



International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236 807 342
Fax: +43 2236 71313
E-mail: publications@iiasa.ac.at
Web: www.iiasa.ac.at

Interim Report

IR-01-058

Correlation Analysis of Fitness Landscapes

Hannelore Brandt (brandt@iiasa.ac.at)

Approved by

Ulf Dieckmann (dieckman@iiasa.ac.at)
Project Coordinator, Adaptive Dynamics Network

November 2001

Interim Reports on work of the International Institute for Applied Systems Analysis receive only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.



The Adaptive Dynamics Network at IIASA fosters the development of new mathematical and conceptual techniques for understanding the evolution of complex adaptive systems.

Focusing on these long-term implications of adaptive processes in systems of limited growth, the Adaptive Dynamics Network brings together scientists and institutions from around the world with IIASA acting as the central node.

Scientific progress within the network is reported in the IIASA Studies in Adaptive Dynamics series.

THE ADAPTIVE DYNAMICS NETWORK

The pivotal role of evolutionary theory in life sciences derives from its capability to provide causal explanations for phenomena that are highly improbable in the physico-chemical sense. Yet, until recently, many facts in biology could not be accounted for in the light of evolution. Just as physicists for a long time ignored the presence of chaos, these phenomena were basically not perceived by biologists.

Two examples illustrate this assertion. Although Darwin's publication of "The Origin of Species" sparked off the whole evolutionary revolution, oddly enough, the population genetic framework underlying the modern synthesis holds no clues to speciation events. A second illustration is the more recently appreciated issue of jump increases in biological complexity that result from the aggregation of individuals into mutualistic wholes.

These and many more problems possess a common source: the interactions of individuals are bound to change the environments these individuals live in. By closing the feedback loop in the evolutionary explanation, a new mathematical theory of the evolution of complex adaptive systems arises. It is this general theoretical option that lies at the core of the emerging field of adaptive dynamics. In consequence a major promise of adaptive dynamics studies is to elucidate the long-term effects of the interactions between ecological and evolutionary processes.

A commitment to interfacing the theory with empirical applications is necessary both for validation and for management problems. For example, empirical evidence indicates that to control pests and diseases or to achieve sustainable harvesting of renewable resources evolutionary deliberation is already crucial on the time scale of two decades.

The Adaptive Dynamics Network has as its primary objective the development of mathematical tools for the analysis of adaptive systems inside and outside the biological realm.

IIASA STUDIES IN ADAPTIVE DYNAMICS

- No. 1 Metz JAJ, Geritz SAH, Meszéna G, Jacobs FJA, van Heerwaarden JS:
Adaptive Dynamics: A Geometrical Study of the Consequences of Nearly Faithful Reproduction.
IIASA Working Paper WP-95-099.
In: van Strien SJ, Verduyn Lunel SM (eds.): *Stochastic and Spatial Structures of Dynamical Systems*, Proceedings of the Royal Dutch Academy of Science (KNAW Verhandelingen), North Holland, Amsterdam, pp. 183–231 (1996).
- No. 2 Dieckmann U, Law R:
The Dynamical Theory of Coevolution: A Derivation from Stochastic Ecological Processes.
IIASA Working Paper WP-96-001.
Journal of Mathematical Biology (1996) 34, 579–612.
- No. 3 Dieckmann U, Marrow P, Law R:
Evolutionary Cycling of Predator-Prey Interactions: Population Dynamics and the Red Queen.
IIASA Preprint (1995).
Journal of Theoretical Biology (1995) 176, 91–102.
- No. 4 Marrow P, Dieckmann U, Law R:
Evolutionary Dynamics of Predator-Prey Systems: An Ecological Perspective.
IIASA Working Paper WP-96-002.
Journal of Mathematical Biology (1996) 34, 556–578.
- No. 5 Law R, Marrow P, Dieckmann U:
On Evolution under Asymmetric Competition.
IIASA Working Paper WP-96-003.
Evolutionary Ecology (1997) 11, 485–501.
- No. 6 Metz JAJ, Mylius SD, Diekmann O:
When Does Evolution Optimise? On the Relation between Types of Density Dependence and Evolutionarily Stable Life History Parameters.
IIASA Working Paper WP-96-004.
- No. 7 Ferrière R, Gatto M:
Lyapunov Exponents and the Mathematics of Invasion in Oscillatory or Chaotic Populations.
Theoretical Population Biology (1995) 48, 126–171.

- No. 8 Ferrière R, Fox GA:
Chaos and Evolution.
IIASA Preprint (1996).
Trends in Ecology and Evolution (1995) 10, 480–485.
- No. 9 Ferrière R, Michod RE:
The Evolution of Cooperation in Spatially Heterogeneous Populations.
IIASA Working Paper WP-96-029.
American Naturalist (1996) 147, 692–717.
- No. 10 Van Dooren TJM, Metz JAJ:
Delayed Maturation in Temporally Structured Populations with Non-Equilibrium Dynamics.
IIASA Working Paper WP-96-070.
Journal of Evolutionary Biology (1998) 11, 41–62.
- No. 11 Geritz SAH, Metz JAJ, Kisdi É, Meszéna G:
The Dynamics of Adaptation and Evolutionary Branching.
IIASA Working Paper WP-96-077.
Physical Review Letters (1997) 78, 2024–2027.
- No. 12 Geritz SAH, Kisdi É, Meszéna G, Metz JAJ:
Evolutionarily Singular Strategies and the Adaptive Growth and Branching of the Evolutionary Tree.
IIASA Working Paper WP-96-114.
Evolutionary Ecology (1998) 12, 35–57.
- No. 13 Heino M, Metz JAJ, Kaitala V:
Evolution of Mixed Maturation Strategies in Semelparous Life-Histories: the Crucial Role of Dimensionality of Feedback Environment.
IIASA Working Paper WP-96-126.
Philosophical Transactions of the Royal Society of London Series B (1997) 352, 1647–1655.
- No. 14 Dieckmann U:
Can Adaptive Dynamics Invade?
IIASA Working Paper WP-96-152.
Trends in Ecology and Evolution (1997) 12, 128–131.
- No. 15 Meszéna G, Czibula I, Geritz SAH:
Adaptive Dynamics in a Two-Patch Environment: a Simple Model for Allopatric and Parapatric Speciation.
IIASA Interim Report IR-97-001.
Journal of Biological Systems (1997) 5, 265–284.

- No. 16 Heino M, Metz JAJ, Kaitala V:
The Enigma of Frequency-Dependent Selection.
IIASA Interim Report IR-97-061.
Trends in Ecology and Evolution (1998) 13, 367–370.
- No. 17 Heino M:
Management of Evolving Fish Stocks.
IIASA Interim Report IR-97-062.
Canadian Journal of Fisheries and Aquatic Sciences (1998) 55, 1971–1982.
- No. 18 Heino M:
Evolution of Mixed Reproductive Strategies in Simple Life-History Models.
IIASA Interim Report IR-97-063.
- No. 19 Geritz SAH, van der Meijden E, Metz JAJ:
Evolutionary Dynamics of Seed Size and Seedling Competitive Ability.
IIASA Interim Report IR-97-071.
Theoretical Population Biology (1999) 55, 324–343.
- No. 20 Galis F, Metz JAJ:
Why are there so many Cichlid Species? On the Interplay of Speciation and Adaptive Radiation.
IIASA Interim Report IR-97-072.
Trends in Ecology and Evolution (1998) 13, 1–2.
- No. 21 Boerlijst MC, Nowak MA, Sigmund K:
Equal Pay for all Prisoners. / The Logic of Contrition.
IIASA Interim Report IR-97-073.
American Mathematical Society Monthly (1997) 104, 303–307.
Journal of Theoretical Biology (1997) 185, 281–294.
- No. 22 Law R, Dieckmann U:
Symbiosis without Mutualism and the Merger of Lineages in Evolution.
IIASA Interim Report IR-97-074.
Proceedings of the Royal Society of London Series B (1998) 265, 1245–1253.
- No. 23 Klinkhamer PGL, de Jong TJ, Metz JAJ:
Sex and Size in Cosexual Plants.
IIASA Interim Report IR-97-078.
Trends in Ecology and Evolution (1997) 12, 260–265.
- No. 24 Fontana W, Schuster P:
Shaping Space: The Possible and the Attainable in RNA Genotype-Phenotype Mapping.
IIASA Interim Report IR-98-004.
Journal of Theoretical Biology (1998) 194, 491–515.

- No. 25 Kisdi É, Geritz SAH:
Adaptive Dynamics in Allele Space: Evolution of Genetic Polymorphism by Small Mutations in a Heterogeneous Environment.
IIASA Interim Report IR-98-038.
Evolution (1999) 53, 993-1008.
- No. 26 Fontana W, Schuster P:
Continuity in Evolution: On the Nature of Transitions.
IIASA Interim Report IR-98-039.
Science (1998) 280, 1451–1455.
- No. 27 Nowak MA, Sigmund K:
Evolution of Indirect Reciprocity by Image Scoring. / The Dynamics of Indirect Reciprocity.
IIASA Interim Report IR-98-040.
Nature (1998) 393, 573–577.
Journal of Theoretical Biology (1998) 194, 561-574.
- No. 28 Kisdi É:
Evolutionary Branching Under Asymmetric Competition.
IIASA Interim Report IR-98-045.
Journal of Theoretical Biology (1999) 197, 149-162.
- No. 29 Berger U:
Best Response Adaptation for Role Games.
IIASA Interim Report IR-98-086.
- No. 30 Van Dooren TJM:
The Evolutionary Ecology of Dominance-Recessivity.
IIASA Interim Report IR-98-096.
Journal of Theoretical Biology (1999) 198, 519-532.
- No. 31 Dieckmann U, O’Hara B, Weisser W:
The Evolutionary Ecology of Dispersal.
IIASA Interim Report IR-98-108.
Trends in Ecology and Evolution (1999) 14, 88–90.
- No. 32 Sigmund K:
Complex Adaptive Systems and the Evolution of Reciprocation.
IIASA Interim Report IR-98-100.
Ecosystems (1998) 1, 444-448.
- No. 33 Posch M, Pichler A, Sigmund K:
The Efficiency of Adapting Aspiration Levels.
IIASA Interim Report IR-98-103.
Proceedings of the Royal Society of London Series B (1999) 266, 1427-1435.

- No. 34 Mathias A, Kisdi É:
Evolutionary Branching and Coexistence of Germination Strategies.
IIASA Interim Report IR-99-014.
- No. 35 Dieckmann U, Doebeli M:
On the Origin of Species by Sympatric Speciation.
IIASA Interim Report IR-99-013.
Nature (1999) 400, 354–357.
- No. 36 Metz JAJ, Gyllenberg M:
How Should We Define Fitness in Structured Metapopulation Models? Including an Application to the Calculation of Evolutionarily Stable Dispersal Strategies.
IIASA Interim Report IR-99-019.
Proceedings of the Royal Society of London Series B (2001) 268, 499–508.
- No. 37 Gyllenberg M, Metz JAJ:
On Fitness in Structured Metapopulations.
IIASA Interim Report IR-99-037.
Research Report A38 (1999), University of Turku, Institute of Applied Mathematics, Turku, Finland.
- No. 38 Meszéna G, Metz JAJ:
Species Diversity and Population Regulation: The Importance of Environmental Feedback Dimensionality.
IIASA Interim Report IR-99-045.
- No. 39 Kisdi É, Geritz SAH:
Evolutionary Branching and Sympatric Speciation in Diploid Populations.
IIASA Interim Report IR-99-048.
- No. 40 Ylikarjula J, Heino M, Dieckmann U:
Ecology and Adaptation of Stunted Growth in Fish.
IIASA Interim Report IR-99-050.
Evolutionary Ecology (1999) 13, 433–453.
- No. 41 Nowak MA, Sigmund K:
Games on Grids.
IIASA Interim Report IR-99-038.
In: Dieckmann U, Law R, Metz JAJ (eds.): The Geometry of Ecological Interactions: Simplifying Spatial Complexity, Cambridge University Press, Cambridge, UK, pp. 135–150 (2000).
- No. 42 Ferrière R, Michod RE:
Wave Patterns in Spatial Games and the Evolution of Cooperation.
IIASA Interim Report IR-99-041.
In: Dieckmann U, Law R, Metz JAJ (eds.): The Geometry of Ecological Interactions: Simplifying Spatial Complexity, Cambridge University Press, Cambridge, UK, pp. 318–332 (2000).

- No. 43 Kisdi É, Jacobs FJA, Geritz SAH:
Red Queen Evolution by Cycles of Evolutionary Branching and Extinction.
IIASA Interim Report IR-00-030.
- No. 44 Meszéna G, Kisdi É, Dieckmann U, Geritz SAH, Metz JAJ:
Evolutionary Optimisation Models and Matrix Games in the Unified Perspective of Adaptive Dynamics.
IIASA Interim Report IR-00-039.
- No. 45 Parvinen K, Dieckmann U, Gyllenberg M, Metz JAJ:
Evolution of Dispersal in Metapopulations with Local Density Dependence and Demographic Stochasticity.
IIASA Interim Report IR-00-035.
- No. 46 Doebeli M, Dieckmann, U:
Evolutionary Branching and Sympatric Speciation Caused by Different Types of Ecological Interactions.
IIASA Interim Report IR-00-040.
The American Naturalist (2000) 156, S77–S101.
- No. 47 Heino M, Hanski I:
Evolution of Migration Rate in a Spatially Realistic Metapopulation Model.
IIASA Interim Report IR-00-044.
The American Naturalist (2001) 157, 495–511.
- No. 48 Gyllenberg M, Parvinen K, Dieckmann U:
Evolutionary Suicide and Evolution of Dispersal in Structured Metapopulations.
IIASA Interim Report IR-00-056.
- No. 49 Van Dooren TJM:
The Evolutionary Dynamics of Direct Phenotypic Overdominance: Emergence Possible, Loss Probable.
IIASA Interim Report IR-00-048.
Evolution (2000) 54, 1899–1914.
- No. 50 Nowak MA, Page KM, Sigmund K:
Fairness Versus Reason in the Ultimatum Game.
IIASA Interim Report IR-00-057.
Science (2000) 289, 1773–1775.
- No. 51 De Feo O, Ferrière R:
Bifurcation Analysis of Population Invasion: On-Off Intermittency and Basin Riddling.
IIASA Interim Report IR-00-074.
International Journal of Bifurcation and Chaos (2000) 10, 443–452.

- No. 52 Laaka-Lindberg S, Heino M:
Clonal Dynamics and Evolution of Dormancy in the leafy hepatic Lophozia silvicola.
IIASA Interim Report IR-01-018.
Oikos (2001) 94, 525–532.
- No. 53 Sigmund K, Hauert C, Nowak MA:
Reward and Punishment in Minigames.
IIASA Interim Report IR-01-031.
Proceedings of the National Academy of Sciences of the USA (2001) 98, 10757–10762.
- No. 54 Hauert C, De Monte S, Sigmund K, Hofbauer J:
Oscillations in Optional Public Good Games.
IIASA Interim Report IR-01-036.
- No. 55 Ferrière R, Le Galliard JF:
Invasion Fitness and Adaptive Dynamics in Spatial Population Models.
IIASA Interim Report IR-01-043.
- No. 56 de Mazancourt C, Loreau M, Dieckmann U:
Can the Evolution of Plant Defense Lead to Plant-Herbivore Mutualism?
IIASA Interim Report IR-01-053.
The American Naturalist (2001) 158, 109-123.
- No. 57 Claessen D, Dieckmann U:
Ontogenetic Niche Shifts and Evolutionary Branching in Size-structured Populations
IIASA Interim Report IR-01-056.
- No. 58 Brandt H:
Correlation Analysis of Fitness Landscapes
IIASA Interim Report IR-01-058.

Issues of the IIASA Studies in Adaptive Dynamics series can be obtained free of charge.
Please contact:

Adaptive Dynamics Network
International Institute for Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg
Austria

Telephone +43 2236 807, Telefax +43 2236 71313, E-Mail adn@iiasa.ac.at,
Internet <http://www.iiasa.ac.at/Research/ADN>

Abstract

The notion of a fitness landscape has permeated the analysis of evolutionary processes for more than 60 years. Introduced by Sewall Wright for discussing biological evolution and speciation, the concept has recently been transferred to the study of abstract genotypes of various evolutionary algorithms. The features of high-dimensional fitness landscapes can vary to a high degree, and the question by which means they can be described has turned out to be a challenging problem. Even though some statistics have been suggested for this purpose and are already well-analyzed, presently discussed statistics do not seem appropriate for obtaining sufficiently accurate predictions of evolutionary dynamics at the level of fitness.

In this study, an analysis of three different types of fitness landscapes is presented. I introduce a new correlation measure, and show by comparing the actual evolutionary waiting times to those predicted when only taking into account the correlation statistics, that these statistics seem to capture salient information of the underlying fitness landscapes.

Based on one-dimensional correlation statistics, very accurate predictions of evolutionary waiting times are achieved for the fitness landscape of the Travelling Salesman Problem and NKp landscapes with low degree of neutrality. Both for NKp landscapes with high neutrality and RSF landscapes, which in a similar way involve large-scale neutrality, higher-dimensional correlation statistics provide enough information to estimate evolutionary waiting times. Finally I present an approach towards analytic descriptions of evolutionary dynamics for the analyzed fitness landscapes of low neutrality.

About the Author

Hannelore Brandt
Institute of Mathematics
Strudlhofgasse 4
A-1090 Vienna, Austria
and
Adaptive Dynamics Network
International Institute for Applied Systems Analysis
A-2361 Laxenburg, Austria

Acknowledgement

The author would like to thank Ulf Dieckmann for excellent support of the work and for the possibility of collaborating with the Adaptive Dynamics Network at the International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria.

Contents

1	Introduction	3
1.1	Simple evolutionary models	3
1.2	Evolutionary algorithms	4
1.3	Fitness landscapes	7
1.4	Correlation functions	9
1.5	Structure of this work	11
2	One-dimensional correlation	13
2.1	TSP fitness landscapes	14
2.1.1	General information	14
2.1.2	An evolutionary algorithm using point mutation	15
	Without correlation: Percolation	21
	Monomorphic correlation	24
	Polymorphic correlation	29
2.1.3	Other mutation operators	31
	Reverse mutation	31
	Remove-and-reinsert mutation	34
2.1.4	A second TSP and a variant evolutionary algorithm	36
2.2	NKp fitness landscapes	46
2.2.1	Low neutrality	49
2.2.2	High neutrality	53
2.3	The Royal Staircase Fitness function	56
2.3.1	The fitness function	57
2.3.2	The evolutionary algorithm	59
3	Higher-dimensional correlation	65
3.1	NKp fitness landscapes	65
3.1.1	Two-dimensional correlation	65

<i>CONTENTS</i>	2
3.1.2 Three-dimensional correlation	69
3.2 The Royal Staircase Fitness function	73
3.2.1 Two-dimensional correlation	73
3.2.2 Three-dimensional correlation	76
4 Towards analytic descriptions	82
4.1 TSP fitness landscapes	84
4.1.1 Prerequisites	84
4.1.2 Approximation of fitness distribution dynamics	85
Mutation	86
Selection	89
Recursion relations	90
4.1.3 Approximation of maximal fitness	91
4.1.4 Testing the predictions	91
4.2 NKp fitness landscapes	93
4.2.1 Testing the predictions	94
4.2.2 Possible improvements and extensions	96
5 Summary	98

Chapter 1

Introduction

1.1 Simple evolutionary models

Since Darwin's theory of evolution, mathematical models have become indispensable for understanding evolutionary processes. Even though individuals are recognized as the fundamental level of evolution, competition is also observed among genes within a genome, groups of organisms and species, and in that sense, evolution simultaneously takes place on various levels. In order to study evolutionary phenomena, mathematical models often describe dynamics directly at the level which is of immediate interest. Unfortunately, evolutionary dynamics at a single level cannot be exhaustively understood without accounting for the interactions with other levels. These interactions can result in very complex dynamics which are impossible to be described by straightforward models.

