Reeves. Instead of studying the dynamics of evolutionary algorithms, they focus on the representation of search problems and the fitness landscapes induced by the algorithm.

When choosing representation and genetic operators, many implicit assumptions about the epistatic structure of the fitness function to be optimized are made. For example, one-point crossover assumes that adjacent bit positions are highly correlated. F. Burkowski proposes nullifying these assumptions, and learning the epistatic structure on-line.

J. Garnier and L. Kallel develop statistical techniques to estimate the number and sizes of the basins of attraction in a fitness landscape.

Can fitness functions be classified according to the difficulty they present to a genetic algorithm? T. Jansen comments on this issue and details some of the flaws that a number of measures, introduced as predictors of problem difficulty, suffer from.

The subject of symmetry in the representation of the search space has been studied in three independent contributions. First, there is the paper of P. Stagge and C. Igel, which deals with redundancy caused by isomorphic structures in the context of topology optimization of neural networks. In a second contribution, A. Marino introduces specific genetic operators for the graph coloring problem, in order to search the highly symmetric solution space more efficiently. Finally, C. Van Hoyweghen introduces a technique to detect spin-flip symmetry in arbitrary search spaces.

The last contribution to this book, and also the largest in size, is by P. Del Moral and L. Miclo. They first propose a model for GAs which generalizes Vose's model (introduced in the tutorial by Rowe) to continuous search spaces, inhomogeneous fitness environments, and inhomogeneous mutation operators. In this model a population is seen as a distribution (or measure) on the search space, and the genetic algorithm as a measure-valued dynamical system. Then, they study the convergence properties of the finite population process when the population size tends to infinity, and prove that it converges uniformly with respect to time to the infinite population

process. Finally, they show how their model of GAs can be used in the context of nonlinear filtering. They prove the efficiency of finite population GAs to approximate an observed signal in the presence of noise.

Acknowledgments

We would like to thank all the participants of the summer school for attending and sending in valuable contributions afterwards. This book is their work. We were also pleased with the fact that so many senior researchers were willing to write a 20+ page tutorial on their subject.

Finding sponsorship to organize a summer school on a theoretical subject was a lot easier than we thought. We would like to thank the Network of Excellence in Evolutionary Computing (EvoNet), the Doctoral Study Programme Science–Pharmacy of the University of Antwerpen (DOCOP), the Foundation for European Evolutionary Computation Conferences (FEECC), and the Leiden Institute for Advanced Computer Science (LIACS) for their helpful cooperation. The editing of this book was done while B. Naudts was funded as a Postdoctoral Fellow of the Fund for Scientific Research – Flanders, Belgium (F.W.O.).

But above all, we are greatly indebted to our summer school co-organizers: Jennifer Willies of EvoNet, Jano van Hemert of the LIACS, and Ives Landrieu, Luk Schoofs, and Clarissa Van Hoyweghen of the University of Antwerpen. They did a great job!

*Paris, Antwerpen, and Southampton,
Spring 2001*

*L. Kallel
B. Naudts
A. Rogers*

Table of Contents

Part I: Tutorials

Introduction to Evolutionary Computing in Design Search and Optimisation <i>A. J. Keane</i>	1
Evolutionary Algorithms and Constraint Satisfaction: Definitions, Survey, Methodology, and Research Directions <i>A. E. Eiben</i>	13
The Dynamical Systems Model of the Simple Genetic Algorithm <i>J. E. Rowe</i>	31
Modelling Genetic Algorithm Dynamics <i>A. Prügel-Bennett and A. Rogers</i>	59
Statistical Mechanics Theory of Genetic Algorithms <i>J. L. Shapiro</i>	87
Theory of Evolution Strategies – A Tutorial <i>H.-G. Beyer and D. V. Arnold</i>	109
Evolutionary Algorithms: From Recombination to Search Distributions <i>H. Mühlenbein and T. Mahnig</i>	135
Properties of Fitness Functions and Search Landscapes <i>L. Kallel, B. Naudts, and C. R. Reeves</i>	175

Part II: Technical Papers

A Solvable Model of a Hard Optimisation Problem <i>A. Rogers and A. Prügel-Bennett</i>	207
---	-----

