

AN INTRODUCTION TO GENETIC ALGORITHMS



MELANIE MITCHELL

An Introduction to Genetic Algorithms

Mitchell Melanie

A Bradford Book The MIT Press

Cambridge, Massachusetts • London, England

Fifth printing, 1999

First MIT Press paperback edition, 1998

Copyright © 1996 Massachusetts Institute of Technology

All rights reserved. No part of this publication may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Set in Palatino by Windfall Software using ZzT_EX.

Library of Congress Cataloging-in-Publication Data

Mitchell, Melanie.

An introduction to genetic algorithms / Melanie Mitchell.

p. cm.

"A Bradford book."

Includes bibliographical references and index.

ISBN 0-262-13316-4 (HB), 0-262-63185-7 (PB)

1. Genetics—Computer simulation. 2. Genetics—Mathematical models. I. Title.

QH441.2.M55 1996

575.1'01'13—dc20 95-24489

CIP

Table of Contents

An Introduction to Genetic Algorithms.....	1
Mitchell Melanie.....	1
Chapter 1: Genetic Algorithms: An Overview.....	2
Overview.....	2
1.1 A BRIEF HISTORY OF EVOLUTIONARY COMPUTATION.....	2
1.2 THE APPEAL OF EVOLUTION.....	4
1.3 BIOLOGICAL TERMINOLOGY.....	5
1.4 SEARCH SPACES AND FITNESS LANDSCAPES.....	6
1.5 ELEMENTS OF GENETIC ALGORITHMS.....	7
Examples of Fitness Functions.....	7
GA Operators.....	8
1.6 A SIMPLE GENETIC ALGORITHM.....	8
1.7 GENETIC ALGORITHMS AND TRADITIONAL SEARCH METHODS.....	10
1.9 TWO BRIEF EXAMPLES.....	12
Using GAs to Evolve Strategies for the Prisoner's Dilemma.....	13
Hosts and Parasites: Using GAs to Evolve Sorting Networks.....	16
1.10 HOW DO GENETIC ALGORITHMS WORK?.....	21
THOUGHT EXERCISES.....	23
COMPUTER EXERCISES.....	24
Chapter 2: Genetic Algorithms in Problem Solving.....	27
Overview.....	27
2.1 EVOLVING COMPUTER PROGRAMS.....	27
Evolving Lisp Programs.....	27
Evolving Cellular Automata.....	34
2.2 DATA ANALYSIS AND PREDICTION.....	42
Predicting Dynamical Systems.....	42
Predicting Protein Structure.....	47
2.3 EVOLVING NEURAL NETWORKS.....	49
Evolving Weights in a Fixed Network.....	50
Evolving Network Architectures.....	53
Direct Encoding.....	54
Grammatical Encoding.....	55
Evolving a Learning Rule.....	58
THOUGHT EXERCISES.....	60
COMPUTER EXERCISES.....	62
Chapter 3: Genetic Algorithms in Scientific Models.....	65
Overview.....	65
3.1 MODELING INTERACTIONS BETWEEN LEARNING AND EVOLUTION.....	66
The Baldwin Effect.....	66
A Simple Model of the Baldwin Effect.....	68
Evolutionary Reinforcement Learning.....	72
3.2 MODELING SEXUAL SELECTION.....	75
Simulation and Elaboration of a Mathematical Model for Sexual Selection.....	76
3.3 MODELING ECOSYSTEMS.....	78
3.4 MEASURING EVOLUTIONARY ACTIVITY.....	81
Thought Exercises.....	84
Computer Exercises.....	85