Nevertheless it is feasible to define evolutionary dynamics at some lower level and to deduce the dynamics at a higher level of interest. In this thesis I demonstrate applications of this methodology to some simple and artificial evolutionary systems for which the dynamics are precisely defined at the level of genotypes. Of interest for us, however, are only the resulting dynamics at the level of fitness. The benefit of such an approach is that dynamics at the level of interest do not have to be stipulated but are properly founded on some lower level evolution.

In particular, abstract individuals are assumed to have fitness values which are only determined by their genotypes. Most commonly, fitness is identified with the probability for reproductive success an individual has.

The assumption that genotypes alone determine fitness values is a rough simplification of the complex process which actually determines the reproductive success. Every individual, although it can to some extent be identified with its genotype, undergoes a complicated process of development, involving interactions with other individuals and the environment, which is in turn influenced by the individuals that live therein. As all these complications are not studied in this work, our models can be mainly compared with specific evolutionary systems like experiments with bacteria or replicating biomolecules. The dynamics observed in these systems are nonetheless very elaborate and provide a basis for understanding more involved evolutionary scenarios.

The evolutionary systems studied in this work are generally called evolutionary algorithms and are simulated on computers. Genotypes are represented by vectors and a function maps each genotype to a certain fitness value. Starting with a population of such individuals, they produce offspring having genotypes different from their parent individuals: mutation and crossing-over operators transform genotypes according to specific stochastic schemes. These offspring individuals therefore possibly also have fitness values different from those in the old population. A selection operator then selects the next generation of individuals out of the offspring population, with fitter individuals having a higher chance to be chosen. The individuals of the new generation then reproduce again. As fitness is only depending on the genotypes of individuals, this procedure amounts to a successive optimization of fitness values present in a population. It is therefore common to take objective functions of optimization tasks as fitness functions. Apart from functions of this type I also choose fitness maps which are specifically designed to represent commonly observed properties of evolutionary systems, like neutrality in the genotype-to-fitness map.

1.2 Evolutionary algorithms

An optimization task is the search for the solution of a problem of the form

$$\max_{x \in G} f(x) \quad \text{with } f : X \rightarrow Z, G \subseteq X$$

where X and Z are topological spaces with Z having a total order. The function f is called the objective function and G is the feasible region. The

elements $x \in G$ are termed feasible points. All $\hat{x} \in G$ with

$$f(\hat{x}) \geq f(x) \quad \forall x \in G$$

are called optimal solution or solution of the problem. A feasible point \bar{x} is called a local optimum if there exists a neighborhood U of \bar{x} in G such that

$$f(\bar{x}) \geq f(x) \quad \forall x \in U$$

This definition is not restricted to problems of maximizing a given function, as for all problems of minimization there exists an equivalent optimization task in the above sense. The complexity of optimization tasks is in general very high. Most optimization problems except local optimization tasks belong to the class of NP-hard problems for which no algorithm with polynomial solution time exists.

However, for many problems of practical interest the focus is not on detecting the optimal solution but on finding a feasible point sufficiently close to it within short computation time. For this purpose it is convenient to let candidate solutions simply evolve towards better ones. Such algorithms, which make use of basic principles of evolution like mutation and selection (sometimes crossover is also considered), are called *evolutionary algorithms* and have been established as efficient tools for finding quasi-optimal solutions of many optimization problems (Michalewicz 1992, Beasley 1997).

In these algorithms, feasible points of the optimization task are seen as genotypes of abstract individuals. Most problems of optimization allow for a binary or real number vector representation of these genotypes. Each individual has a certain fitness determined by the objective function at the considered feasible point. Starting with a whole population of such individuals, stochastic mutation and crossover operators are used to produce an offspring population different from the parent generation: When a genotype is mutated, random entries of the genotype vector are modified. At crossing-over, parts of the genotypes of two individuals are interchanged. Both types of operators produce new individuals, which can be regarded as offspring, having genotypes and thus possibly also fitness values different from those of their parent individuals. Out of the offspring population, a new generation of individuals is selected according to a specific scheme biased towards fitter individuals. Such a selection scheme can either be deterministic, e.g. the fitter half of all individuals is taken over to the next generation, or stochastic, when fitter individuals are assigned a greater probability to be taken over.

Computation time to find individuals of certain fitness can be reduced a lot when using clever mutation, crossing-over and selection operators specifically designed for a problem. But regardless of the complexity of operators, evolutionary algorithms succeed in finding fitter and fitter individuals and thus points close to the optimal solutions.

The particular interest in evolutionary algorithms derives from the fact that the principle of evolving genotypes can be applied to nearly any optimization task. Problems of operations research like the Travelling Salesman Problem (TSP), scheduling or transportation problems, as well as designing neural networks (Pasemann *et al.* 1999), problems on graphs, prediction of RNA molecule folding (Schuster *et al.* 1994, Schuster 1997, Fontana *et al.* 1993) and many other tasks have successfully been tackled.

In this study the following concept of an evolutionary algorithm is used:

1. Generate a population of P random genotypes.
2. Create offspring by mutating the genotypes of individuals.
3. Evaluate the fitness of all offspring and stop if one reaches a certain final value.
4. Select according to a selection scheme a new generation of P individuals.

Go to step 2.

This algorithm does not include crossing-over as a second genetic operator besides mutation. Although crossing-over is in some cases decisive for the evolutionary search, many studies of evolutionary algorithms have also demonstrated that the qualitative features of the dynamics are not necessarily influenced when only considering mutation.

Although most of the interest in evolutionary algorithms is focused on their performance as methods for solving optimization tasks, they can also be regarded as simple models for evolutionary dynamics. This is the way evolutionary algorithms are studied in this work. An analysis of the resulting dynamics will of course also lead to a better understanding of their property as search algorithms for problems of optimization.

Since individuals of an evolutionary algorithm become more and more fit over generations, we can imagine that they climb fitness peaks. For many years the metaphor of a fitness landscape has served as a powerful concept

for discussing biological evolution. In this thesis I investigate the highly complicated structures of specific fitness landscapes for an improvement of our understanding of evolutionary processes on the level of fitness.

1.3 Fitness landscapes

The notion of a fitness landscape has permeated the analysis of evolutionary processes for more than 60 years (Gavrilets 1997). Introduced by Sewall Wright (Wright 1932) for discussing biological evolution and speciation, the concept has recently been transferred to the study of abstract genotypes of various evolutionary algorithms. Arranging genotypes in a topological space with each genotype situated next to those which can be reached by a single mutation, the genotype's mutants or neighbors, and adding one dimension to include the fitness values of genotypes leads to the picture of a fitness landscape. A fitness landscape is therefore determined by a finite but usually very large set V of genotypes, a definition of neighborhood of genotypes, and a fitness function $f : V \rightarrow \mathbb{R}$. The set V can be interpreted as the vertex set of a graph, with its edge set given by the neighborhood relations. A population of individuals can then be seen as a cloud of points on the fitness landscape, with the combined effect of mutation and selection forcing the population to perform a hill-climbing process towards fitness peaks.

According to this idea, every evolutionary algorithm can be regarded as a dynamical process on a fitness landscape. Understanding the dynamics of evolutionary algorithms, which is the aim of this work, therefore amounts to understanding the structural features of the analyzed fitness landscapes. Stating the properties of fitness landscapes which are crucial for the resulting evolutionary dynamics is unfortunately not a straightforward process. The discussion by which means they can be characterized is still ongoing.

For a long time, the notion of rugged fitness landscapes, involving many local peaks separated by fitness valleys, has dominated the discussion of adaptive processes (Kauffman and Levin 1987). However, for many evolutionary processes the intuitive image of a three-dimensional landscape with its emphasis on peaks and valleys (see Figure 1) may be inappropriate. Hill-climbing on such a rugged fitness landscape, a population will soon end up at a local peak: selective pressures will prevent it from crossing the surrounding adaptive valleys to reach a higher fitness peak. Yet, biological evolution and evolutionary algorithms do not appear to become entrapped in local fitness

peaks as often as this intuitive picture suggests (Schuster 1996).

The reason why this picture of a three-dimensional landscape can be misleading is that most genotypes differ in much more than two properties, and the number of genotypes that can be reached by a single mutational step, i.e. the number of neighbors, is very large. The resulting fitness landscapes are of dimensions much higher than three. It turns out that the structural features of these high-dimensional landscapes are very different from those of low-dimensional ones. Every dimension offers a connection to neighbors possibly having higher fitness. Consequently, for many high-dimensional landscapes the problem of being stuck in a local adaptive peak far away from the global optimum might even be non-existent (Gavrilets 1997). Ruggedness should therefore not be related to the frequency of local optima in a landscape. I refer of a fitness landscape to be rugged, if a substantial range of different fitness values is found in the neighborhood of a typical genotype.

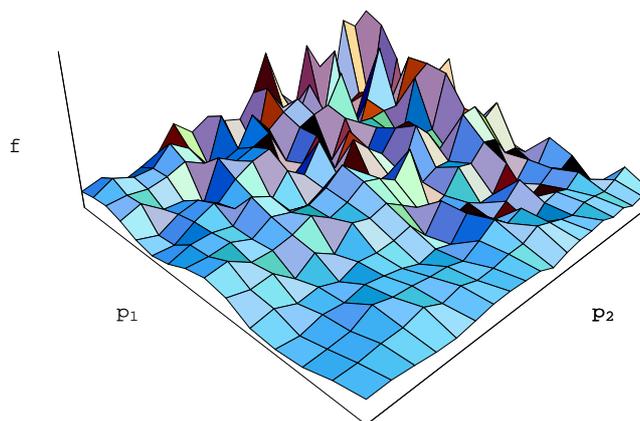


Figure 1 *If genotypes only differ in two properties (p_1, p_2) and if an appropriate mutation operator is chosen, they can be arranged in a two-dimensional space, with each individual next to those which can be reached by a single mutation. Assigning a certain fitness value f to all genotypes leads to a fitness landscape. Evolution in areas where the surface is rather smooth may result in shorter waiting times, whereas in rugged regions a population can easily end up in a local peak. The features of high-dimensional fitness landscapes, however, are certainly very different from the properties we observe for three-dimensional ones.*

From investigating special types of fitness landscapes it is known that their features can vary to a high degree. It is therefore natural to ask which

structural properties of fitness landscapes determine the evolutionary dynamics. Having identified the crucial features would allow for classifying fitness landscapes in such a way that the performance of evolutionary processes can be predicted. Even though some statistics have been suggested for this purpose and are already well-analyzed (Weinberger 1990, Stadler 1992, Stadler 1996, Barnett 1997), presently discussed statistics do not seem to be appropriate for obtaining sufficiently accurate predictions of evolutionary waiting times.

Most of the analysis of fitness landscape so far has been concentrated on the feature of ruggedness, particularly measured by two types of correlation functions.

1.4 Correlation functions

In order to describe complex and high-dimensional fitness landscapes, mainly two types of correlation functions have been studied so far. Let a fitness landscape be given by a set V of genotypes, a definition of neighborhood of genotypes, and a fitness function $f : V \rightarrow \mathbb{R}$. If we define the mean and variance of the fitness landscape,

$$\bar{f} = \frac{1}{|V|} \sum_{x \in V} f(x) \quad \text{and} \quad \sigma_f^2 = \frac{1}{|V|} \sum_{x \in V} (f(x) - \bar{f})^2 = \overline{f^2} - \bar{f}^2,$$

the (direct) correlation function $\rho(d)$ introduced by Eigen *et al.* (1989) is given by

$$\rho(d) = \frac{\langle f(x)f(y) \rangle_d - \bar{f}^2}{\sigma_f^2},$$

where $\langle f(x)f(y) \rangle_d$ denotes the mean value of the product of fitness values of all pairs of genotypes x and y , which are at a distance of d mutational steps from each other in genotype space.

Another measure of correlation proposed by Weinberger (1990), the auto-correlation function $r(s)$, is based on time series $\{f(x_0), f(x_1), \dots\}$ of fitness values along random walks on the fitness landscape:

$$r(s) = \frac{\langle f(x_t)f(x_{t+s}) \rangle - \langle f(x_t) \rangle \langle f(x_{t+s}) \rangle}{\sqrt{(\langle f(x_t)^2 \rangle - \langle f(x_t) \rangle^2) (\langle f(x_{t+s})^2 \rangle - \langle f(x_{t+s}) \rangle^2)}}$$

The mean values, denoted by $\langle \rangle$, are taken over all initial conditions x_0 and all times t . If all genotypes have the same number of neighbors, i.e. the set V interpreted as a graph is regular, the autocorrelation function simplifies to

$$r(s) = \frac{\langle f(x_t)f(x_{t+s}) \rangle - \bar{f}^2}{\sigma_f^2}$$

The direct correlation function describes how related the fitness values of two genotypes at certain mutational distances are, whereas the autocorrelation function relates the fitness values of genotypes which can be reached from each other by making a certain number of mutational steps. Both measures are closely related according to

$$r(s) = \sum_{d=0}^n P_{s,d} \rho(d),$$

where $P_{s,d}$ is the probability that a random walk of s steps ends at distance d from its starting point, and n is the maximal mutational distance between two genotypes in V (Happel and Stadler 1996). These two correlation functions thus carry equivalent information. In particular, for the correlation between neighboring genotypes we get the following:

$$\varrho = r(1) = \rho(1)$$

For a large class of landscapes, among them important landscapes of combinatorial optimization tasks like the Travelling Salesman Problem, the autocorrelation function is exponential,

$$r(s) = \varrho^s.$$

Both types of correlation functions are used to specify the degree of ruggedness of a fitness landscape. In particular, the number of local optima can be estimated using these correlation functions, but such estimations are difficult to verify if the set V of genotypes is large. Based on these two types of correlation functions, only rough approximations of evolutionary dynamics have been available. Moreover, it was shown by Barnett (1997), that the autocorrelation function is invariant on NKp fitness landscapes with varying degree of neutrality, see Section 2.2. As a measure intended to describe ruggedness, the autocorrelation and the direct correlation function appear to

be inappropriate for characterizing fitness landscapes involving both features, ruggedness and neutrality.

The correlation measures presented in this work are designed to capture the combined effect of ruggedness and neutrality. The new correlation measures enhance and reduce the information provided by the two conventional correlation functions: Whereas the focus is only on single mutational steps, the initial fitness of a genotype is taken into account. The correlation functions introduced here relate the fitness values of genotypes that can be transformed into each other by a single mutation. For every fitness value, or for a narrow range of fitness values, this information thus consists of a whole probability distribution of neighboring fitness values. I call the correlation measure one-dimensional, as it considers only one property of genotypes, i.e. their fitness. For highly neutral fitness landscapes, more properties of genotypes have to be respected, and one-dimensional correlation information is extended to higher-dimensional correlations. Based on this new type of correlation information, critical features of fitness landscapes are captured.

1.5 Structure of this work

In this study I focus on specific fitness landscapes of well-known problems and investigate the durations (or waiting times) of evolutionary processes on these landscapes. I introduce a new type of correlation statistics, different from those used so far, and show, by comparing observed and predicted waiting time distributions, that these statistics are very useful for understanding, predicting, and classifying evolutionary processes on high-dimensional fitness landscapes.

In Chapter 2 the focus is on one-dimensional correlation statistics, which are used to predict waiting times on various specific fitness landscapes of the Travelling Salesman Problem (TSP), on NK_p fitness landscapes, and on the landscape of the Royal Staircase Fitness (RSF) function. After some general information on the Travelling Salesman Problem and the utilized evolutionary algorithm, I show that a percolation approximation of the fitness landscape, which neglects the correlation between the fitness of neighboring genotypes, is not successful in predicting the evolutionary dynamics at the level of fitness values. In contrast, one-dimensional correlation information, based on a sufficiently fine-grained classification of fitness values, allows for very accurate predictions of evolutionary waiting times. In the

following, similarly encouraging results are obtained for other mutation operators for the TSP, resulting in differently structured fitness landscapes, and for NKp landscapes with low neutrality. For NKp landscapes involving high degrees of neutrality, and for the landscape of the Royal Staircase Fitness function, which in a similar way decomposes into few but very extended neutral clusters, only relatively rough estimates of evolutionary waiting times are achieved.

Chapter 3 extends the analysis to higher-dimensional correlation statistics, which capture the salient features even of landscapes with high degrees of neutrality. Both for NKp landscapes with high neutrality and RSF landscapes, the inclusion of information about genotypes that goes beyond their fitness values, is decisive for obtaining good approximations of the evolutionary algorithm. Based on three-dimensional correlation information, which describes genotypes by three crucial properties, it is possible to predict the process of neutral evolution of a population evolving on clusters of equal fitness. The conceptual simplicity of a successful approximation of the dynamics on these highly neutral fitness landscapes is surprising.

An approach towards an analytic description of evolutionary processes on fitness landscapes is presented in Chapter 4. The structure of one-dimensional correlation information allows for an approximation based on normal distributions. By applying a stochastic selection scheme based on an exponential selection function, recursion equations for the change of mean and variance of fitness in large populations evolving over generations are derived. These yield good estimations of the actual dynamics.

A summary of my findings and a sketch of open questions resulting from this study is provided in Chapter 5.

Chapter 2

One-dimensional correlation

In this chapter I study several specific evolutionary algorithms and show that one-dimensional correlation statistics of the analyzed fitness landscapes allow for a very accurate description of the dynamics on the level of fitness. This is proved by comparing actual evolutionary waiting times to those predicted by such statistics.

The dynamics of an evolutionary algorithm which are relevant for us, are the dynamics on the level of fitness. The underlying highly complicated dynamics on the genotype level are not of direct interest. Thinking of a population on a fitness landscape we want to know how the next mutational and selective step changes the fitness distribution in the population. When the genotype of an individual is mutated, the new genotype is a neighbor of the old one in genotype space. After mutating all genotypes, the selection procedure, which is only dependent on the fitness values of individuals, is a straightforward process. Consequently, the neighborhood structure of all individuals plays the key role of evolutionary dynamics. For a condensed description of this structure, individuals having a similar fitness and a similar neighborhood in terms of fitness values are thought to belong to the same class of individuals. For each class we determine a mean fitness value and a mean distribution of all the classes present in the neighborhood. If the number of different classes is much lower than the number of genotypes, we can approximate the complicated evolutionary process on the fitness landscape by a much simpler process based on transitions between the classes. The power of such an approximation of an evolutionary algorithm depends on the number of classes necessary to describe the neighborhood structure of individuals with sufficient accuracy.

In the simplest case, the neighborhood for individuals of certain fitness ranges is very similar. The genotypes are then divided into several fitness classes and a population evolving on the fitness landscape can be approximated by means of the transition probabilities between fitness classes under mutational steps. We call this correlation between fitness classes *one-dimensional correlation*, as it concerns the correlation within one property of genotypes.