Bimodal Performance Profile of Evolutionary Search and the Effects of Crossover <i>M. Oates, J. Smedley, D. Corne, and R. Loader</i>	223
Evolution Strategies in Noisy Environments – A Survey of Existing Work <i>D. V. Arnold</i>	239
Cyclic Attractors and Quasispecies Adaptability <i>J. E. Rowe</i>	251
Genetic Algorithms in Time-Dependent Environments <i>C. Ronnewinkel, C.O. Wilke, and T. Martinetz</i>	261
Statistical Machine Learning and Combinatorial Optimization <i>A. Berny</i>	287
Multi-Parent Scanning Crossover and Genetic Drift <i>C. A. Schippers</i>	307
Harmonic Recombination for Evolutionary Computation <i>F. J. Burkowski</i>	331
How to Detect all Maxima of a Function <i>J. Garnier and L. Kallel</i>	343
On Classifications of Fitness Functions <i>T. Jansen</i>	371
Genetic Search on Highly Symmetric Solution Spaces: Preliminary Results <i>A. Marino</i>	387
Structure Optimization and Isomorphisms <i>P. Stagge and C. Igel</i>	409
Detecting Spin-Flip Symmetry in Optimization Problems <i>C. Van Hoyweghen</i>	423
Asymptotic Results for Genetic Algorithms with Applications to Nonlinear Estimation <i>P. Del Moral and L. Miclo</i>	439
Index	495

An Introduction to Evolutionary Computing in Design Search and Optimisation

A. Keane

Lecture notes produced by A. Rogers

Department of Mechanical Engineering

University of Southampton

Southampton SO17 1BJ, UK

E-mail: andy.keane@soton.ac.uk

Abstract. Evolutionary computing (EC) techniques are beginning to find a place in engineering design. In this tutorial we give a brief overview of EC and discuss the techniques and issues which are important in design search and optimisation.

Keywords

Evolutionary computing, design search, optimisation, problem solving environment

1 Introduction

The purpose of this introduction is to present, to a wide audience, some of the ideas which are important in the application of evolutionary computing (EC) algorithms to engineering design. This leads on to an area known as *design search and optimisation* where evolutionary computing techniques are beginning to find real applications.

Perhaps the first question to ask is: why do we do this at all? The general answer is that life is becoming more computationally complex and in the field of engineering we are presented with an increasing number of computationally complex problems to solve. The computational complexity of optimising these systems challenges the capabilities of existing techniques and we seek some other heuristics. When we look around at the real world, we see many examples of complex systems, ourselves included, which have arisen through evolution. This naturally leads to some desire to emulate evolution in our attempts at problem solving.

As an example of a computationally complex problem, we can consider designing aircraft wings – an area from which several of the examples in this tutorial will be drawn. Alternatively, we may be trying to schedule some process in a factory or simply trying to fit curves to data. In some senses, all these things are about design, and by design, we mean the idea that we want to create something. The whole point to evolutionary computing is the creation of designs, particularly in the domain of complex problems and environments.

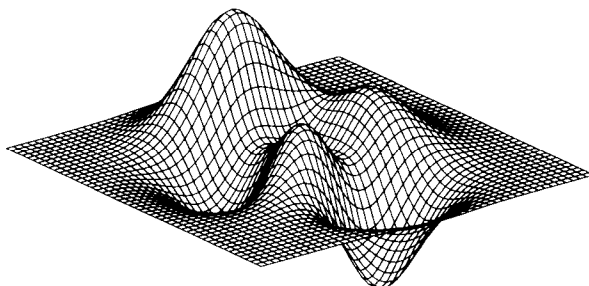


Figure 1. Example of an optimisation landscape

It is useful to briefly consider what design is, how it is done and where evolutionary computing fits in. We are all surrounded every day by extremely complicated technology which we generally take for granted. As an example, consider the cars we drive. Should we have a crash, the airbags will explode and the car will crumple safely around us because someone has thought about and modeled all the processes and components involved in that crash. To do this nowadays, automotive engineers use high-performance computer networks which are highly developed into sophisticated environments.

Designing complex things such as cars requires us to put lots of things together. It requires us to know something about what it is we are trying to design, to have some way of interfacing it to the user and perhaps some type of search and optimisation routine with resource management. The whole thing may be called a *problem solving environment* (PSE) and it is into this environment that evolutionary computing techniques fit.

In summary, designers in the broadest sense – simply people trying to create something – increasingly see the need for a PSE. Ideally, these PSEs contain some sort of search and optimisation and it is here that evolutionary computing techniques find an application. Thus, evolutionary computing is not used in isolation but as a part of a design tool which has to be of some utility.

What follows are some introductory remarks on several topics which will be expanded on later in the course of this tutorial.

Search and Optimisation. Two terms which are often used in evolutionary computing are search and optimisation. We can illustrate what these mean by considering a simple landscape – see figure 1. Optimisation is about trying to walk up a hill and hopefully, through effectively searching the landscape, walking up the highest hill.

At its simplest this is all we are trying to do. We have some landscape which we are trying to walk through. Usually we cannot see the landscape, so the whole business of evolutionary computing is trying to do this blindfold. This introduces questions

such as: how do we know we are at the top of the hill and how do we know it is the right hill?