Table of Contents

Chapter 4: Theoretical Foundations of Genetic Algorithms.....	87
Overview.....	87
4.1 SCHEMAS AND THE TWO-ARMED BANDIT PROBLEM.....	87
The Two-Armed Bandit Problem.....	88
Sketch of a Solution.....	89
Interpretation of the Solution.....	91
Implications for GA Performance.....	92
Deceiving a Genetic Algorithm.....	93
Limitations of "Static" Schema Analysis.....	93
4.2 ROYAL ROADS.....	94
Royal Road Functions.....	94
Experimental Results.....	95
Steepest-ascent hill climbing (SAHC).....	96
Next-ascent hill climbing (NAHC).....	96
Random-mutation hill climbing (RMHC).....	96
Analysis of Random-Mutation Hill Climbing.....	97
Hitchhiking in the Genetic Algorithm.....	98
An Idealized Genetic Algorithm.....	99
4.3 EXACT MATHEMATICAL MODELS OF SIMPLE GENETIC ALGORITHMS.....	103
Formalization of GAs.....	103
Results of the Formalization.....	108
A Finite-Population Model.....	108
4.4 STATISTICAL-MECHANICS APPROACHES.....	112
THOUGHT EXERCISES.....	114
COMPUTER EXERCISES.....	116
5.1 WHEN SHOULD A GENETIC ALGORITHM BE USED?.....	116
5.2 ENCODING A PROBLEM FOR A GENETIC ALGORITHM.....	117
Binary Encodings.....	117
Many-Character and Real-Valued Encodings.....	118
Tree Encodings.....	118
5.3 ADAPTING THE ENCODING.....	118
Inversion.....	119
Evolving Crossover "Hot Spots".....	120
Messy Gas.....	121
5.4 SELECTION METHODS.....	124
Fitness-Proportionate Selection with "Roulette Wheel" and "Stochastic Universal"	
Sampling.....	124
Sigma Scaling.....	125
Elitism.....	126
Boltzmann Selection.....	126
Rank Selection.....	127
Tournament Selection.....	127
Steady-State Selection.....	128
5.5 GENETIC OPERATORS.....	128
Crossover.....	128
Mutation.....	129
Other Operators and Mating Strategies.....	130
5.6 PARAMETERS FOR GENETIC ALGORITHMS.....	130
THOUGHT EXERCISES.....	132
COMPUTER EXERCISES.....	133

Table of Contents

Chapter 6: Conclusions and Future Directions.....	135
Overview.....	135
Incorporating Ecological Interactions.....	136
Incorporating New Ideas from Genetics.....	136
Incorporating Development and Learning.....	137
Adapting Encodings and Using Encodings That Permit Hierarchy and Open-Endedness.....	137
Adapting Parameters.....	137
Connections with the Mathematical Genetics Literature.....	138
Extension of Statistical Mechanics Approaches.....	138
Identifying and Overcoming Impediments to the Success of GAs.....	138
Understanding the Role of Schemas in GAs.....	138
Understanding the Role of Crossover.....	139
Theory of GAs With Endogenous Fitness.....	139
Appendix A: Selected General References.....	140
Appendix B: Other Resources.....	141
SELECTED JOURNALS PUBLISHING WORK ON GENETIC ALGORITHMS.....	141
SELECTED ANNUAL OR BIENNIAL CONFERENCES INCLUDING WORK ON GENETIC ALGORITHMS.....	141
INTERNET MAILING LISTS, WORLD WIDE WEB SITES, AND NEWS GROUPS WITH INFORMATION AND DISCUSSIONS ON GENETIC ALGORITHMS.....	142
Bibliography.....	143

Chapter 1: Genetic Algorithms: An Overview

Overview

Science arises from the very human desire to understand and control the world. Over the course of history, we humans have gradually built up a grand edifice of knowledge that enables us to predict, to varying extents, the weather, the motions of the planets, solar and lunar eclipses, the courses of diseases, the rise and fall of economic growth, the stages of language development in children, and a vast panorama of other natural, social, and cultural phenomena. More recently we have even come to understand some fundamental limits to our abilities to predict. Over the eons we have developed increasingly complex means to control many aspects of our lives and our interactions with nature, and we have learned, often the hard way, the extent to which other aspects are uncontrollable.