In the most difficult case, each genotype has a very specific neighborhood structure, and number of classes necessary to describe the possible neighborhood structures is as big as the genotype space itself. Important is the question of which case fitness landscapes of actual interest are. We therefore turn our attention to fitness landscapes which have received much attention during the last years.

In the following, the fitness landscape of a specific optimization task, the Travelling Salesman Problem, will be analyzed. I demonstrate that a characterization of the fitness landscape which neglects the correlation of neighboring fitness values is not successful, and show that a sufficiently accurate correlation-based description captures the salient features of the landscape.

2.1 TSP fitness landscapes

2.1.1 General information

A salesman who has to visit each city on a given list, knowing the distances between all pairs of cities, will try to minimize the length of his tour. This optimization task is called *The Travelling Salesman Problem* (TSP) and has received attention for many decades. The TSP belongs to the class of NP-hard optimization problems and the solution time therefore grows more than exponentially with the number of cities.

In many cases of actual interest, however, the focus is not on detecting the shortest possible tour, but on finding a tour that is sufficiently close to the optimum within feasible computation time. For this purpose it is convenient to implement an evolutionary algorithm and let candidate solutions evolve towards better ones.

In this study the landscape of the TSP is chosen as a benchmark problem because of its canonical genotype-to-fitness map and the attention that it has received in recent studies of fitness landscapes (Stadler and Schnabl 1992,

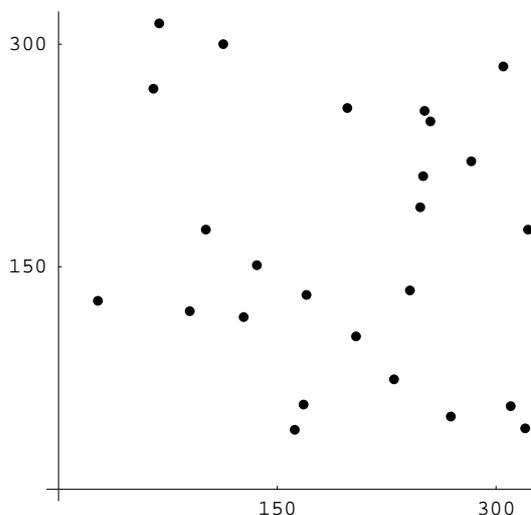


Figure 2.1: *The 25 cities of a sample TSP.*

Reidys and Stadler 2001).

2.1.2 An evolutionary algorithm using point mutation

In order to test correlation statistics for describing fitness landscapes, at first a TSP with rather few cities and comparatively short evolutionary waiting times is chosen. Later I test these statistics for a TSP with 40 cities and a different evolutionary algorithm, see Section 2.1.4. For the following example of the TSP, 25 cities have been distributed randomly according to a uniform distribution over the square $[0, 327]^2$, see Figure 2.1.

The biological terms describing an evolutionary algorithm for the TSP are set to the following (later some different definitions are used):

- Genotypes

Each possible tour starting in City 1, visiting all other cities, and ending again in City 1, corresponds to a genotype; its vector representation is given by the sequence of cities. For every genotype, the length of the respective tour is calculated according to a distance table involving all pairs of cities. For 25 cities, the total number of different genotypes is $24!$, which is of order 10^{23} . Because of this extremely large number of

possibilities a strategy of testing all tours in order to find the best one is infeasible.

- Fitness

The fitness f of a specific genotype g is given by

$$f(g) := \frac{1}{l(g)}$$

where $l(g)$ is the length of the tour corresponding to genotype g .

- Mutation

To mutate a genotype, every operator that changes the vector in a way that the mutated genotype is still a possible tour can be considered. Here, three frequently used mutation operators are chosen (Manderick 1997). For each of these mutation operators, two positions within the tour, corresponding to two indices of the genotype vector, are chosen at random, see Figure 2.2.

- point mutation: the cities at two indices of the vector are swapped. If the fourth and the eighth index are chosen, a mutant of the tour $[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, \dots]$ is the vector $[1, 2, 3, \mathbf{8}, 5, 6, 7, \mathbf{4}, 9, 10, \dots]$.
- reverse mutation: the order of cities between two indices is reversed. A mutant of the tour $[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, \dots]$ is the vector $[1, 2, 3, \mathbf{8,7,6,5,4}, 9, 10, \dots]$.
- remove-and-reinsert: the city at the first index is taken out and reinserted at the second index. A mutant of the tour $[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, \dots]$ is either the vector $[1, 2, 3, \mathbf{5,6,7,8,4}, 9, 10, \dots]$ or $[1, 2, 3, \mathbf{8,4,5,6,7}, 9, 10, \dots]$, depending on the order in which the two indices were chosen.

Two individuals that can be transformed into each other by a single mutation are called neighbors. For the following we use point mutation and show later that the results of a correlation-based description are similar for other mutation operators.

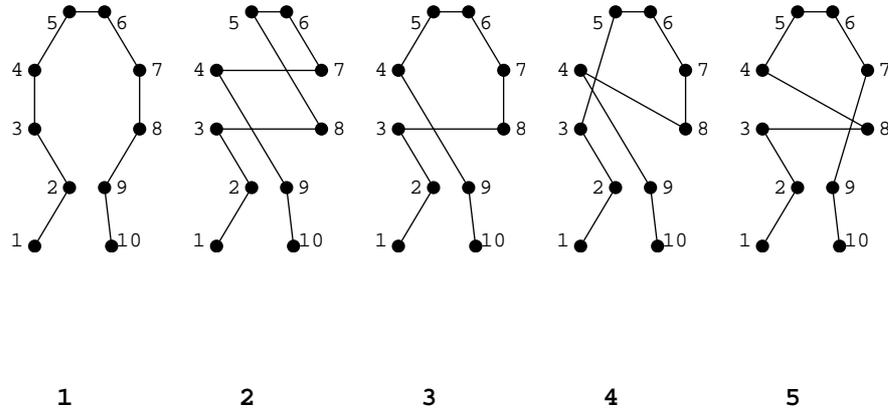


Figure 2.2: *A comparison of different mutation operators for the Travelling Salesman Problem. A part of the original tour (1) and the resulting tours after a point mutation (2), reverse mutation (3) and the two possibilities for remove-and-reinsert mutation (4 and 5), depending on the order the indices are chosen.*

- Selection

In every generation, each genotype produces two offspring individuals which are both once mutated. The best third of the union of the old population and the offspring is taken over to the next generation. This process thus keeps the population size (chosen at 15 individuals below) constant.

- Evolutionary waiting times

Evolutionary waiting times are stochastic variables defined as the number of generations necessary for a population to evolve between two given fitness values. In this work we concentrate on waiting times as these provide crucial statistics of an evolutionary process.

Before we focus on the evolutionary waiting times on our specific TSP landscape, some basic information is given concerning the behavior of the utilized evolutionary algorithm using point mutation.

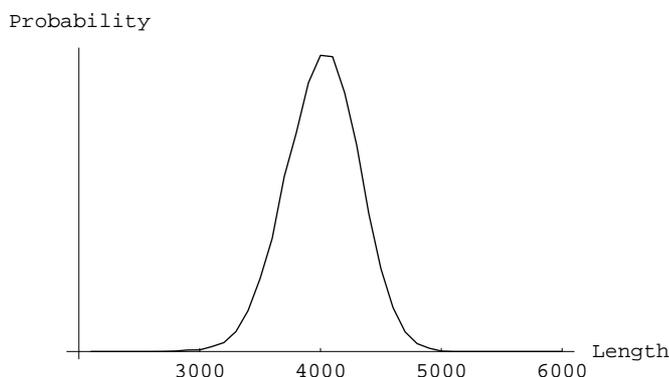


Figure 2.3: *Length distribution of random TSP tours.*

Randomly produced tours normally have a length between 3000 and 5000, the distribution of length values is shown in Figure 2.3.

A population of $P = 15$ individuals, each producing two mutated offspring per generation, rapidly tends to climb the fitness landscape, see Figure 2.4, where the shortest tour length decreases from about 4000 to about 1700 in 500 generations. The best fitness in the population normally remains constant for a number of generations and then suddenly jumps to a higher level. This is a common property of evolutionary algorithms and is referred to as epochal evolution.

The best tour found by all different mutation operators discussed in this paper has a length of 1369 and typically is found within 200 generations, see Figure 2.5. This solution is certainly very close to the global optimum of this TSP.

To test the different reduced descriptions of fitness landscapes studied in this paper, it is necessary to obtain statistics of evolutionary waiting times for different fitness intervals. For this purpose, initial and final fitness values are chosen from the interval $1/5000$ to $1/2900$; producing random tours with fitness values in this range is relatively easy, and this is a prerequisite for obtaining the correlation statistics investigated later. In particular, initial fitness values are chosen $1/5000$, $1/4300$, and $1/3600$, and final fitness values $1/4300$, $1/3600$ and $1/2900$. To construct the distribution of evolutionary waiting times from a certain initial fitness f_i to a final value f_f we proceed as follows. In generation 0, the population is initialized with a random genotype

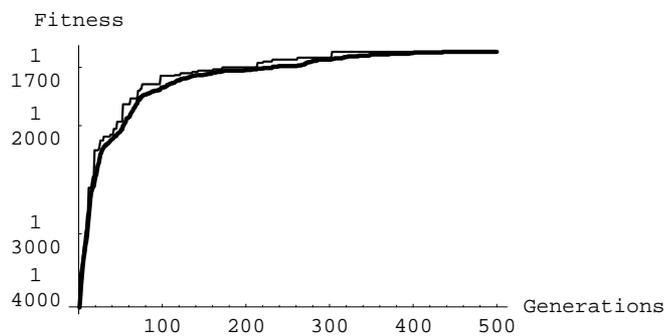


Figure 2.4: *The evolution of highest (thin line) and mean fitness (thick line) in a population. Periods of constant fitness are interspersed with sporadic jumps, a characteristic property of evolutionary algorithms.*

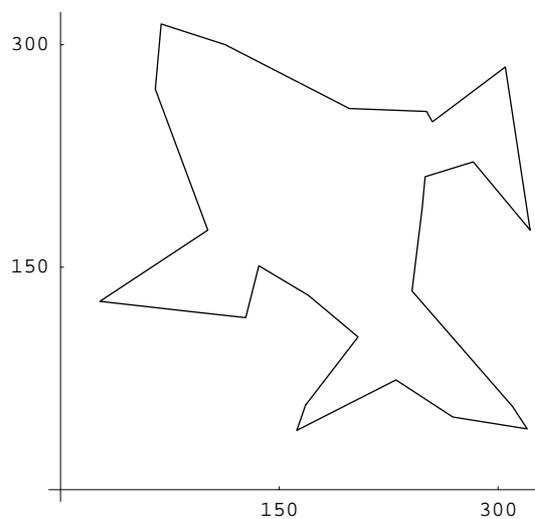


Figure 2.5: *The shortest tour found for our sample TSP has a length of 1369.*

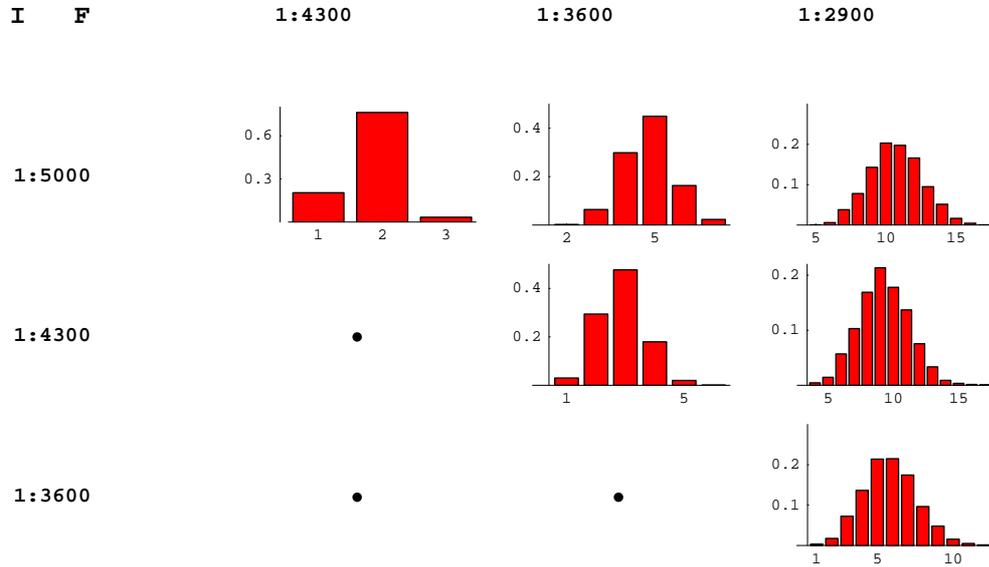


Figure 2.6: A matrix of waiting time distributions in generations for 3 initial (\mathbf{I}) and 3 final (\mathbf{F}) fitness values, based on an evolutionary algorithm that uses point mutation. This graph serves as a target for results derived from different reduced descriptions of fitness landscapes, and allows to assess their performance. As for three cases initial fitness is equal or higher than the final fitness, waiting times are not relevant and the corresponding distributions are replaced with filled circles.

that has a fitness of approximately f_i . The number of generations necessary until one individual of the population reaches fitness f_f is stored as the waiting time of a run. For a given pair of initial and final values, results of 2000 such runs are combined to obtain the distribution of waiting times; an example is shown in Figure 2.6.

In the following subsections I consider landscape statistics of increasing complexity and compare their suitability for predicting the actual distributions of evolutionary waiting times.

Without correlation: Percolation

The basic concept of percolation theory is a grid in a multidimensional space, with each lattice site being independently filled with probability p (Kesten 1982, Grimmett 1989). If p exceeds a certain critical value, a subset of the filled sites forms a connected infinite cluster that percolates through the entire grid. Cluster statistics have been used to study a wide variety of problems (Sahimi 1994, Stauffer and Aharony 1995). Here I use a percolation approximation to provide a reduced description of the TSP fitness landscape presented in the last section. I demonstrate that this approximation, which is ignoring the correlation between the fitness values, does not result in good approximations of the evolutionary algorithm.

If the fitness values of neighboring genotypes are not correlated, the neighborhood structure can be regarded as to be the same for all individuals. In the absence of correlation therefore the distribution of waiting times for a population to evolve from a certain initial to a final fitness value only depends on the frequency of genotypes whose fitness values exceed the final fitness threshold. As the neighborhood structure of an individual is assumed to be independent of this individual's fitness value, initial fitness is ignored in this percolation approximation.

For each final fitness threshold all elements of the genotype space are divided into two classes. Those genotypes with fitness beyond the threshold are in class 1, all others in class 0. A percolation approximation of a fitness landscape then amounts to (i) considering only the labels 0 or 1, while ignoring actual fitness values and genotypes, and (ii) assuming that individuals independently belong to either of these two clusters with probability p and $1 - p$. The percolation probability p is estimated from a large number of random TSP tours.

The evolutionary algorithm is then imitated as follows: The entire population is in class 0 initially. In each generation, and for every individual, two offspring are chosen out of 300 neighbors of this individual (for 25 cities each TSP genotype has $25 * 24/2 = 300$ neighbors under point mutation). If at least one of these 30 offspring individuals (15 individuals produce two offspring each) belongs to the higher fitness class (class 1), the process stops. Otherwise it continues with the next generation.

In order to calculate the probability $p_{stop}(g)$ for the described process to end in a certain generation g , first some variables have to be defined:

A ... total number of possible individuals

P ... population size

N ... number of neighbors per individual

m ... number of offspring per generation and individual

p ... probability for an individual to be in the higher fitness class (class 1)

A_1 ... number of all individuals in class 1 ($\approx Ap$)

The probability p_k for an individual to have k neighbors in class 1 out of N possible is then given by

$$\begin{aligned} p_k &= \frac{\binom{A_1}{k} \binom{A-A_1}{N-k}}{\binom{A}{N}} \\ &= \frac{A_1! (A - A_1)! N! (A - N)!}{k! (N - k)! (A_1 - k)! (A - A_1 - N + k)! A!} \\ &= \binom{N}{k} \frac{A_1! (A - A_1)! (A - N)!}{(A_1 - k)! (A - A_1 - N + k)! A!} \end{aligned}$$

If N and k are relatively small compared to A (which is true for this TSP landscape), this equation is well approximated by

$$p_k \approx \binom{N}{k} p^k (1 - p)^{N-k}.$$

Now, m offspring are chosen out of the N neighbors; these are not necessarily different. The probability that none of m offspring of an individual belongs to class 1 is $\sum_{i=0}^{N-1} p_i \left(\frac{N-i}{N}\right)^m$. As the events, that no offspring of all different individuals belongs to class 1, are independent from each other, the probability p^* that at least one offspring of the whole population is of higher fitness can now be calculated:

$$p^* = 1 - \left(\sum_{i=0}^{N-1} p_i \left(\frac{N-i}{N} \right)^m \right)^P.$$

The probability $p_{stop}(g)$ that after g generations one offspring finally belongs to class 1, is then given by

$$p_{stop}(g) = (1 - p^*)^{g-1} p^*.$$

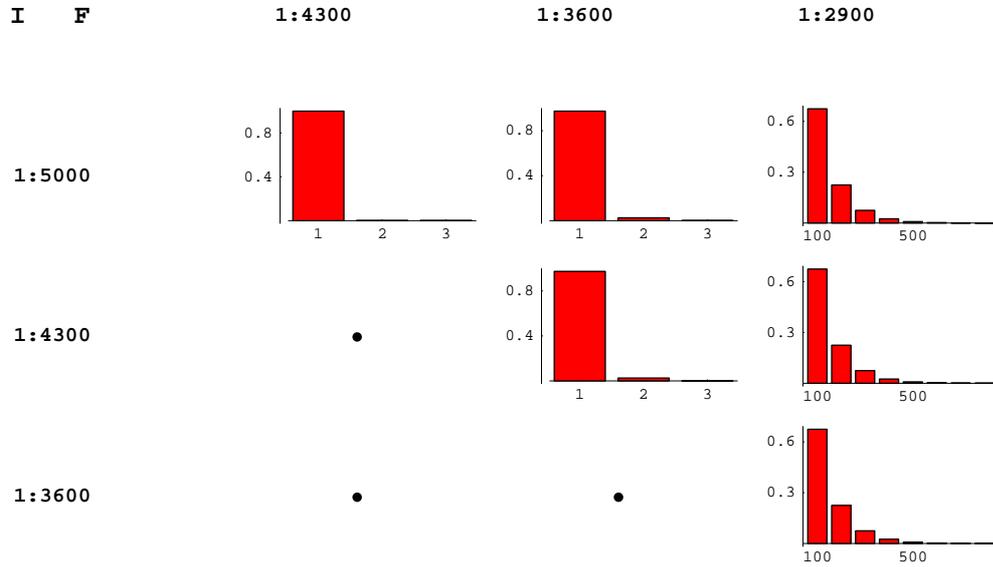


Figure 2.7: *Waiting time distributions for reaching certain final fitness thresholds as predicted by the percolation approximation. Initial fitness values are ignored in this approximation. Compare the depicted distributions to those in Figure 2.6.*

For the analyzed TSP fitness landscape and the utilized evolutionary algorithm the parameters are fixed to $A = 24!$, $P = 15$, $N = 300$, and $m = 2$. The percolation probabilities p , denoting the probability for an arbitrary individual to be of fitness above a threshold, are, depending on the three final fitness thresholds used before, approximately given by $p(1/4300) = 0.86429$, $p(1/3600) = 0.114685$, and $p(1/2900) = 0.000367$.