Resources. Search and optimisation always requires some other piece of code to be plugged into it – we apply evolutionary computing to something. It is usually a piece of analysis code and we are trying to capture the intentions of the designer in this code. Given the computational complexity of the problems we are considering, computing power dominates this process. Given overwhelming computing power, most problems can be solved. With finite computing resources we have to be more careful.

Current Techniques. What is the current state of the art in engineering design? Most PSEs now include some sort of hill-climbing optimisation and parallel computation is common. Evolutionary search is known about but it is not generally in everyday use. The use of meta-computing techniques and resource scheduling are beginning to be considered as approaches to the problems of limited computing power.

Representation. Representation is one of the key issues in this work. How do we represent what it is that we are trying to design? Inevitably people represent their designs with numbers, so the method of encoding between the two is important. However, an engineer also has a selection of models which can be used which range in sophistication and cost. It is not simply a question of representation and method but representation, method, and model which seem to be the most important. Designers have aspirations, constraints, and varying and multiple objectives. They often don't know what they want or cannot express what it is exactly they want. But in general, they want robust designs.

Classical search methods have been around for at least fifty years and they are well tried. Crucially, people don't believe they are likely to lead to any more innovation. On the other hand, evolutionary computing methods have some advantages but they have a problem – computational expense. If we come up with an evolutionary computing paradigm which is really very powerful but is too computationally expensive, it will not be used.

2 A Brief Overview of Evolutionary Computing

The history of evolutionary computing goes back to the 1960s with the introduction of ideas and techniques such as genetic algorithms, evolutionary strategies and evolutionary programming [3]. Whilst all differ slightly in their actual implementations, all these evolutionary computing techniques use the same metaphor of mapping problem solving onto a simple model of evolution.

EVOLUTION		PROBLEM SOLVING
Individual	↔	Candidate Solution
Fitness	↔	Quality
Environment	↔	Problem

We have a population which we want to evolve from time step t to time step $t + 1$ by selection, recombination and mutation. There is some contention about which of the two, mutation or recombination, is the more powerful, with there being a whole history in the German academic community of using only mutation in evolutionary strategies. To a certain extent though, the ingredients can be mixed and matched. One is trying to introduce diversity in the designs through mutation and recombination, and then exploit this exploration through selection of better solutions; balancing the growth of ideas alongside the selection of ideas. One interesting aspect is that this balance occurs all the way through but the nature of the algorithms we use implies that there tends to be a large diversity at the beginning and very little towards the end.

Hopefully through the correct choice of operators, we achieve a balance between exploration and exploitation and our solutions steadily get better.

Although a great deal of the evolutionary computing literature concerns comparisons of one algorithm to another, some general points about the advantages of evolutionary computing techniques can be made:

- Widely applicable.
- Low development and application cost.
- Easily incorporated into other methods.
- Solutions are interpretable.
- Can be run interactively and allows incorporation of user-proposed solutions.
- Provide many alternative solutions.

Whilst they have been shown to be useful in many areas, evolutionary computing techniques have several disadvantages when compared to some other techniques:

- No guarantee for optimal solution within finite time.
- Weak theoretical basis.
- May need parameter tuning for good performance.
- Often computationally expensive and thus slow.

Apart from the computational expense, perhaps that which most directly effects the acceptance of evolutionary computing is the weakness of their theoretical basis. With no theory, we are left with a set of rules of thumb for choosing and tuning the algorithms.

3 Evolutionary Computing in Design Search and Optimisation

When actually implementing an evolutionary computing algorithm we are presented with a wide range of choices. Which algorithm should we use? Which operators should we use? How do we set the parameters? There are a number of books available which help with these questions [1–4,7]; however, when using these algorithms in engineering design we have a number of other additional issues which must be considered.

3.1 Representation

Deciding on a good representation is fundamental to the performance of evolutionary computing techniques. The algorithms work on numbers, usually binary, but we are not trying to design a string of numbers. We may be trying to design a wing, a communications network or a schedule, and in some way we have to link the two together. This point turns out to be absolutely critical to the success of the application.

It is self-evidently true that we can trade-off the ability of any representation to be compact, against its ability to represent all possible designs. This is a really fundamental decision. Do we want to describe any possible design or do we want to describe something in a small region around where we currently stand?

If we are using an evolutionary algorithm to design the wing for a jet aircraft, do we want to think about bi-planes? Probably not. Do we want to consider a wing of alternative materials? Maybe. Someone has to make these decisions and it cannot be done without some domain expertise. Experience shows you need a domain expert and an evolutionary computation expert together, as it is this critical stage which often determines the outcome of the project.

If we again consider our wing design problem suggested above, the classical aerodynamics of the wing are based on concepts of wing span, sweep, chord and camber. The interesting thing about these words is that these are ways which aerodynamicists have worked out to describe wings. They are a particular representation of a wing which is relevant when considering the flow of air over it.