The advent of electronic computers has arguably been the most revolutionary development in the history of science and technology. This ongoing revolution is profoundly increasing our ability to predict and control nature in ways that were barely conceived of even half a century ago. For many, the crowning achievements of this revolution will be the creation—in the form of computer programs—of new species of intelligent beings, and even of new forms of life.

The goals of creating artificial intelligence and artificial life can be traced back to the very beginnings of the computer age. The earliest computer scientists—Alan Turing, John von Neumann, Norbert Wiener, and others—were motivated in large part by visions of imbuing computer programs with intelligence, with the life-like ability to self-replicate, and with the adaptive capability to learn and to control their environments. These early pioneers of computer science were as much interested in biology and psychology as in electronics, and they looked to natural systems as guiding metaphors for how to achieve their visions. It should be no surprise, then, that from the earliest days computers were applied not only to calculating missile trajectories and deciphering military codes but also to modeling the brain, mimicking human learning, and simulating biological evolution. These biologically motivated computing activities have waxed and waned over the years, but since the early 1980s they have all undergone a resurgence in the computation research community. The first has grown into the field of neural networks, the second into machine learning, and the third into what is now called "evolutionary computation," of which genetic algorithms are the most prominent example.

1.1 A BRIEF HISTORY OF EVOLUTIONARY COMPUTATION

In the 1950s and the 1960s several computer scientists independently studied evolutionary systems with the idea that evolution could be used as an optimization tool for engineering problems. The idea in all these systems was to evolve a population of candidate solutions to a given problem, using operators inspired by natural genetic variation and natural selection.

In the 1960s, Rechenberg (1965, 1973) introduced "evolution strategies" (*Evolutionsstrategie* in the original German), a method he used to optimize real-valued parameters for devices such as airfoils. This idea was further developed by Schwefel (1975, 1977). The field of evolution strategies has remained an active area of research, mostly developing independently from the field of genetic algorithms (although recently the two communities have begun to interact). (For a short review of evolution strategies, see Back, Hoffmeister, and Schwefel 1991.) Fogel, Owens, and Walsh (1966) developed "evolutionary programming," a technique in

Chapter 1: Genetic Algorithms: An Overview

which candidate solutions to given tasks were represented as finite-state machines, which were evolved by randomly mutating their state-transition diagrams and selecting the fittest. A somewhat broader formulation of evolutionary programming also remains an area of active research (see, for example, Fogel and Atmar 1993). Together, evolution strategies, evolutionary programming, and genetic algorithms form the backbone of the field of evolutionary computation.

Several other people working in the 1950s and the 1960s developed evolution-inspired algorithms for optimization and machine learning. Box (1957), Friedman (1959), Bledsoe (1961), Bremermann (1962), and Reed, Toombs, and Baricelli (1967) all worked in this area, though their work has been given little or none of the kind of attention or followup that evolution strategies, evolutionary programming, and genetic algorithms have seen. In addition, a number of evolutionary biologists used computers to simulate evolution for the purpose of controlled experiments (see, e.g., Baricelli 1957, 1962; Fraser 1957 a,b; Martin and Cockerham 1960). Evolutionary computation was definitely in the air in the formative days of the electronic computer.