It is not surprising that this drastic simplification of the fitness landscape is too coarse. The waiting time statistics in Figure 2.7 show that the actual evolutionary process towards high fitness values is much faster than this reduced description suggests. In the actual process, the probability for choosing a neighbor with fitness above the final value increases over the generations as the population successively attains higher fitness values. The percolation approximation cannot capture this critical effect.

As a next step I thus incorporate a critical landscape feature: the neighborhood of a genotype strongly depends on its own fitness. Thus, the fitness values of neighboring genotypes are correlated.

Monomorphic correlation

The results obtained for the percolation approximation suggest dividing all genotypes into more than just two fitness classes, with each class having a different distribution of mutant fitness. The correlation c_{ij} between classes i and j is the probability for a random neighbor, or mutant, of an individual of class i to be in class j . As class frequencies vary, this correlation function is not symmetric. The correlation matrix $C = (c_{ij})$ is then used to define the transition matrix T of a Markov chain that approximates the evolutionary algorithm (Rudolph 1997). In this approximation, the whole population is still considered to reside in the same fitness class, and is thus assumed to be monomorphic at this level of description. Given that the population is in class i , a transition from class i to a higher fitness class j occurs, if at least one offspring of all the individuals belongs to fitness class j , but no offspring is in one of the classes higher than j . Transition to lower fitness classes are not possible; the population will therefore remain in the same class if no offspring possesses a higher fitness. To calculate the probability t_{ij} for a transition from i to j we need the following variables, assuming that i and j are fixed:

c_+ ... union of all fitness classes higher than j

c_0 ... class j

c_- ... union of all fitness classes lower than j

p_l ... probability for a random neighbor of an individual of class i to belong to $c_l, l = +, 0, -$

P ... population size

N ... number of neighbors per individual

m ... number of offspring per generation and individual

The probability w_{lk} for an individual of the population, which thus belongs to class i , to have k neighbors in c_l is determined by a binomial distribution, $w_{lk} = \binom{N}{k} p_l^k (1 - p_l)^{N-k}$. For each individual, m offspring are independently chosen out of N neighbors. The probability t_{ij} for a transition from class i to j is the probability, that no offspring of all individuals is in c_+ , but at least one in c_0 . Let A be the event that no offspring is in c_+ , and B the event that at least one offspring is in c_0 . Thus t_{ij} is given by the probability $P(A \cap B)$ that both events A and B occur. If \bar{B} is the complement of B and thus the event that no offspring is in c_0 , we get

$$\begin{aligned} P(A \cap B) &= P(A) - P(A \cap \bar{B}) \\ &= P(A) - P(C) \end{aligned}$$

where C denotes the event that all offspring are in class c_- . Thus,

$$t_{ij} = \left(\sum_{n=0}^{N-1} w_{+n} \left(\frac{N-n}{N} \right)^m \right)^P - \left(\sum_{n=1}^N w_{-n} \left(\frac{n}{N} \right)^m \right)^P.$$

The resulting transition matrix $T = (t_{ij})$ defines a Markov chain and allows to derive distributions of waiting times between the three pairs of initial and final fitness values already used in the last sections. Moreover, the mean number of generations needed for attaining the absorbing state (final fitness) from different starting classes (initial fitness) can be computed analytically (see Kemeny 1960). These results are presented below.

Coarse-grained correlation.

We begin by introducing four fitness classes, separated by the three final fitness values used, e.g., in Figure 2.6. The correlation matrix $C = (c_{ij})$, where c_{ij} is the probability for a random neighbor of an individual of class i to belong to class j , is estimated by randomly mutating random genotypes of class i , see Figure 2.8. The resulting 4×4 transition matrix defines a Markov chain for which waiting time distributions are computed. These turn out to be closer to the actual ones; yet, systematic differences of mean values and variances illustrate the need for further refinement of this correlation-based approach.

Figure 2.9 shows that even if the population already is in the class next to the final one, producing an offspring the fitness of which exceeds the final fitness threshold takes too much time. The many transitions within classes,

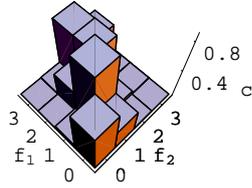


Figure 2.8: The correlation matrix for 4 fitness classes. $c(f_1, f_2)$ denotes the probability for a random offspring of an individual of fitness class f_1 to belong to class f_2 . The classes range from 0 (lowest fitness) to 3 (highest fitness). These statistics require refinement.

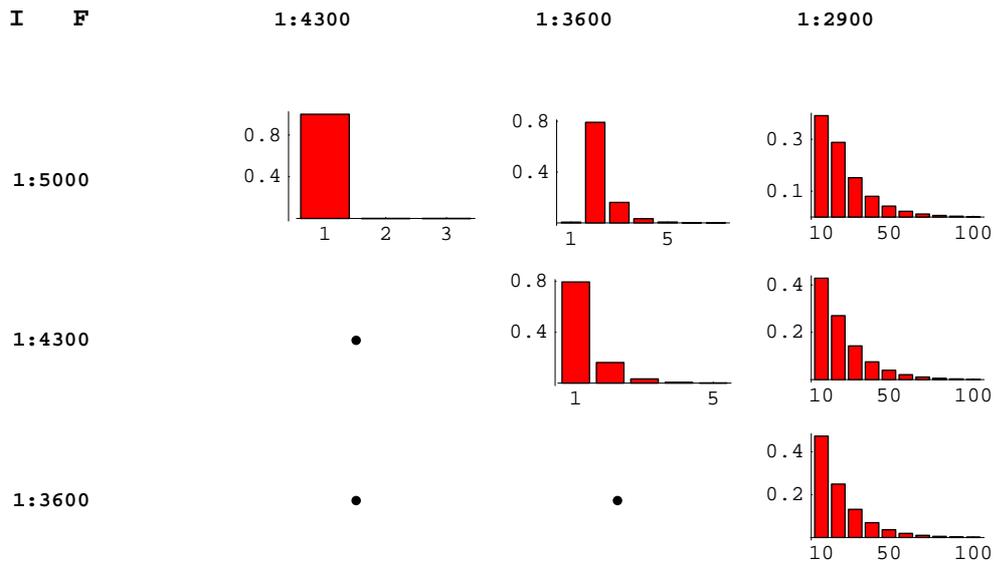


Figure 2.9: Waiting time distributions for transitions between given initial and final fitness values as predicted by monomorphic evolution based on a coarse-grained correlation approximation. Comparison of depicted distributions to those in Figure 2.6 shows that waiting times are overestimated by this approach.

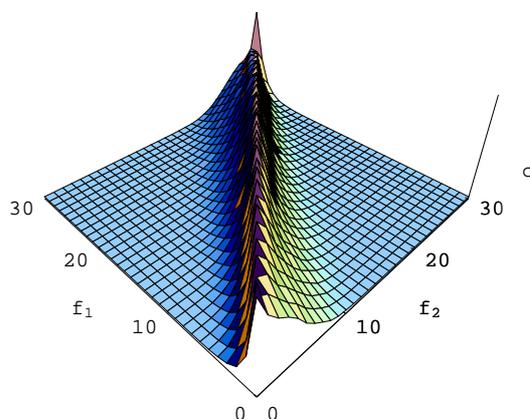


Figure 2.10: *The correlation matrix for 31 fitness classes. $c(f_1, f_2)$ denotes the probability for a random offspring of an individual of fitness class f_1 to belong to class f_2 . The classes range from 0 (lowest fitness) to 31 (highest fitness). For all classes neighboring genotypes tend to have the same or similar fitness values.*

leading from the lower bound of a class' fitness range to the upper bound are neglected by only allowing for a small number of classes. This observation suggests to introduce a fine-grained classification of fitness values.

Fine-grained correlation. To improve the predictive accuracy of the correlation approximation, we consider a 31×31 correlation matrix. 30 equally spaced fitness thresholds between tour lengths 2900 and 5000 serve as the boundaries of a fine-grained classification. As before, the correlation matrix is obtained by randomly generating neighbors of random genotypes, see Figure 2.10. In this manner, small changes in fitness values, which can be decisive for the dynamics of an evolutionary algorithm, are no longer neglected.

The fact that most of the probability mass of the correlation matrix is concentrated around its diagonal indicates that neighboring genotypes tend to possess similar fitness values. Yet, it is possible that they belong to distant fitness classes. Using the same principles for constructing a Markov chain as described above, Figure 2.11 shows the resulting transition matrix $T = (t_{ij})$ of the process.

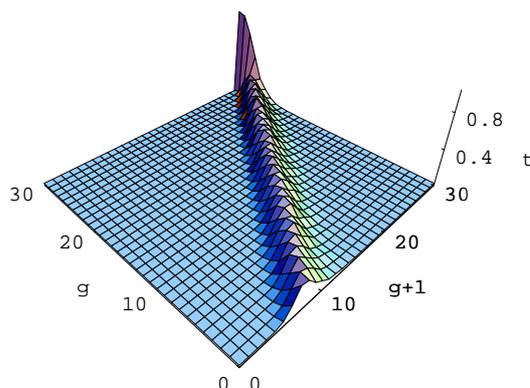


Figure 2.11: *The transition matrix of the Markov chain that provides a correlation-based approximation of the evolutionary algorithm. A population at generation g jumps to higher fitness classes in the next generation with probability t . Class 31 is called an absorbing state; if the population reaches it, the process stops.*

The fine-grained correlation matrix provides a detailed summary of the adjacency relations between the different fitness classes; the neighborhood structure for different TSP genotypes should therefore be described with sufficient accuracy. We thus might expect that the waiting time distributions derived from this simplified process are a close match to the actual ones. And, indeed, the fine-grained monomorphic correlation approximation is the first approach presented here that succeeds in capturing many of the qualitative and quantitative features of the evolutionary algorithm as summarized in Figure 2.6. The results presented in Figure 2.12 therefore underline that a fine-grained correlation matrix as defined above carries salient information about the fitness landscape's structure.

Although correlation approximation of the fitness landscape results in predictions that are not far away from observations on the actual evolutionary algorithm, it is interesting to ask why the simplified process is always about twice as fast as the evolutionary algorithm itself. An investigation of fitness distributions for all generations of the evolving population reveals that in each generation fitness values of genotypes in the population are spread over a wide range of classes. For this reason, the assumption of monomorphism (implying that all individuals of the population belong to a single fitness

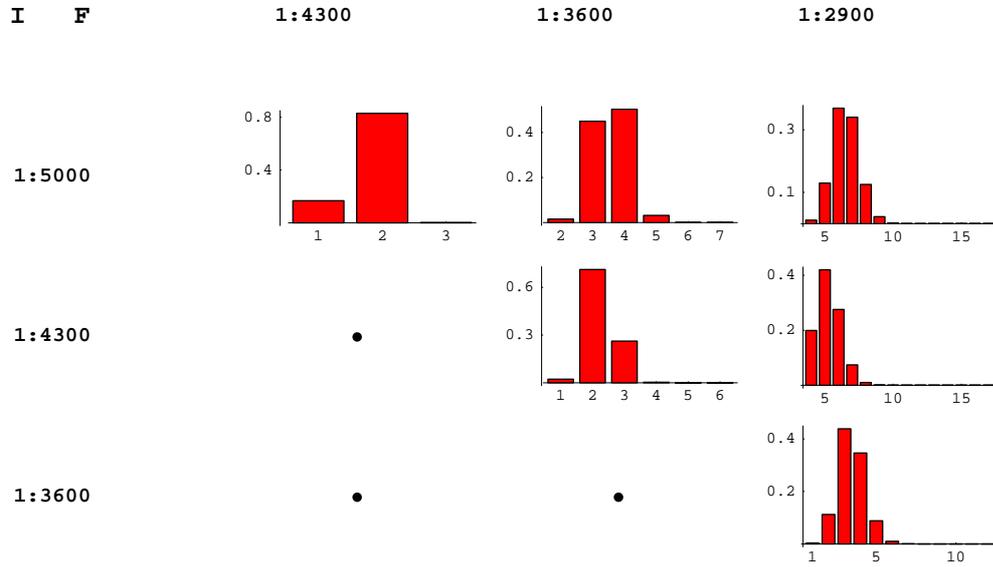


Figure 2.12: *Waiting time distributions as predicted by monomorphic evolution based on a fine-grained correlation approximation. Comparing these distributions with those in Figure 2.6 shows that a relatively good approximation of the actual evolutionary algorithm has been achieved.*

class) is inaccurate. Instead of assuming transitions of the whole population's fitness from one class to another, individuals in the population need to be envisaged as belonging to different fitness classes.

Polymorphic correlation

In this subsection the approximation of the evolutionary algorithm of the TSP is further improved. We use the same statistics as before, namely the 31×31 matrix of correlation probabilities for each pair of fitness classes. However, instead of considering only one class that represents the fitness of the whole population, individuals of the population can now belong to different fitness classes in each generation. The population's state in one generation is no longer a specific fitness class, but is given by a frequency distribution over all 31 possible fitness classes. In other words, we allow the popula-

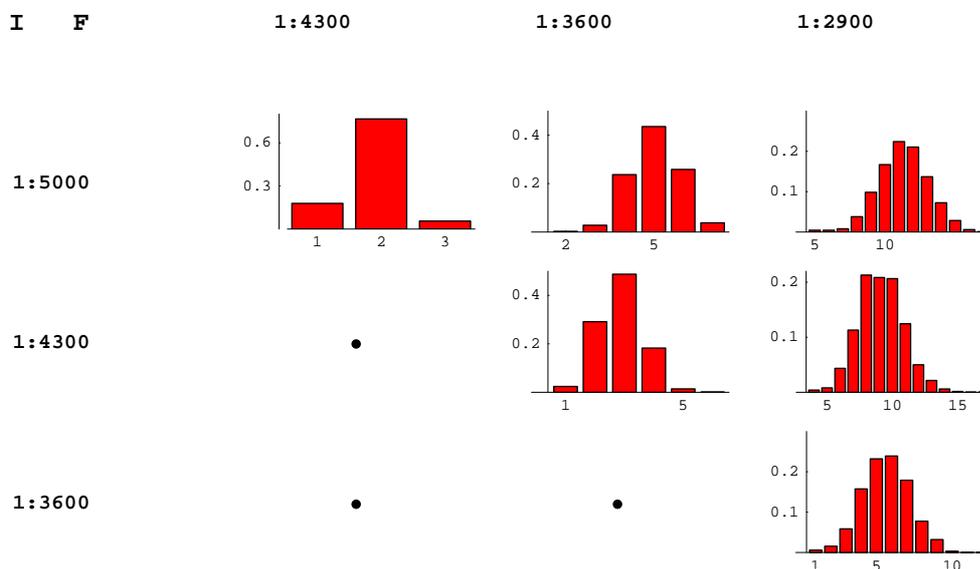


Figure 2.13: *Waiting time distributions as predicted by polymorphic evolution based on a fine-grained correlation information. A comparison with Figure 2.6 shows that this approach yields almost exact predictions of waiting times. Although the full TSP is approximated by a 31×31 matrix, the match with the actual process is remarkably good.*

tion to be polymorphic. As before, offspring from a given fitness class are produced according to the probabilities provided by the correlation matrix. The transition matrix of this Markov process describes the probability for a population with a certain frequency distribution of fitness values to jump to another composition of fitness classes in the next generation. As there are $\binom{45}{15}$ such population states, the transition matrix was not calculated. Instead, I have directly implemented the stochastic process based on the correlation matrix and have combined the outcome of 2000 trials to construct the distribution of waiting times, shown in Figure 2.13. This stochastic process is a simplification of the actual evolutionary algorithm: individuals are reduced to only their fitness classes, and the fitness classes of their offspring are determined according to the correlation matrix. Selection is based on

fitness only and thus works in the simplified process in the same way as in the actual evolutionary algorithm. The complicated process of evolution on the high-dimensional TSP fitness landscape is therefore described by means of the correlation matrix.

Comparison of mean waiting times from different initial to final fitness values from the outcome of 2000 runs of the actual evolutionary algorithm (first number) and 2000 runs of the stochastic approximation algorithm based on polymorphic correlation (second number), corresponding to the waiting times distributions depicted in Figure 2.6 and 2.13:

initial / final fitness	1:4300	1:3600	1: 2900
1:5000	1.8 / 1.9	4.8 / 5	10.7 / 11.3
1:4300	*	2.9 / 2.9	9.3 / 9.1
1:3600	*	*	5.8 / 5.6

It turns out that, at least for this special fitness landscape, the polymorphic correlation approximation is an excellent way to predict the time scales of evolution. Comparing these results with those derived by applying the monomorphic correlation approximation shows that allowing for the specific composition of fitness values within a population is crucial for obtaining accurate predictions. These findings raise the question whether this approach will also perform for other fitness landscapes in a similarly accurate and equally successful way. A first test is to investigate other fitness landscapes of this TSP which result from considering new mutation operators.

2.1.3 Other mutation operators

Reverse mutation

In this section, we use the same configuration of 25 cities and maintain all other parameters, only the mutation operator is changed. For producing an individual's offspring by reverse mutation, we choose two indices. The cities between the smaller and the larger index are now rearranged in reverse order. As two genotypes can be neighbors under reverse mutation while being separated by a large distance in genotype space under other mutation operators, resulting fitness landscapes may have very different features. The importance of the mutation operator for determining the structural features of fitness landscapes is underscored by the fact that under reverse mutation

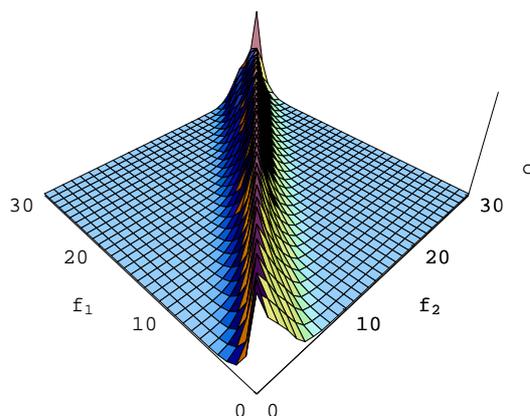


Figure 2.14: *The correlation matrix for reverse mutation. Producing an offspring with significantly higher fitness appears to be easier under point mutation, at least in the depicted range of fitness values.*

the evolutionary algorithm needs many more generations for reaching the three final fitness classes, than were required under point mutation. On the other hand, tours with length under 1400 are hardly ever found within 100 generations under point mutation, whereas this task seems to be achieved much more rapidly under reverse mutation. Even the tour of length 1369, see Figure 2.5, was found within this time limit.

As the polymorphic correlation approach provided the best approximation of the actual waiting times under point mutation, here we focus on an evaluation of this way of imitating the evolutionary process under reverse mutation. Like in the case of point mutation, the correlation matrix is estimated by generating random neighbors of random genotypes. The structure of neighborhoods is similar to the one obtained for point mutation and yet possesses some different properties, see Figure 2.14 .

The waiting time distributions obtained from 2000 simulation runs of the correlation-based stochastic process exhibit a very close match with the waiting times of the actual evolutionary algorithm, see Figure 2.15. Again, the polymorphic correlation approximation is successful in describing the evolutionary algorithm. This implicates that the correlation matrix provides all the information necessary to predict evolutionary waiting times on TSP fitness landscapes.