An alternative would be to use the x, y and z coordinates for the surfaces. This presents the problem that whilst we can describe a wing, we can also describe virtually anything else as well. Moving down the trade-off in compactness, we are increasing the complexity of the object which we can describe and are thus increasing the domain in which the potential designs can exist. We are making the problem potentially harder and if we do not need the complexity at this level, we run the risk of making the problem almost insoluble.

The opposite danger comes from sticking too rigidly to an existing representation. Through existing design techniques the representation may have become so specific that it traps the possible solutions into a local optimum and does not have the generality to describe wing shapes which may indeed be better.

Choices across this sort of domain make a real difference and it is clear that they cannot be made without some domain knowledge.

3.2 Constraints and Multiple Objectives

Often the objectives of a design are not completely defined. An engineer will often have multiple objectives and in some sort of hierarchy. Unless this appraisal of potential designs is somehow formulated and included in the analysis code, the expectations of the engineer will differ from the designs which our algorithm is actually producing.

Constraints are another issue which is yet to be fully addressed. How do we handle a list of certain criteria which invalidate a particular design? There are a number of techniques, some of which are discussed by Eiben in this volume (page 13), and we must decide which one is most relevant.

3.3 Mutation

Mutation is critical to a lot of these results. It is often said to be there just to ensure that every part of the search space may be reached. However, successful search algorithms are often run just using mutation and selection without crossover. Indeed evolutionary strategies are predominantly used in this fashion.

The idea behind mutation is that we are making local steps in our landscape. Our choice of mutation operator and representation should mean that we are making small changes to our design and not leaping to a radically different solution.

3.4 Recombination

Recombination or crossover allows parents to pass on some of their characteristics to their children. The motivation being that we can use good parents to develop even better children. There are many ways of doing this but all involve the combination of parts of one parent with the complementary parts from another.

There is a trade-off between recombination and mutation which reflects some of the history of the trade-off between the various communities. Recombination tends to be seen as more of an exploitation operator and mutation as more of a exploration operator. Some people would tend to suggest mixing these in, so we do more exploration at the beginning and more exploitation towards the end. These are more examples of parameter and algorithm tuning which are dictated by rules of thumb rather than a solid theoretical basis.

3.5 Niching

When we initially discussed optimisation, we considered trying to find the single global optimum in the problem. In engineering applications this may not be the case and the question has to be asked whether we want the best solution or do we want to know about a number of areas in the problem space where there are good candidate solutions? Depending on our aspirations, we may want to impose some form of niching techniques to allow the population to divide into smaller sub-populations, each focused around a different part of the problem space.

3.6 Repeatability and Elitism

Evolutionary computing techniques are stochastic and thus no two runs will necessarily produce the same results. Whilst in some applications this is acceptable, it can be quite discouraging to produce a particular good solution once and never to be able to find it again.

This leads to many issues concerning the repeatability of evolutionary computing algorithms. Do we just want to get one very good solution or do we want a good solution every time? If we were doing process scheduling we might settle for the former but an automotive engineer designing a car would probably prefer the latter.

Another concern which stems from the stochastic nature of the algorithm, is that a good solution once found within the population may be lost later in the evolution. Whilst it always makes sense to keep a record of the best solutions, elitism strategies ensure that these solutions stay in the population.

3.7 When to Stop?

Having set up the representation and run the evolutionary computing algorithm, the final decision is when to stop it. Without any prior knowledge, it is impossible to tell whether the best solution has been reached. The decision most often comes down to one of time and computing resources.

4 An Example of EC Design

As an example of the techniques and problems discussed we will consider an actual case of an evolutionary computing algorithm being used to solve an engineering problem. Figure 2 shows a photograph from a NASA mission in 1987 as part of a proof of concept programme to show that astronauts could build structures in space.

One of the problems with these structures is that they tend to have very severe vibration problems. They are light regular alloy structures in an environment where there



Figure 2. NASA photograph showing an astronaut constructing a boom in space

is no air to provide damping. We started to look at the vibrational characteristics of structures like these when subject to some vibrational noise at one end [6,5]. The objective was to use evolutionary computing techniques to design the geometry of the beam such that the vibrational noise does not travel through it. This is a straightforward engineering job which is actually critical to both ESA and NASA who plan to launch future missions with booms of this type fifty metres in length.

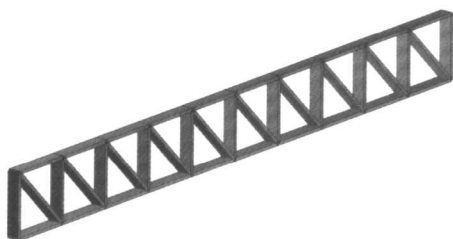


Figure 3. Diagram of simple 2D boom design