Genetic algorithms (GAs) were invented by John Holland in the 1960s and were developed by Holland and his students and colleagues at the University of Michigan in the 1960s and the 1970s. In contrast with evolution strategies and evolutionary programming, Holland's original goal was not to design algorithms to solve specific problems, but rather to formally study the phenomenon of adaptation as it occurs in nature and to develop ways in which the mechanisms of natural adaptation might be imported into computer systems. Holland's 1975 book *Adaptation in Natural and Artificial Systems* presented the genetic algorithm as an abstraction of biological evolution and gave a theoretical framework for adaptation under the GA. Holland's GA is a method for moving from one population of "chromosomes" (e.g., strings of ones and zeros, or "bits") to a new population by using a kind of "natural selection" together with the genetics-inspired operators of crossover, mutation, and inversion. Each chromosome consists of "genes" (e.g., bits), each gene being an instance of a particular "allele" (e.g., 0 or 1). The selection operator chooses those chromosomes in the population that will be allowed to reproduce, and on average the fitter chromosomes produce more offspring than the less fit ones. Crossover exchanges subparts of two chromosomes, roughly mimicking biological recombination between two single-chromosome ("haploid") organisms; mutation randomly changes the allele values of some locations in the chromosome; and inversion reverses the order of a contiguous section of the chromosome, thus rearranging the order in which genes are arrayed. (Here, as in most of the GA literature, "crossover" and "recombination" will mean the same thing.)

Holland's introduction of a population-based algorithm with crossover, inversion, and mutation was a major innovation. (Rechenberg's evolution strategies started with a "population" of two individuals, one parent and one offspring, the offspring being a mutated version of the parent; many-individual populations and crossover were not incorporated until later. Fogel, Owens, and Walsh's evolutionary programming likewise used only mutation to provide variation.) Moreover, Holland was the first to attempt to put computational evolution on a firm theoretical footing (see Holland 1975). Until recently this theoretical foundation, based on the notion of "schemas," was the basis of almost all subsequent theoretical work on genetic algorithms

In the last several years there has been widespread interaction among researchers studying various evolutionary computation methods, and the boundaries between GAs, evolution strategies, evolutionary programming, and other evolutionary approaches have broken down to some extent. Today, researchers often use the term "genetic algorithm" to describe something very far from Holland's original conception. In this book I adopt this flexibility. Most of the projects I will describe here were referred to by their originators as GAs; some were not, but they all have enough of a "family resemblance" that I include them under the rubric of genetic algorithms.

1.2 THE APPEAL OF EVOLUTION

Why use evolution as an inspiration for solving computational problems? To evolutionary-computation researchers, the mechanisms of evolution seem well suited for some of the most pressing computational problems in many fields. Many computational problems require searching through a huge number of possibilities for solutions. One example is the problem of computational protein engineering, in which an algorithm is sought that will search among the vast number of possible amino acid sequences for a protein with specified properties. Another example is searching for a set of rules or equations that will predict the ups and downs of a financial market, such as that for foreign currency. Such search problems can often benefit from an effective use of parallelism, in which many different possibilities are explored simultaneously in an efficient way. For example, in searching for proteins with specified properties, rather than evaluate one amino acid sequence at a time it would be much faster to evaluate many simultaneously. What is needed is both computational parallelism (i.e., many processors evaluating sequences at the same time) and an intelligent strategy for choosing the next set of sequences to evaluate.

Many computational problems require a computer program to be *adaptive*—to continue to perform well in a changing environment. This is typified by problems in robot control in which a robot has to perform a task in a variable environment, and by computer interfaces that must adapt to the idiosyncrasies of different users. Other problems require computer programs to be innovative—to construct something truly new and original, such as a new algorithm for accomplishing a computational task or even a new scientific discovery. Finally, many computational problems require complex solutions that are difficult to program by hand. A striking example is the problem of creating artificial intelligence. Early on, AI practitioners believed that it would be straightforward to encode the rules that would confer intelligence on a program; expert systems were one result of this early optimism. Nowadays, many AI researchers believe that the "rules" underlying intelligence are too complex for scientists to encode by hand in a "top-down" fashion. Instead they believe that the best route to artificial intelligence is through a "bottom-up" paradigm in which humans write only very simple rules, and complex behaviors such as intelligence emerge from the massively parallel application and interaction of these simple rules. Connectionism (i.e., the study of computer programs inspired by neural systems) is one example of this philosophy (see Smolensky 1988); evolutionary computation is another. In connectionism the rules are typically simple "neural" thresholding, activation spreading, and strengthening or weakening of connections; the hoped-for emergent behavior is sophisticated pattern recognition and learning. In evolutionary computation the rules are typically "natural selection" with variation due to crossover and/or mutation; the hoped-for emergent behavior is the design of high-quality solutions to difficult problems and the ability to adapt these solutions in the face of a changing environment.