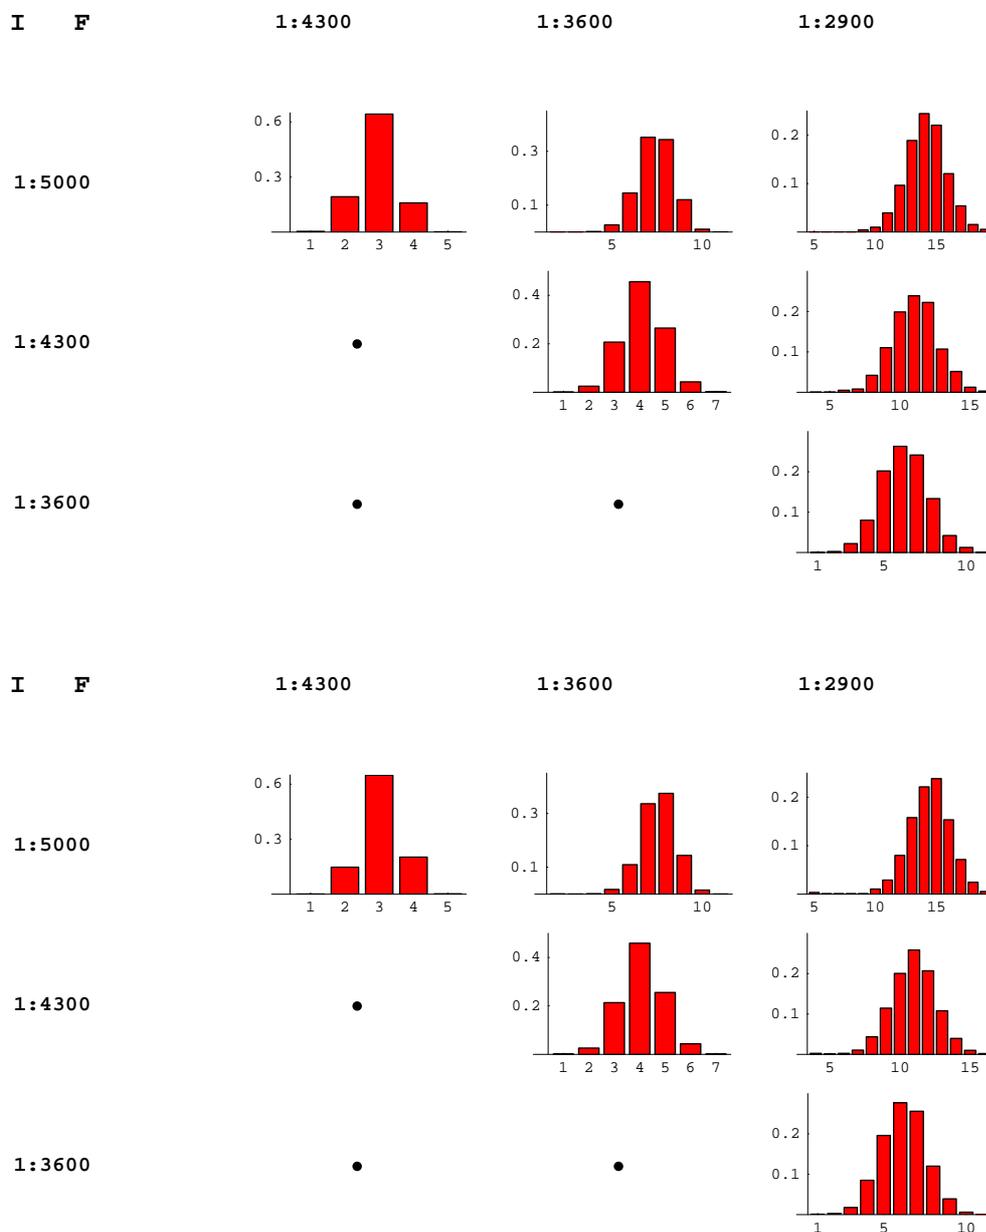


Figure 2.15: *Waiting times of the evolutionary algorithm with reverse mutation (top) and distributions predicted by the polymorphic correlation approximation (bottom). The predictive accuracy is remarkable.*

Corresponding to the distributions presented in 2.15, the following table shows a comparison of mean waiting times of the actual evolutionary algorithm (first number) using reverse mutation and the predictions based on polymorphic correlation:

initial / final fitness	1:4300	1:3600	1: 2900
1:5000	3.0 / 3.1	7.4 / 7.6	14.2 / 14.4
1:4300	*	4.1 / 4.1	11.1 / 11
1:3600	*	*	6.3 / 6.2

We extend our investigations of the TSP by investigating a third mutation operator, resulting in yet another fitness landscape.

Remove-and-reinsert mutation

To apply the remove-and-reinsert mutation operator, we choose two indices of the genotype vector successively. The city at the first index will now be taken out and inserted at the second index. The cities in between move backwards or forwards by one index. Notice that for such remove-and-reinsert mutations the order of the two indices is important, whereas point mutations and reverse mutations are symmetric in this respect. Using the same configuration of the TSP, each genotype now has 600 neighbors. Nonetheless, the correlation matrix has the same basic features as before, characterized by a strong emphasis of correlations along the diagonal, see Figure 2.17. Again we construct a stochastic process based on the correlation matrix only and combine the outcome of 2000 runs to the predicted waiting time distribution.

A comparison of actual and predicted waiting times for remove-and-reinsert mutation, see Figure 2.16, demonstrates again that the correlation matrix captures all the information necessary to describe the complex behavior of the evolutionary algorithm. Comparison of actual and predicted mean waiting times for remove-and-reinsert mutation:

initial / final fitness	1:4300	1:3600	1: 2900
1:5000	2.5 / 2.5	6.3 / 6.3	13.1 / 12.8
1:4300	*	3.6 / 3.5	10.7 / 10.0
1:3600	*	*	6.3 / 5.9

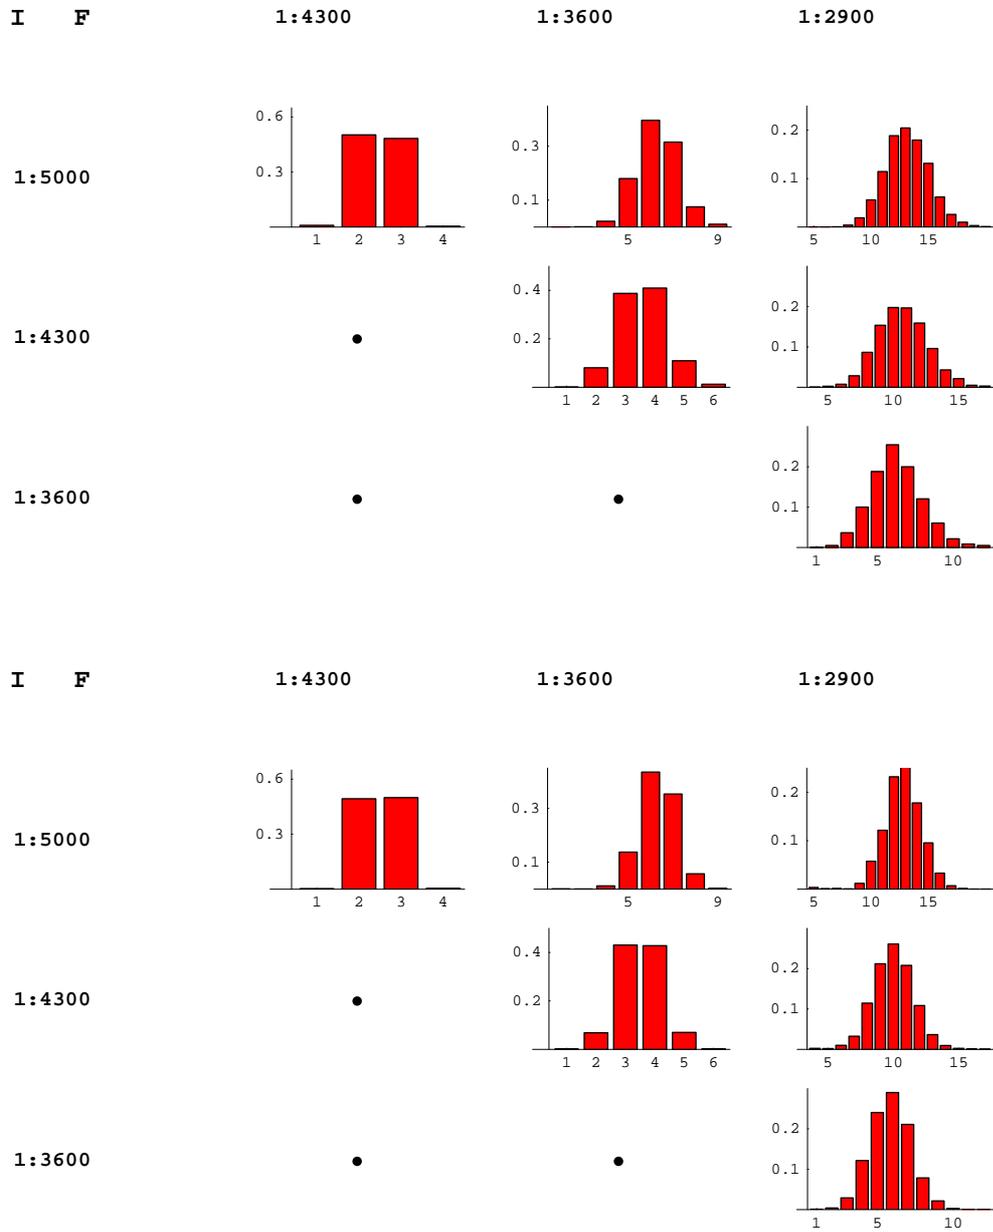


Figure 2.16: *Actual (top) and predicted (bottom) waiting times for remove-and-reinsert mutation.*

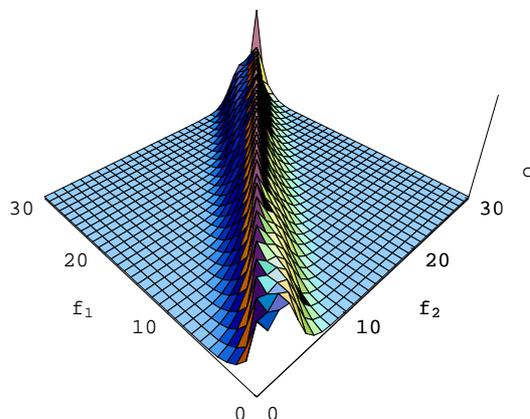


Figure 2.17: *The correlation matrix for remove-and-reinsert mutation.*

In the next section we test the correlation-based description for another Travelling Salesman Problem and a different evolutionary algorithm, and show that the one-dimensional correlation matrix again captures the crucial features of a fitness landscape that determine evolutionary waiting times.

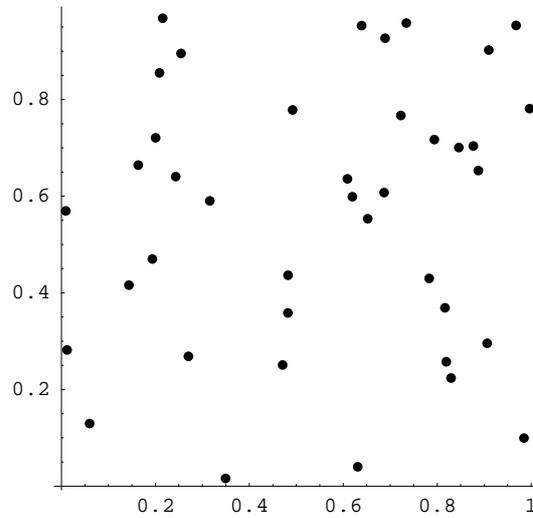
2.1.4 A second TSP and a variant evolutionary algorithm

In the last chapter we have seen that an evolutionary process on a specific TSP fitness landscape was very well approximated by a polymorphic description based on one-dimensional correlation statistics. Now we test this stochastic approximation for another Travelling Salesman Problem. For this purpose, 40 cities are distributed randomly according to a uniform distribution over the square $[0, 1]^2$ and the distance table is calculated.

We choose a new definition of fitness and construct an evolutionary algorithm with stochastic selection:

- Genotypes

Same definition as for the first TSP: each possible tour starting in City 1, visiting all remaining 39 cities of the TSP, and ending again in City 1, corresponds to a genotype. The number of different genotypes is $39!$ and thus of order 10^{46} .

Figure 2.18: *The 40 cities of the new sample TSP.*

- Fitness

According to the central limit theorem, lengths of TSP tours are approximately normally distributed, as they are sums of random distances between pairs of cities. In order to make the corresponding fitness values follow a normal distribution as well, a linear transformation of the length is chosen. As selection will be no longer ranking-based but stochastic, it is moreover convenient to make fitness values belong to the interval $[0, 1]$. For calculating the fitness therefore an upper bound $L_{sup} = 30$ for the length of all possible TSP tours and a lower bound $L_{inf} = 5$ is chosen. The fitness $f(g)$ of genotype g is then defined as follows:

$$f(g) := \frac{L_{sup} - l(g)}{L_{sup} - L_{inf}} \in [0, 1]$$

where $l(g)$ is the length of the tour represented by genotype g .

- Mutation

For this TSP we choose reverse mutation: a genotype is mutated by reversing the order of cities between two randomly chosen indices.

- Selection

In every generation, each genotype produces one offspring individual which undergoes a single reverse mutation. From the whole population of P offspring we select P times independently an individual according to a stochastic selection scheme with selection coefficient $S \geq 0$. Individual i having fitness $f(i)$ is chosen with probability

$$p(i) = \frac{e^{Sf(i)}}{\sum_{j=1}^P e^{Sf(j)}}$$

The higher the selection coefficient, the lower is the probability for individuals of smaller fitness values to be taken over to the next generation.

- Evolutionary waiting times

Again we focus on the number of generations necessary for a population to evolve between to given fitness thresholds for a comparison of the actual evolutionary algorithm and the correlation-based description. In generation 0 the whole population of P individuals is initialized with a random genotype with a fitness value lying in the interval $[f_i - 0.01, f_i]$, if f_i is the initial fitness. Over generations, individuals undergo mutation and selection until one genotype's fitness exceeds the final fitness threshold. The number of generations necessary is then stored, and the outcome of 1000 such runs of the evolutionary algorithm is combined to the distribution of waiting times.

The definition of fitness we use for this TSP implies that not only the lengths of TSP tours, but also fitness values of corresponding genotypes follow a normal distribution, see Figure 2.19. Figure 2.20 shows the evolution of maximal and mean fitness in single run of the evolutionary algorithm with a population of $P = 100$ individuals and the selection coefficient set to $S = 250$. The best tour ever found by this evolutionary algorithm is presented in Figure 2.21 and was detected within 100 generations for a parameter setting of $P = 100$ and $S = 250$.

For a test of one-dimensional correlation statistics describing this TSP fitness landscapes, we proceed as we did in the last chapter and construct a matrix of 6 actual waiting time distributions between several initial and final

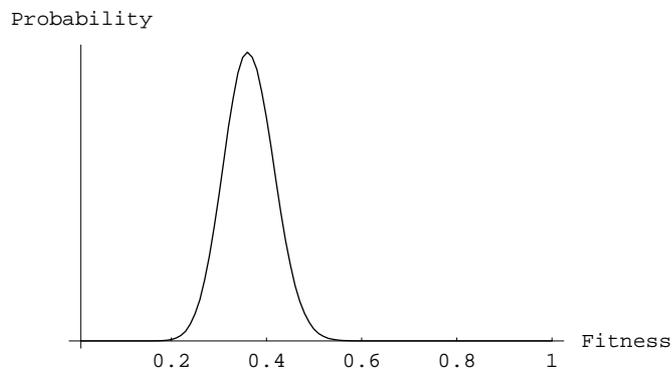


Figure 2.19: *The fitness distribution of TSP genotypes is very similar to a normal distribution.*

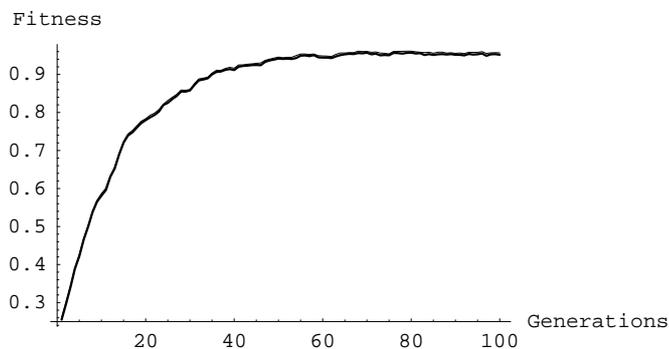


Figure 2.20: *The evolution of the highest and mean fitness in a population with 100 individuals and a selection coefficient of 250, derived from a single run. Due to the high selection coefficient, highest and mean fitness are practically identical.*

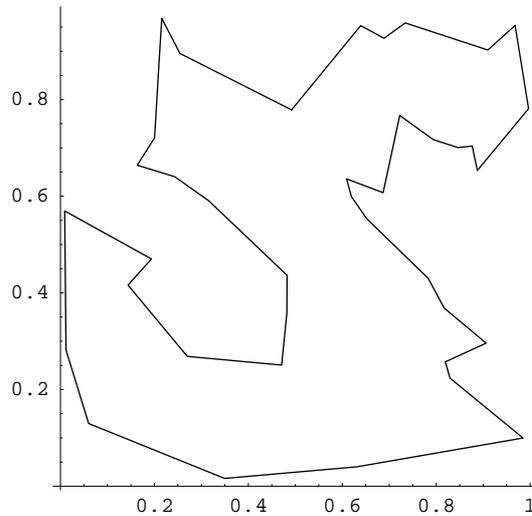


Figure 2.21: *The shortest tour found for our sample TSP has length 5.2 and fitness 0.99.*

fitness values, which then serves as a target for correlation-based predictions. As in the case of the first TSP, it is not possible to obtain correlation statistics for very high and very low fitness, because the probability for random genotypes to be of these fitness values is practically zero. The focus is therefore at first on waiting times between fitness values for which it is possible to obtain correlation statistics. Subsequently I present a way to extend correlation statistics to the whole fitness range, and show by comparing actual and predicted waiting times that this method is successful.

At first, however, the focus is on waiting times within relatively small fitness intervals. For this purpose the initial fitness thresholds are set to 0.2, 0.3, and 0.4 and the final thresholds to 0.4, 0.5, and 0.6. The short waiting times between 0.3 and 0.4, 0.4 and 0.4, and 0.4 and 0.5 are neglected. The population size is set to $P = 10$ and as selection coefficient $S = 20$ is chosen. The resulting 6 waiting times distributions are derived from 2000 runs of the evolutionary algorithm for each pair of initial and final fitness.

In order to test if predictions based on correlation statistics are sufficiently accurate for this TSP as well, we divide the interval $[0, 1]$ of possible fitness values into 100 classes with boundaries $0.01, 0.02, \dots, 0.99$. By randomly mutating random genotypes, the correlation matrix is estimated, see

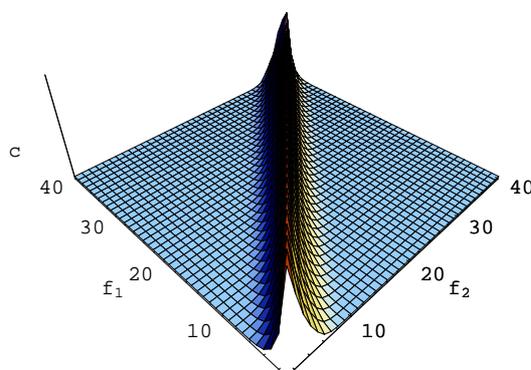


Figure 2.22: *The correlation matrix for the TSP landscape with 40 cities. The fitness classes range from class 20 to 60 which corresponds to fitness values 0.2 - 0.6. For lower and higher fitness it is difficult to obtain valid statistics.*

Figure 2.22. Valid correlation statistics are only obtained for the range of fitness classes for which random genotypes have a reasonable probability to belong to. Nevertheless, the structure of the correlation matrix shows, that mutant fitness is approximately normal distributed, and thus the correlation statistics allow for a continuous extension towards very low and very high fitness classes. Some results of this approach are presented later in this section and in Chapter 4. The reason for normal mutant fitness distribution is that under reverse mutation, two random distances of a TSP tour are replaced by random others. According to the central limit theorem mutant fitness is approximately normally distributed.

As a first examination of the correlation-based predictions, we use the actually obtained matrix of correlation statistics and construct a simplified stochastic process in which the matrix, the only information about the fitness landscapes, provides the information for mutating individuals. The outcome of 1000 runs of this process is used to predict waiting time distributions, see Figure 2.23 for a comparison with the actual waiting times of the evolutionary algorithm.

A comparison of actual (first number) and predicted (second number) mean waiting times, corresponding to the distributions in Figure 2.23, for evolution from certain initial to final fitness thresholds, shows that the one-

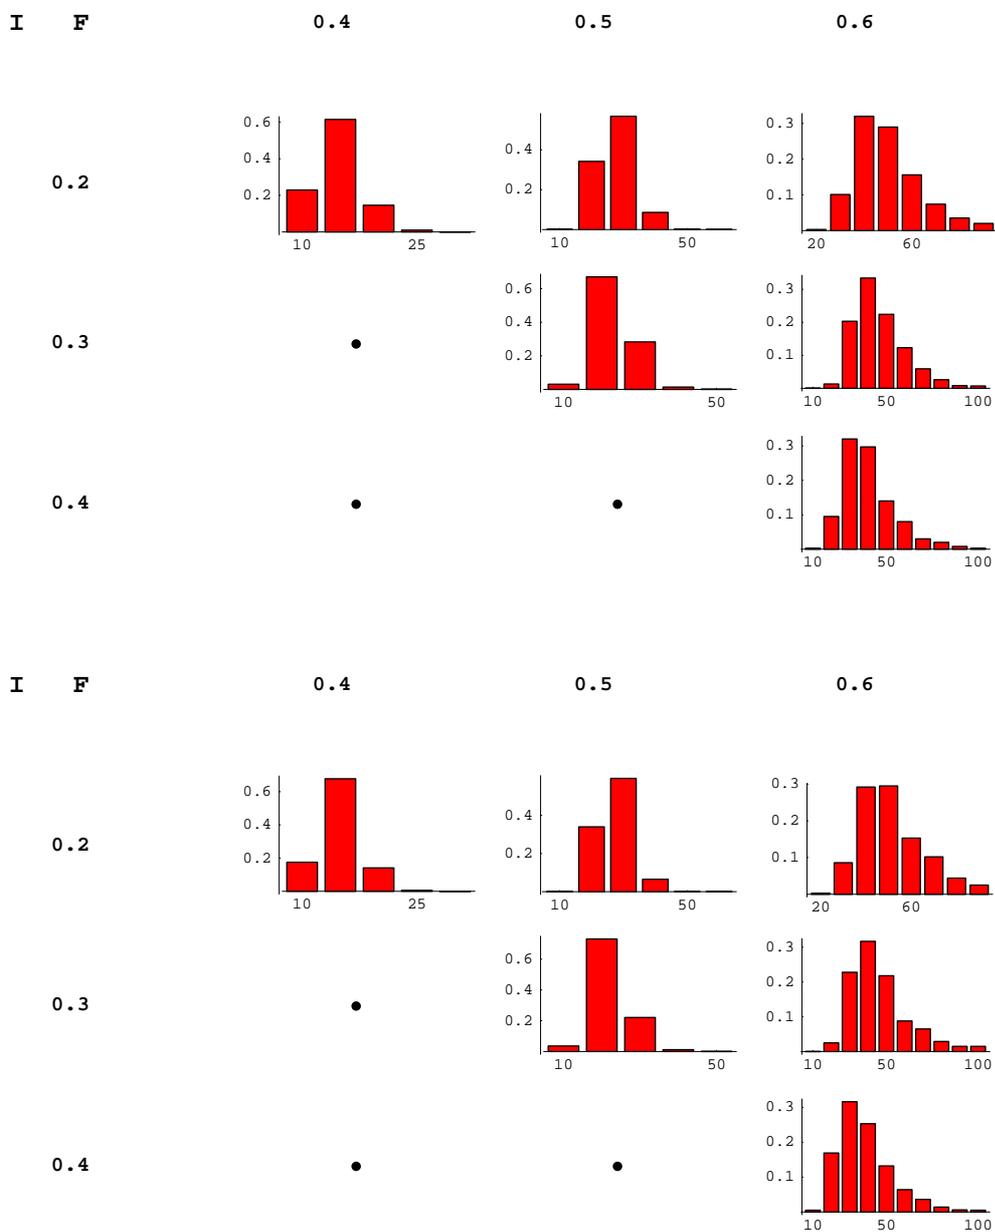


Figure 2.23: *Actual (top) and predicted (bottom) waiting times for the TSP with 40 cities. The one-dimensional correlation matrix basically provides all the information necessary to predict the evolution of fitness in these ranges.*

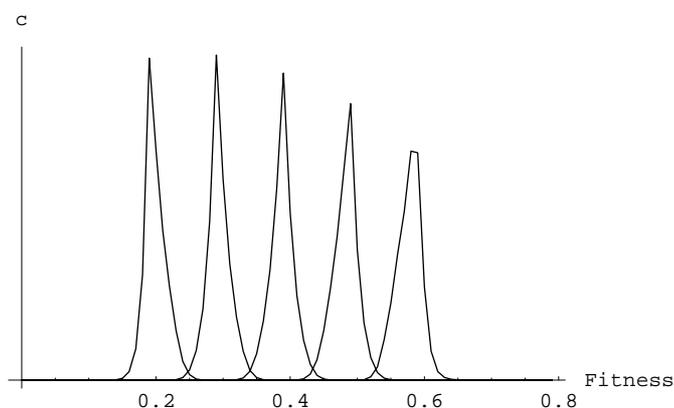


Figure 2.24: *Distribution of mutant fitness for parent fitness 0.2, 0.3, 0.4, 0.5, and 0.6, seen from the left, can be well approximated by normal distributions.*

dimensional correlation matrix essentially provides all the information necessary to specify the hill-climbing process of individuals on the TSP fitness landscape.

initial / final fitness	0.4	0.5	0.6
0.2	11.6 / 11.9	22.0 / 21.9	44.1 / 45.7
0.3	*	17.3 / 16.5	40.7 / 40.5
0.4	*	*	34.5 / 32.6

A detailed investigation of the correlation statistics shows that mutant fitness is approximately normal distributed with the mean value and variance depending linearly on the parent fitness values in a very good approximation, see Figure 2.24 and 2.25. We therefore extended the correlation matrix for the fitness classes for which no statistics were obtained, due to very low class frequencies, by normal distributions with parameters defined by two linear functions. In order to verify if this approximation of correlation statistics is valid, a new evolutionary algorithm is constructed, with population size $P = 100$ and the selection coefficient set to $S = 250$ in order to enable evolution to very high fitness values within a reasonable number of generations. Whereas initial fitness values are kept, new final fitness thresholds are chosen: 0.7, 0.8, and 0.9. Figure 2.26 shows a comparison of actual and predicted waiting

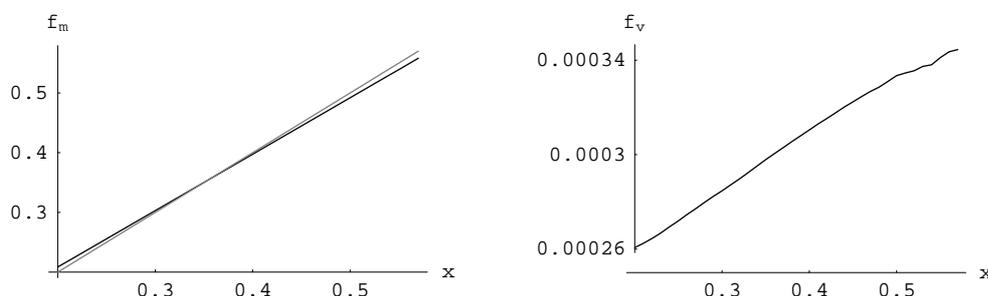


Figure 2.25: The functions $f_m(x)$ and $f_v(x)$ for mean and variance of neighbor fitness depending on parent fitness x . In the first picture the identity function is depicted in gray. Genotypes of lower fitness tend to have higher mean value of mutant fitness whereas it is the opposite case for genotypes of higher fitness. For parent fitness lower than 0.2 and higher than 0.6 no valid statistics were obtained. Both functions are linear in a good approximation.

times distributions, both derived from 1000 runs per pair of initial and final fitness.

Actual (first number) and predicted mean values of waiting times for the distributions shown in Figure 2.26:

initial / final fitness	0.7	0.8	0.9
0.2	14.1 / 14.3	20.4 / 21.3	34.3 / 35.1
0.3	*	18.7 / 19.1	32.6 / 32.8
0.4	*	*	30.8 / 30.3

The accuracy of these results indicates that TSP fitness landscapes can be described successfully by two functions which determine the mean and variance of mutant fitness for individuals of certain fitness. The one-dimensional correlation statistics are then approximated by normal distributions with parameters determined by these two functions. The simplicity of such a description of TSP landscapes is surprising, and certainly a remarkable result in the study of fitness landscapes.

Statistics describing the structure of neighborhoods for genotypes of different fitness have proved to carry the appropriate information for predicting evolutionary waiting times on TSP landscapes. We now leave the realm of TSP landscapes and demonstrate successful applications of this method

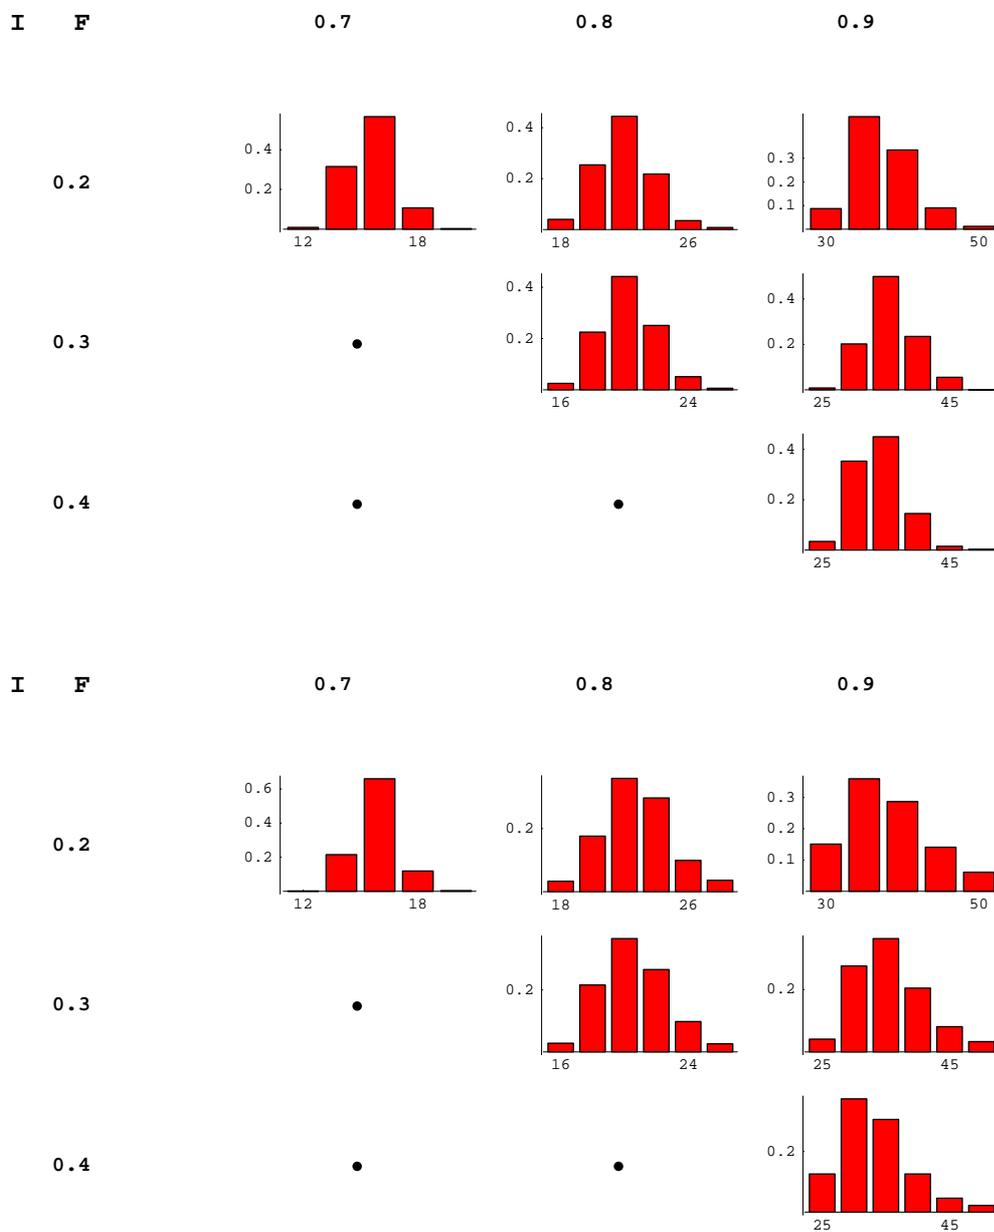


Figure 2.26: *Actual (top) and predicted (bottom) waiting times for evolution to higher fitness. For fitness classes for which no correlation statistics are obtained, mutant fitness is assumed to be normally distributed with mean and variance linearly depending on parent fitness. The accuracy of the predictions shows that estimating correlation information is possible.*

to other landscapes with widely different structural features. In certain respects, the TSP fitness landscape is exceptional; in particular, neighbors that have exactly the same fitness, so-called neutral neighbors, are occurring very rarely. The impact of neutral networks on TSP landscapes is thus negligible (Huynen *et al.* 1996). In the next section we therefore turn our attention to fitness landscapes that allow for tuning the degree of neutrality.

2.2 NKp fitness landscapes

The family of NK landscapes was first introduced by Kauffman (1993) and later extended to the family of NKp landscapes (Weinberger 1990, Fontana *et al.* 1993). The original idea was to envisage a bit string of length N as a genotype. All loci (positions on the genotype vector) therefore carry one of the alleles (entries) 0 or 1. For each locus i of the genotype, K different other loci are chosen at random, which are then called epistatically linked to locus i . The fitness contribution of locus i thus depends on the entries of K other loci and on the locus i itself. To each of the resulting 2^{K+1} combinations a random number is assigned from the interval $[0, 1]$, which determines the contribution to fitness by locus i . Once the fitness table for each locus and for each combination of linked loci is fixed, the fitness contribution of every locus of a genotype can be determined. The final fitness value of the genotype is the sum of the fitness contributions of all its loci, divided by N .

Mutating a genotype of the NK (or NKp) model is changing a random vector entry from 1 to 0 and vice versa. A mutation is called neutral if it does not change the genotype's fitness, the corresponding two neighbors in genotype space are called neutral neighbors. Accordingly a fitness landscape is called neutral, if a substantial fraction of all possible mutations is neutral. From the construction of NK landscapes it is clear that practically no mutation is neutral, which is true for TSP landscapes as well.

Adding the parameter $p \in [0, 1]$ is a proximate way of incorporating neutrality into the model: The contribution of a specific combination of loci is 0 with probability p , and is assigned randomly from the range $[0, 1]$ with probability $1 - p$. This NKp model is even easier to justify biologically than the original NK model. Many combinations of alleles do not influence a genotype's fitness. The probability for neutral neighbors on the resulting NKp fitness landscape can now be adjusted by simply changing the parameter p . For $p = 0$, the landscape corresponds to a normal NK landscape. With

increasing p , neighboring genotypes are more likely to have the same fitness and thus to belong to a whole network of neutral genotypes. For very high p the landscape consists of only few but very extended neutral networks and for $p = 1$ we get a completely flat landscape.

Averaged over all NK p landscapes with fixed N , K , and p , the fitness contribution of Np loci per genotype is zero. As for the remaining $N(1 - p)$ loci the contribution is a random number following a uniform distribution over the interval $[0, 1]$, the mean fitness over all NK p landscapes with fixed parameters N , K and p is given by:

$$f_{mean} = \frac{1}{N} \frac{N(1-p)}{2} = \frac{1-p}{2}$$

The probability for a neutral mutation p_{neutr} , which is the probability that two neighbors in genotype space belong to the same neutral network is calculated as follows: We assume that a genotype is mutated at locus i . This mutation is neutral if the fitness contributions of all loci remain the same. The probability that two of the randomly generated fitness contributions $\neq 0$ are equal is assumed to be zero. Otherwise the following equations only hold almost surely. The fitness contribution of locus $j \neq i$ does not change if either j is not epistatically linked to locus i , which has probability

$$p_1 = 1 - \frac{K}{N-1}$$

or if it is linked to locus i , but the contributions of the combinations before and after mutation are 0. This happens with probability

$$p_2 = \frac{K}{N-1} p^2.$$

The contribution of locus i itself stays the same if the contributions of the combinations before and after mutation are 0, which has probability $p_3 = p^2$. The probability for a mutation to be neutral is now given by

$$\begin{aligned} p_{neutr} &= (p_1 + p_2)^{N-1} p_3 \\ &= \left(1 + \frac{K(p^2 - 1)}{N-1}\right)^{N-1} p^2. \end{aligned}$$

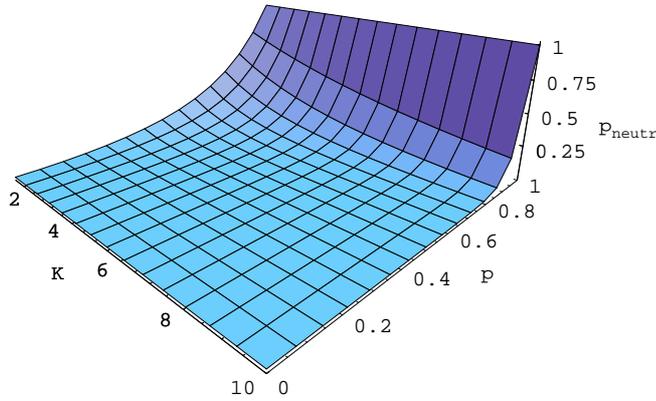


Figure 2.27: Probability p_{neutr} for neutral mutation in NKp landscapes for $N = 40$ and different K and p .