Biological evolution is an appealing source of inspiration for addressing these problems. Evolution is, in effect, a method of searching among an enormous number of possibilities for "solutions." In biology the enormous set of possibilities is the set of possible genetic sequences, and the desired "solutions" are highly fit organisms—organisms well able to survive and reproduce in their environments. Evolution can also be seen as a method for *designing* innovative solutions to complex problems. For example, the mammalian immune system is a marvelous evolved solution to the problem of germs invading the body. Seen in this light, the mechanisms of evolution can inspire computational search methods. Of course the fitness of a biological organism depends on many factors—for example, how well it can weather the physical characteristics of its environment and how well it can compete with or cooperate with the other organisms around it. The fitness criteria continually change as creatures evolve, so evolution is searching a constantly changing set of possibilities. Searching for solutions in the face of changing conditions is precisely what is required for adaptive computer programs. Furthermore, evolution is a massively parallel search method: rather than work on one species at a time, evolution tests and changes millions of species in parallel. Finally, viewed from a high level the "rules" of evolution are remarkably simple: species evolve by means of random variation (via mutation, recombination, and other operators), followed by natural selection in which the fittest tend to

survive and reproduce, thus propagating their genetic material to future generations. Yet these simple rules are thought to be responsible, in large part, for the extraordinary variety and complexity we see in the biosphere.

1.3 BIOLOGICAL TERMINOLOGY

At this point it is useful to formally introduce some of the biological terminology that will be used throughout the book. In the context of genetic algorithms, these biological terms are used in the spirit of analogy with real biology, though the entities they refer to are much simpler than the real biological ones.

All living organisms consist of cells, and each cell contains the same set of one or more *chromosomes*—strings of DNA—that serve as a "blueprint" for the organism. A chromosome can be conceptually divided into *genes*—each of which encodes a particular protein. Very roughly, one can think of a gene as encoding a *trait*, such as eye color. The different possible "settings" for a trait (e.g., blue, brown, hazel) are called *alleles*. Each gene is located at a particular *locus* (position) on the chromosome.

Many organisms have multiple chromosomes in each cell. The complete collection of genetic material (all chromosomes taken together) is called the organism's *genome*. The term *genotype* refers to the particular set of genes contained in a genome. Two individuals that have identical genomes are said to have the same genotype. The genotype gives rise, under fetal and later development, to the organism's *phenotype*—its physical and mental characteristics, such as eye color, height, brain size, and intelligence.

Organisms whose chromosomes are arrayed in pairs are called *diploid*; organisms whose chromosomes are unpaired are called *haploid*. In nature, most sexually reproducing species are diploid, including human beings, who each have 23 pairs of chromosomes in each somatic (non-germ) cell in the body. During sexual reproduction, *recombination* (or *crossover*) occurs: in each parent, genes are exchanged between each pair of chromosomes to form a *gamete* (a single chromosome), and then gametes from the two parents pair up to create a full set of diploid chromosomes. In haploid sexual reproduction, genes are exchanged between the two parents' single-strand chromosomes. Offspring are subject to *mutation*, in which single nucleotides (elementary bits of DNA) are changed from parent to offspring, the changes often resulting from copying errors. The *fitness* of an organism is typically defined as the probability that the organism will live to reproduce (*viability*) or as a function of the number of offspring the organism has (*fertility*).