The expected number of neutral mutants per genotype is Np_{neutr} , but the actual number depends on the fitness of the genotype. If $n(g)$ is the number of neutral mutations for genotype g with fitness $f(g)$, and $z(g)$ is the number of loci on g which make zero contribution to the fitness of g we get that $f(g) \leq 1 - \frac{z(g)}{N}$. For a mutation of g at locus i to be neutral it is necessary that the fitness contribution of i is zero. Thus we get

$$n(g) \leq z(g) \leq N(1 - f(g)),$$

which implies that genotypes with higher fitness have fewer neutral neighbors. For high p , a mutant of a genotype with high fitness is likely to have more zeros in its fitness table and therefore much lower fitness. In this sense NKp fitness landscapes become more rugged with increasing p and at the same time more neutral. The autocorrelation function (see Section 1.4), which is a measure of a landscape's ruggedness, is not capable of describing these combined effects: It is apparently nearly invariant with respect to the parameter p (Barnett 1997). In contrast to that, the correlation function presented here is in principle able to capture the interplay between ruggedness and neutrality. In the following we will observe that for low degrees of neutrality, one-dimensional correlation statistics succeed in predicting waiting times. For fitness landscapes involving higher neutrality, one-dimensional correlation statistics do not result in sufficiently accurate predictions. In Chapter 3, however, higher-dimensional correlation statistics prove to provide the essen-

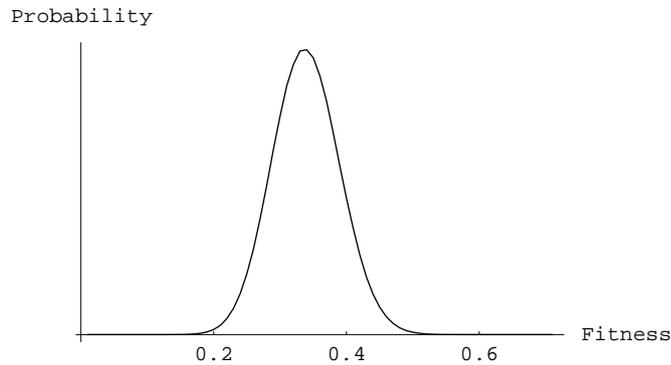


Figure 2.28: *The fitness distribution in an NKp landscape of low neutrality with $N=40$, $K=3$, and $p=0.3$ is approximately normal, and has a mean of about 0.35.*

tial information for approximating evolutionary algorithms even on highly neutral landscapes.

In the subsections below, we present two types of NKp landscapes with different degree of neutrality. As in the case of the TSP, I compare evolutionary waiting times of the actual algorithm with those predicted by a simple stochastic process based on one-dimensional correlation statistics.

2.2.1 Low neutrality

To examine an NKp landscape with very low probability of neutral neighbors, we choose the following parameters: $N = 15$, $K = 3$, and $p = 0.3$. For the parameters chosen, $f_{mean} = 0.35$, and as the fitness of a genotype is on average a sum of $N(1 - p) = 28$ random numbers, fitness in this NKp landscape is approximately normally distributed according to the central limit theorem, similar to fitness distribution in TSP landscapes, see Figure 2.28.

In order to compare actual waiting times on this NKp landscape with predictions based on one-dimensional correlation statistics, the following setting of an evolutionary algorithm is chosen. As selection operator, the stochastic scheme already applied for the second TSP (see last section) is chosen. An individual is therefore taken over to the next generation with probability proportional to its fitness. The selection coefficient is set to $S = 10$, and

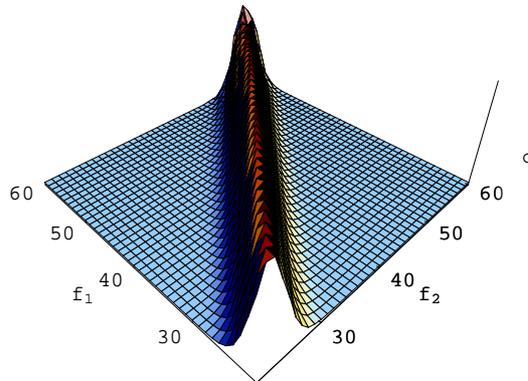


Figure 2.29: *The one-dimensional correlation matrix of NKp genotypes ($N=40$, $K=3$, $p=0.3$) for fitness classes 20 - 60. This range corresponds to the fitness interval 0.2 - 0.6. The matrix shows the same property of emphasis on the diagonal as did the various TSP landscapes.*

the population size is chosen at $P = 5$. The matrix of actual waiting time distributions between 6 pairs of initial (0.2, 0.3, and 0.4) and final fitness values (0.3, 0.4, and 0.45) is presented in Figure 2.31. Each waiting time distribution is derived from 1000 runs of the evolutionary algorithm.

As it was done for the various TSP landscapes in the previous sections, the one-dimensional correlation matrix of this NKp landscape is estimated, see Figure 2.29, by producing random genotypes and their neighbors. The qualitative features of the one-dimensional correlation information are similar to those observed for the TSP landscapes. If a genotype is mutated, e.g. at locus i , the probability that the fitness contribution of locus $j \neq i$ changes is $\frac{K}{N-1}(1-p^2)$: i is linked to j with probability $\frac{K}{N-1}$, and locus j changes its contribution, if it is not zero before and after the mutation, happening with probability $1-p^2$. The probability that the contribution of locus i itself changes under mutation is $1-p^2$. For the current parameter setting due to mutation therefore on average $(K+1)(1-p^2) = 3.64$ bits of a genotype change their fitness contribution. Accordingly, mutant fitness is normally distributed in a good approximation, see Figure 2.30.

As a polymorphic correlation approximation provided nearly accurate predictions of the evolutionary dynamics on TSP landscapes, I use the same approach for this NKp landscape, in order to test the quality of a landscape

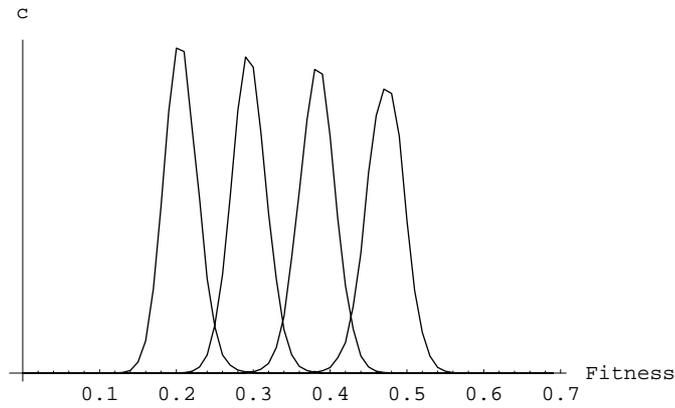


Figure 2.30: *Distribution of mutant fitness for parent fitness 0.2, 0.3, 0.4, and 0.5 shows that the fitness of mutants follows normal distributions in a good approximation. See Figure 2.29 for the matrix of correlation statistics from which these distributions are taken out.*

description based on one-dimensional correlation statistics. Individuals are therefore reduced to their fitness values only, and a stochastic process is constructed, where individuals are mutated to individuals with new fitness values, according to the probabilities provided by one-dimensional correlation information. As the selection scheme is based on fitness values of individuals only, it is applied in the same way the evolutionary algorithm does. Comparing the actual waiting time distributions with those predicted by this stochastic process, shows encouraging results. For low degrees of neutrality the one-dimensional correlation matrix again seems to provide enough information to characterize evolution on NKp fitness landscapes, see Figure 2.31.

Actual (first number) and predicted mean values of waiting times for the distributions shown in Figure 2.31:

initial / final fitness	0.3	0.4	0.45
0.2	3.3 / 3.3	7.9 / 7.4	12 / 10.7
0.3	*	5.1 / 4.7	9.1 / 7.8
0.4	*	*	4.2 / 3.7

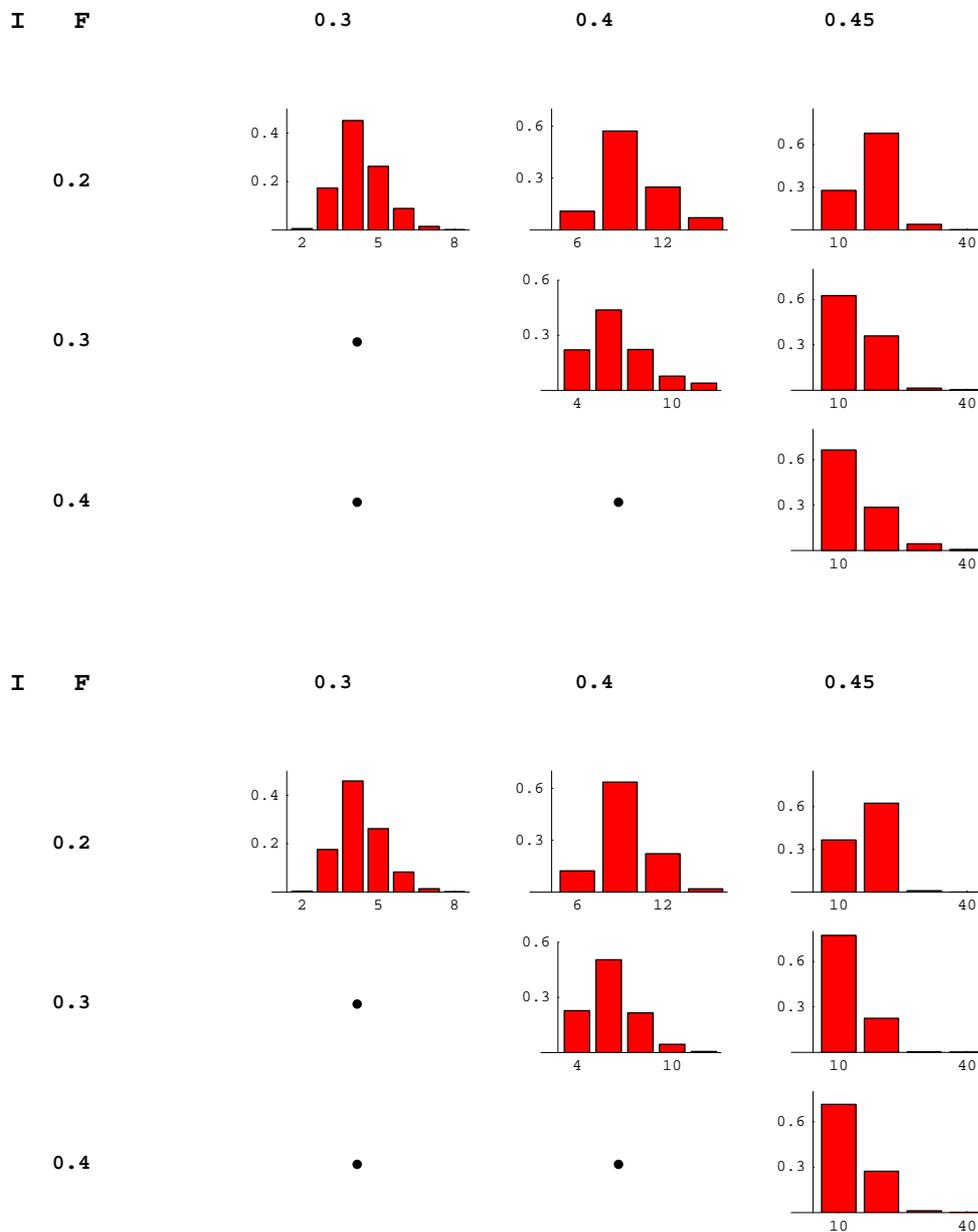


Figure 2.31: *Actual (top) and predicted (bottom) waiting times for evolution on a NKp landscape with low degree of neutrality ($p = 0.3$). The similarity of both results is again very satisfactory.*

2.2.2 High neutrality

In this section, a one-dimensional correlation-based description is tested for an NKp fitness landscape with high neutrality. The parameters are set to $N = 30$, $K = 5$, and $p = 0.99$, resulting in a landscape with $p_{neutr} = 0.89$. The fraction of individuals having zero fitness is very high. $p^N = 74\%$ of all genotypes belong to this class, fitness in this landscape is thus far away from being normally distributed. The mean fitness over all NKp landscapes with this parameter setting is given by $\frac{1-p}{2} = 0.005$. As in the sections before, the one-dimensional correlation matrix is estimated by mutating random genotypes of different fitness classes. The matrix shows that this highly neutral landscape has very different properties from those observed for the landscapes analyzed so far. The probability for a neutral mutation with mean 0.89 varies according to class frequencies. As fitness class 0 is much bigger than all the others, the neighbors of individuals of all fitness classes have a disproportionate high chance to belong to this class. On average, a mutation changes the fitness contributions of only $(K + 1)(1 - p^2) = 0.12$ bits per genotype vector. The distribution of mutant fitness therefore does not resemble a normal distribution as it was the case for TSP and the NKp landscapes with low neutrality.

Despite these differences we construct a stochastic process in which individuals are reduced to their fitness only and where the correlation matrix serves as the basis for the mutation operator. The outcome of 2000 runs of this stochastic process is used to predict the distribution of waiting times. A comparison to the actual distribution of the evolutionary algorithm shows that one-dimensional correlation information is insufficient to provide accurate predictions of the observed dynamics. Even though the scale of waiting times is captured, an error of predictions of about 50% indicates the need for augmenting the correlation information used so far.

A comparison of actual (first number) and predicted mean values of waiting times for the distributions shown in Figure 2.34 shows that waiting times are underestimated by predictions based on one-dimensional correlation:

initial / final fitness	0.02	0.03	0.04
0.01	5.5 / 3.5	9.2 / 5.2	14.7 / 7.2
0.02	*	4.2 / 2.4	10.1 / 4.6
0.03	*	*	8.0 / 3.7

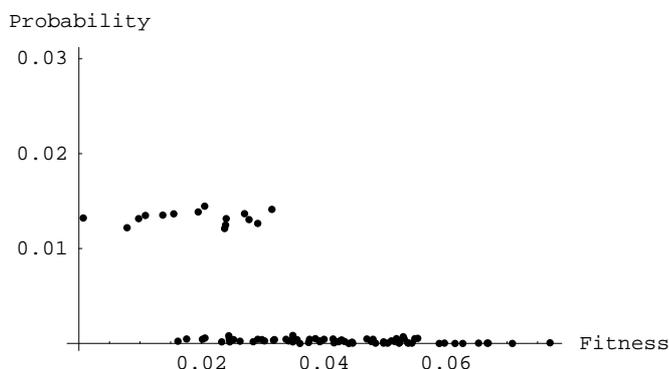


Figure 2.32: In an NKp landscape with high degree of neutrality ($N=30$, $K=5$, and $p=0.99$) the fraction of genotypes with zero fitness is very high ($\approx 75\%$) and genotypes hardly ever have fitness greater than 0.06. In this graph, all different fitness values (less than 100), except fitness 0, that were found when testing 100000 random genotypes are depicted. The landscape thus composes into few but very extended neutral sets. The features of this landscape are completely different from the properties of NKp landscapes with low neutrality.

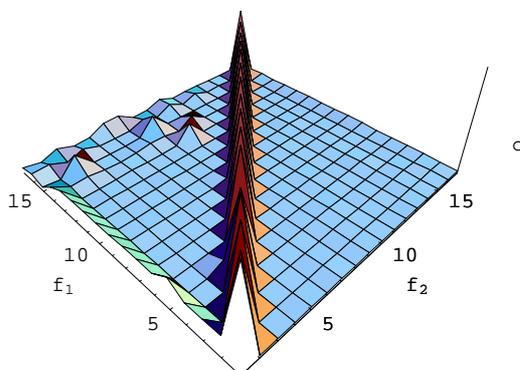


Figure 2.33: The correlation matrix for an NKp landscape with $p=0.99$ shows the high probability for neutral neighbors, and also the relatively high chance for individuals of all different fitness classes to have mutants belonging to class 0, which corresponds to fitness 0.

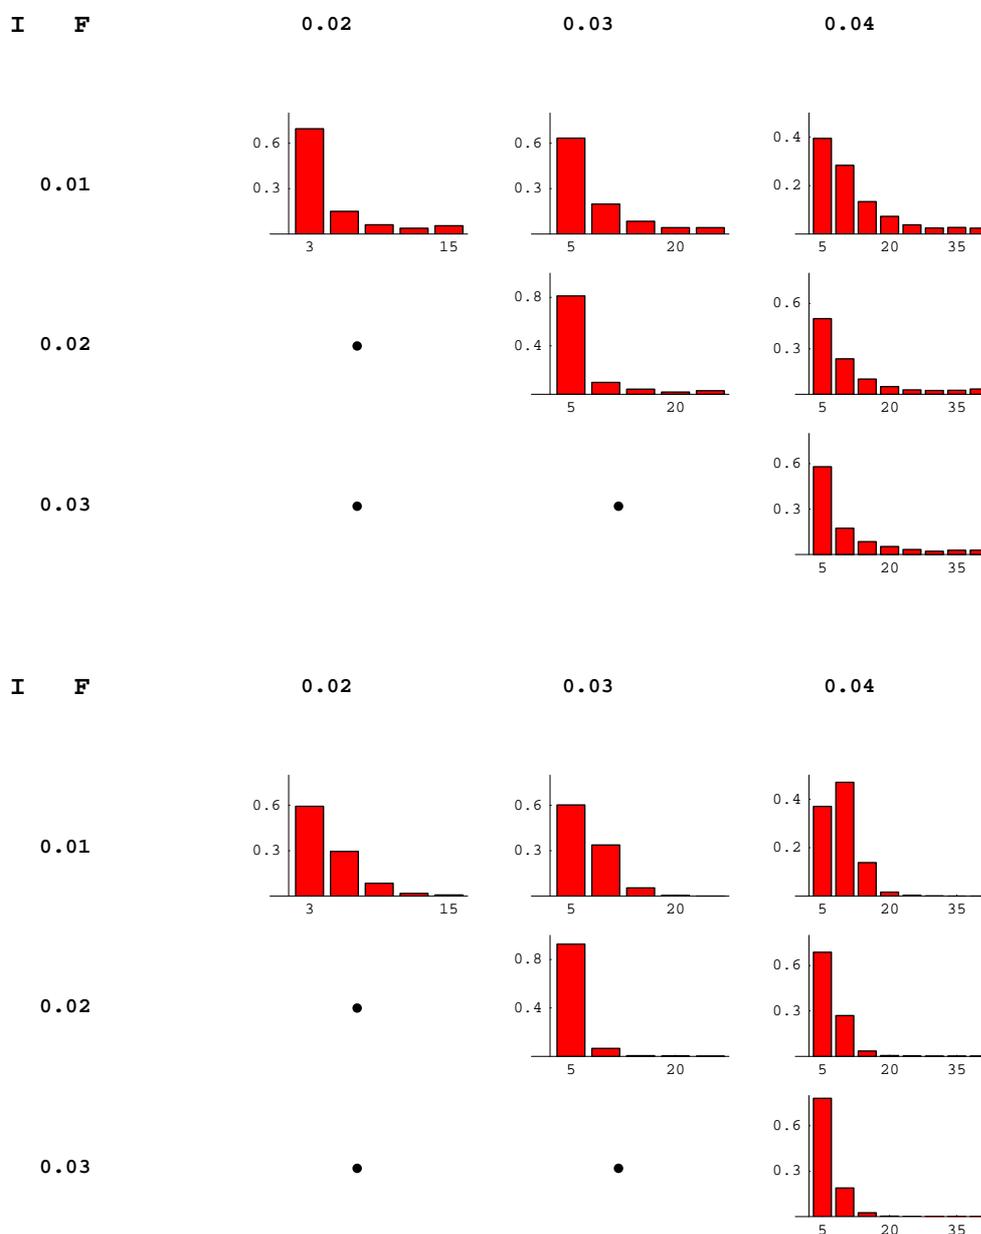


Figure 2.34: *Actual (top) and predicted (bottom) waiting times for evolution on an NKp landscape with a very high degree of neutrality ($p = 0.99$). The comparison shows that actual evolution takes many more generations (about twice as many) to attain a final fitness class; the many steps within one fitness class that probably result from diffusion along neutral networks are not considered by the one-dimensional correlation approximation.*

A refinement of the correlation matrix achieved by introducing smaller fitness classes did not result in more accurately predicted waiting times.