In genetic algorithms, the term *chromosome* typically refers to a candidate solution to a problem, often encoded as a bit string. The "genes" are either single bits or short blocks of adjacent bits that encode a particular element of the candidate solution (e.g., in the context of multiparameter function optimization the bits encoding a particular parameter might be considered to be a gene). An allele in a bit string is either 0 or 1; for larger alphabets more alleles are possible at each locus. Crossover typically consists of exchanging genetic material between two single-chromosome haploid parents. Mutation consists of flipping the bit at a randomly chosen locus (or, for larger alphabets, replacing a the symbol at a randomly chosen locus with a randomly chosen new symbol).

Most applications of genetic algorithms employ haploid individuals, particularly, single-chromosome individuals. The genotype of an individual in a GA using bit strings is simply the configuration of bits in that individual's chromosome. Often there is no notion of "phenotype" in the context of GAs, although more recently many workers have experimented with GAs in which there is both a genotypic level and a phenotypic level (e.g., the bit-string encoding of a neural network and the neural network itself).

1.4 SEARCH SPACES AND FITNESS LANDSCAPES

The idea of searching among a collection of candidate solutions for a desired solution is so common in computer science that it has been given its own name: searching in a "search space." Here the term "search space" refers to some collection of candidate solutions to a problem and some notion of "distance" between candidate solutions. For an example, let us take one of the most important problems in computational bioengineering: the aforementioned problem of computational protein design. Suppose you want use a computer to search for a protein—a sequence of amino acids—that folds up to a particular three-dimensional shape so it can be used, say, to fight a specific virus. The search space is the collection of all possible protein sequences—an infinite set of possibilities. To constrain it, let us restrict the search to all possible sequences of length 100 or less—still a huge search space, since there are 20 possible amino acids at each position in the sequence. (How many possible sequences are there?) If we represent the 20 amino acids by letters of the alphabet, candidate solutions will look like this:

A G G M C G B L...

We will define the distance between two sequences as the number of positions in which the letters at corresponding positions differ. For example, the distance between A G G M C G B L and M G G M C G B L is 1, and the distance between A G G M C G B L and L B M P A F G A is 8. An algorithm for searching this space is a method for choosing which candidate solutions to test at each stage of the search. In most cases the next candidate solution(s) to be tested will depend on the results of testing previous sequences; most useful algorithms assume that there will be some correlation between the quality of "neighboring" candidate solutions—those close in the space. Genetic algorithms assume that high-quality "parent" candidate solutions from different regions in the space can be combined via crossover to, on occasion, produce high-quality "offspring" candidate solutions.

Another important concept is that of "fitness landscape." Originally defined by the biologist Sewell Wright (1931) in the context of population genetics, a fitness landscape is a representation of the space of all possible genotypes along with their fitnesses.

Suppose, for the sake of simplicity, that each genotype is a bit string of length l , and that the distance between two genotypes is their "Hamming distance"—the number of locations at which corresponding bits differ. Also suppose that each genotype can be assigned a real-valued fitness. A fitness landscape can be pictured as an $(l + 1)$ -dimensional plot in which each genotype is a point in l dimensions and its fitness is plotted along the $(l + 1)$ st axis. A simple landscape for $l = 2$ is shown in [figure 1.1](#). Such plots are called landscapes because the plot of fitness values can form "hills," "peaks," "valleys," and other features analogous to those of physical landscapes. Under Wright's formulation, evolution causes populations to move along landscapes in particular ways, and "adaptation" can be seen as the movement toward local peaks. (A "local peak," or "local optimum," is not necessarily the highest point in the landscape, but any small

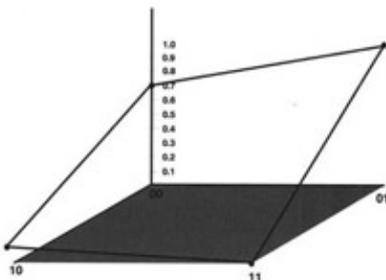


Figure 1.1: A simple fitness landscape for $l = 2$. Here $f(00) = 0.7$, $f(01) = 1.0$, $f(10) = 0.1$, and $f(11) = 0.0$.