For landscapes with a very high degree of neutrality, a one-dimensional correlation approximation may be insufficient. The number of different fitness values is much smaller than the number of possible genotypes. Due to the high-dimensionality of genotype space, sets of genotypes with approximately equal fitness tend to form networks that are connected by single mutational steps. In contrast to the rugged landscapes of combinatorial optimization, these fitness landscapes decompose into sets of neutral networks that are entangled in a very complicated and unstructured way. On those neutral networks selection is completely disabled and neutral evolution dominates.

The consequences of neutral evolution are difficult to understand and one-dimensional correlation statistics cannot specify these dynamics in detail. Neutral networks are too big and too inhomogeneous to be described by the mean neighborhood structure of individuals. A one-dimensional correlation description neglects differences of neighborhood structures within one fitness class. Such differences are expected to occur as a population drifts from the entrance point of a neutral network to an exit, from where it can reach a higher fitness class. Even if the probability for finding a better neighbor is very high at a neutral network's exit, the population probably still needs to evolve for several generations to reach such 'portals'. A one-dimensional correlation approximation does not account for this subtlety and can thus provide only mediocre estimates of evolutionary waiting times.

An improvement of the predictions requires incorporation of new statistics of the landscape, allowing to describe a population's drift along neutral networks.

2.3 The Royal Staircase Fitness function

The fitness function studied in the last sections are quite complicated. Special properties of these functions are difficult to identify without a detailed investigation of the fitness landscapes.

Designing a fitness function in such a way that all its features are evident and analytically available has many consequences. In particular, there is no longer a need to run an evolutionary algorithm, because the optimal solution is known from the beginning. In order to improve our understanding of evolutionary algorithms, it might nonetheless be beneficial to study fit-

ness functions of this type. Even on such well-structured fitness landscapes, the evolutionary dynamics itself are often quite elaborate. Moreover, a detailed knowledge of the structure of the fitness function enables us to make predictions, which are based on analytically derived and therefore accurate correlation statistics.

The class of Royal Staircase Fitness (RSF) functions (van Nimwegen *et al.* 1999, van Nimwegen and Crutchfield 1999, 2000) allow a specific design of fitness landscape architecture. RSF functions make it possible to define fitness landscapes which decompose into sets of few but very extended clusters of genotypes having the same fitness, i.e. neutral genotypes. RSF landscapes therefore permit the study of neutral evolution, which is observed in several natural evolutionary systems as well. Experiments with bacteria colonies and the study of RNA molecule folding have shown that neutrality among genotypes is a frequent phenomenon. Evolutionary dynamics on highly neutral fitness landscapes are difficult to conceive. This process of evolution involving neutrality is often referred to as epochal evolution, as long periods of stasis are interrupted by sudden innovations when the population discovers genotypes of higher fitness (van Nimwegen 1999).

In the following I demonstrate how the dynamics of neutral evolution on RSF landscapes can be approximated by using correlation measures.

2.3.1 The fitness function

The Royal Staircase Fitness (RSF) function involves large-scale neutrality, and is designed in such a way that one has a detailed knowledge about the neutral sets. The definition of the Royal Staircase Fitness function and the utilized evolutionary algorithm (van Nimwegen 1999) are presented in the following:

- Genotypes

A binary string of length $L = NK$ is seen as a genotype. Starting from the left we call the first K bits the first block, the next K bits the second block and so on. The genotype thus consists of N blocks having K bits each. A block is called aligned if all K bits are ones.

- Fitness

To calculate the fitness of genotype g , the number $I(g)$ of ones, starting from the left, until the first zero appears are counted. The fitness f is

then given by

$$f(g) = 1 + \lfloor \frac{I(g)}{K} \rfloor.$$

The fitness values are thus integer numbers ranging from 1 to $N + 1$. A genotype having the first block unaligned has fitness 1, a genotype with the first block aligned but the second unaligned has fitness 2.

$$\underbrace{\overbrace{1111\dots1111}^K \overbrace{1001\dots0111}^K}_{NK} 01001\dots\dots\dots011101001110110$$

This genotype has fitness 2. The single optimal genotype is the bit string consisting of only ones. It is easy to see that the dynamics of the evolutionary algorithm are invariant under permutations of the bit string representation of genotypes, if the fitness function is appropriately modified.

- Mutation

To mutate a genotype, each bit is independently changed from 0 to 1 or vice versa with probability q . Under mutation, an individual has a certain chance to stay the same or to change, even to the optimal genotype.

- Selection

To keep the population size P constant, we choose P times from the set of all individuals, such that each genotype is selected with probability proportional to its fitness. Individuals with higher fitness therefore have a greater chance to be taken over to the next generation than those of low fitness. Individual $i \in \{1, \dots, P\}$ with fitness f_i is chosen with probability $c f_i$. The constant c can be easily determined as $c \sum_{j=1}^P f_j = 1$.

- Evolutionary waiting times

We produce a random string with a certain initial fitness and copy it P times. Evolution starts with a population of P individuals all having the same genotype. In every generation we mutate each genotype, evaluate the new fitness and then select P individuals from the mutated population. Over generations, genotypes get higher fitness and have

more blocks aligned. The process stops if one individual has reached a certain final fitness. The number of generations necessary to reach this fitness threshold is the waiting time of this single run. Several runs then give the distribution of waiting times between an initial and a final fitness value.

The definition of Royal Staircase Fitness function has several implications for the structure of the fitness landscape. The number of different neutral sets is $N + 1$, having fitness $1, \dots, N + 1$. The proportion of genotypes with fitness n is $2^{-K(n-1)} - 2^{-Kn}$, which implicates that the networks of higher fitness are smaller than those of lower fitness.

Genotypes with fitness n are more likely to attain fitness $n + 1$ under mutation than strings with fitness $n - 1$. The probability that mutation increases a genotype's fitness by more than one is very low. The landscape architecture thus forces the population to take practically all of the neutral networks, which can therefore be regarded as the royal stairs, before the optimal genotype is found.

Mutation can also destroy the blocks of an individual. Genotypes may eventually even fall back to fitness 1 if mutation changes a bit of the first block. To fall off a neutral network to a lower one happens quite often to individuals evolving on an RSF fitness landscape, whereas to find the portal to a higher network is much more difficult.

2.3.2 The evolutionary algorithm

To observe population dynamics on the landscape of the Royal Staircase Fitness function, we choose a particular parameter setting: $N = 8, K = 8, P = 200$, and $q = 0.006$.

On average, a population of 200 individuals needs about 2500 generations until the optimal genotype is found. Long periods of stasis are followed by sudden innovations when one individual discovers a network of higher fitness, see Figure 2.35. If this individual is not lost due to the stochastic selection scheme, it may spread and finally drag a great proportion of other individuals onto this network. Then again the best fitness stays the same for many generations, until the next portal that leads to a network of higher fitness is discovered. The higher individuals climb on the fitness landscape, the more difficult it becomes to find such exits from the current cluster of isofitness.

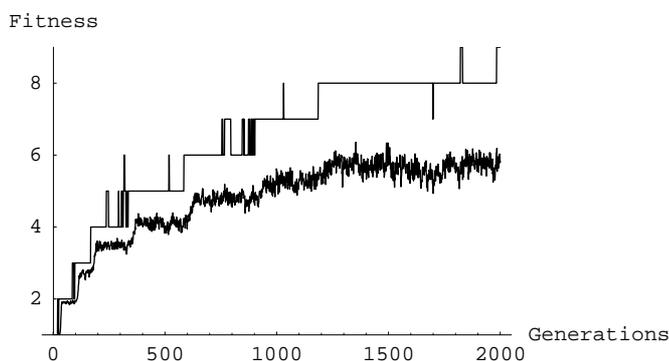


Figure 2.35: *Best and mean fitness in a population of an RSF evolutionary algorithm with $N = 8$, $K = 8$, $P = 200$, and $q = 0.006$, evolving over generations. Periods of stasis are followed by sudden innovations when an individual discovers a network of higher fitness. Due to the stochastic selection scheme, individuals with highest fitness in the population may also be lost.*

In order to provide target statistics for predictions based on reduced landscape descriptions, we construct the actual waiting time distributions of the evolutionary algorithm between several initial and final fitness thresholds. In particular, we choose 1, 5, and 7 as initial and 5, 7, and 9 as final values and only consider waiting times to higher fitness. For each of the resulting 6 pairs of initial and final fitness, the distribution of waiting times is derived from 1000 runs of the actual evolutionary algorithm, see Figure 2.37.

The neighborhood structure of individuals will be the basis for a reduced description of the RSF landscape, as such an approximation has proved to be successful for TSP and NKp landscapes, see last sections. The evolutionary algorithm I use for the Royal Staircase Fitness function involves stochastic mutation, depending on the per bit mutation probability q . For this reason, difficulties with the simple notion 'neighbor' arise: Under the mutation operators studied so far, every genotype has the same fixed number of neighbors, and every link between neighboring genotypes is weighted with the same probability, simply depending on the number of neighbors per genotype. Under a mutation operator depending on a per bit mutation probability, all genotypes in the RSF landscape are neighbors, as mutation can in principle convert one bit string into any other. Every such link between

neighbors is weighted with a certain probability. It is therefore no longer possible to define the correlation between two fitness values f and f' as the probability for a random genotype of fitness f to have a neighbor with fitness f' . An overall definition is: Correlation between fitness values f and f' is the probability for a random genotype of fitness f to end up as a (neighboring) individual with fitness f' after mutation.

A division of fitness values into classes is not necessary in the case of the RSF function, as there are by definition only $N + 1$ different fitness values. The correlation matrix $C = (c_{f,f'})$, $f, f' \in \{1, \dots, N + 1\}$, where $c_{f,f'}$ is the probability that a random string with fitness f is changed to a string of fitness f' under mutation, can now be calculated analytically.

- If f' is smaller than f , the first $(f' - 1)K$ bits have to remain unchanged, whereas in block f' at least one bit has to be mutated in order to get this block unaligned. Thus:

$$c_{f,f'} = (1 - q)^{(f'-1)K} (1 - (1 - q)^K) \quad \text{for } f' < f$$

The probability for decreased fitness after mutation is high compared to the probability for increased fitness. The first case simply means changing any bit of the aligned blocks, whereas the second case requires full alignment of at least one more block of K bits.

- If f and f' equal $N + 1$, $c_{f,f'}$ is the probability that no bit of the whole string is mutated:

$$c_{f,f'} = (1 - q)^{NK} \quad \text{for } f = f' = N + 1$$

The probability p_a that a random string gets an unaligned block aligned under a single mutation process is the probability that at least one bit is changed, and that the resulting string is then aligned:

$$p_a = \frac{1 - (1 - q)^K}{2^K - 1}$$

- If f equals f' , but f' is smaller than the highest possible fitness, all $f - 1$ blocks have to remain unchanged and the f th block is not aligned:

$$c_{f,f'} = (1 - q)^{(f-1)K} (1 - p_a) \quad \text{for } f = f' < N + 1$$

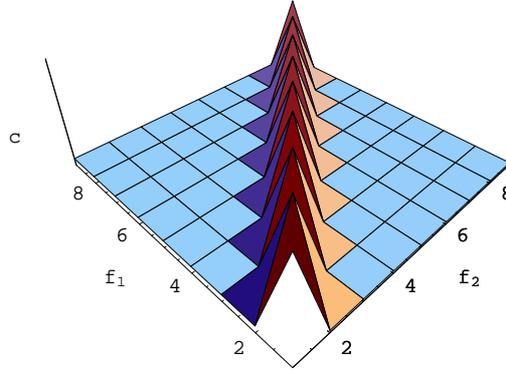


Figure 2.36: *The correlation matrix for the RSF landscape with $N = 8$ and $K = 8$ shows the high probability for neutral mutation.*

- If f is smaller than $f' = N + 1$, $f - 1$ blocks should stay unchanged, the f th block, which has been unaligned, has to become aligned, and the remaining $N - f$ blocks, which are each in any of the 2^K configurations have to be in the only aligned one:

$$c_{f,f'} = (1 - q)^{(f-1)K} p_a \left(\frac{1}{2^K} \right)^{N-f} \quad \text{for } f < f' = N + 1$$

- Finally, for an increase in fitness from f to $f' \neq N + 1$, the f th block has to change from unaligned to aligned, $f' - f - 1$ blocks have to be in the aligned configuration, and the f' th block has to be of unaligned form. This happens with probability:

$$c_{f,f'} = (1 - q)^{(f-1)K} p_a \left(\frac{1}{2^K} \right)^{f'-f-1} (1 - 2^{-K}) \quad \text{for } f < f' < N + 1$$

Using this $(N + 1, N + 1)$ correlation matrix, see Figure 2.36 we can define a simplified evolutionary process where individuals are described only by their fitness. According to the probabilities given by the matrix, individuals are mutated and possibly get another fitness value. Selection works in the same way as for the actual evolutionary algorithm. The outcome of 1000 runs of this stochastic process is combined to the predicted distribution of waiting

times and compared to the actual distribution, see Figure 2.37. Actual and predicted mean values of waiting times for the distributions shown in Figure 2.37:

initial / final fitness	5	7	9
1	618 / 248	1304 / 635	2595 / 1513
5	*	545 / 283	1873 / 1130
7	*	*	960 / 592

The results are encouraging, the scale of evolutionary waiting times is predicted in the right way. The actual process is always about twice as slow as one-dimensional correlation statistics predict, but the main features of the dynamics are approximated very well.

Like in the case of NKp landscapes with high neutrality, it is questionable what assumption underlying this pair-correlation-based description is violated and thus is responsible for the discrepancies. For the evolutionary algorithm on the RSF landscape, it was assumed that all individuals of a certain fitness have the same probability of getting the next blocks aligned. This is not true for those which have evolved over several generations. We have seen that individuals may easily fall back to very low fitness after mutation. These individuals have a lot of bits set correctly, the assumption that they have the bits set randomly results in inaccurate predictions. The number of correctly set bits influences the probabilities for coming back to a higher fitness level.

To give a more precise description of the process we have to incorporate the fact that different genotypes of the same fitness have a different number of bits set correctly in their unaligned blocks. This idea is the equivalent to the idea I will use for NKp landscapes with high neutrality, where I try to incorporate that individuals of the same fitness have different neighborhood structures. Those studies are presented in the following chapter.

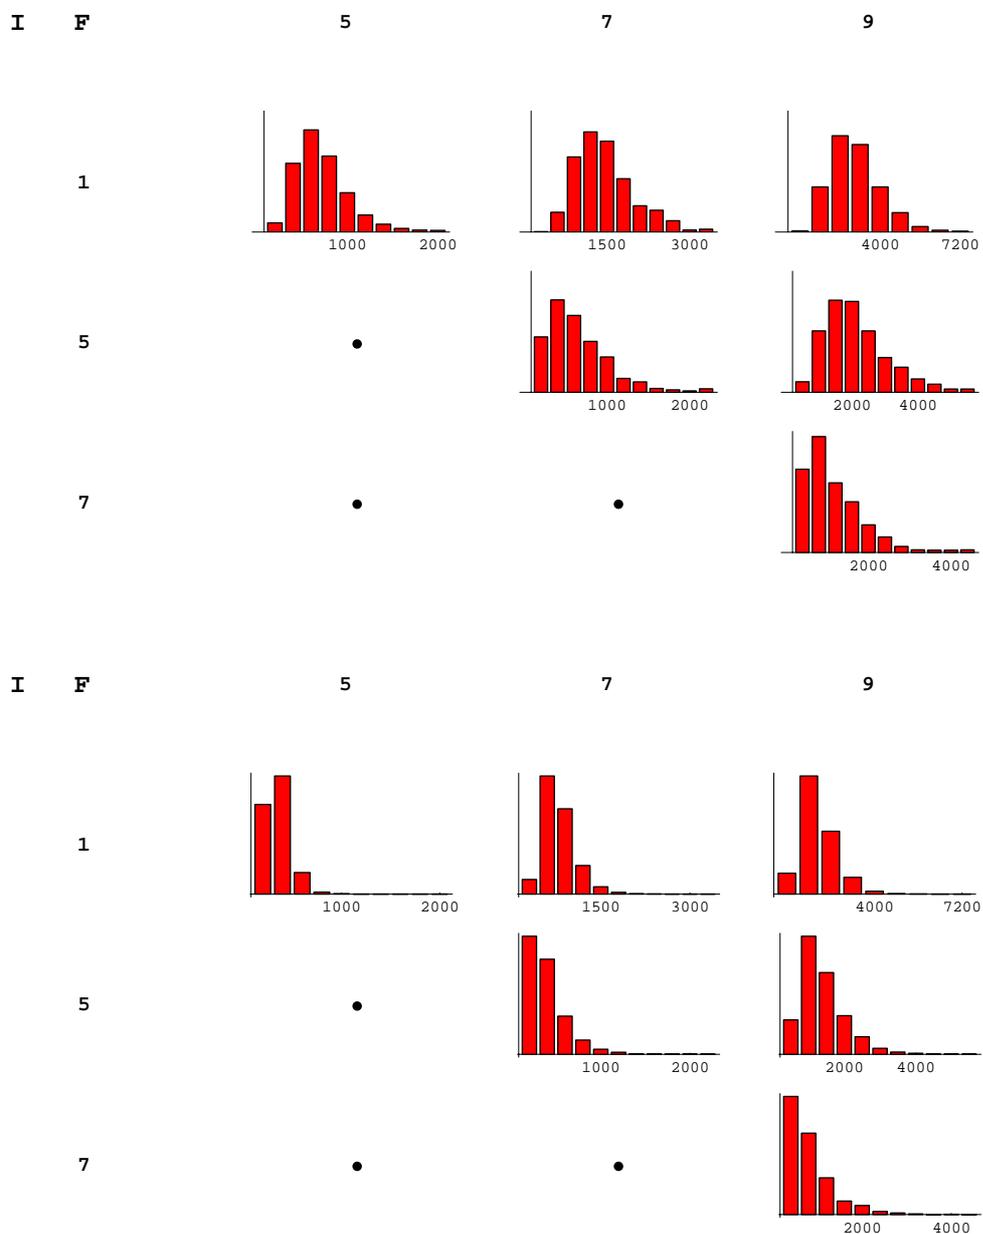


Figure 2.37: *Actual (top) and predicted (bottom) waiting time distributions of an evolutionary algorithm on an RSF landscape, for six pairs of initial and final fitness values. Predictions are based on one-dimensional correlation information. Actual waiting times are underestimated by about 50%.*

Chapter 3

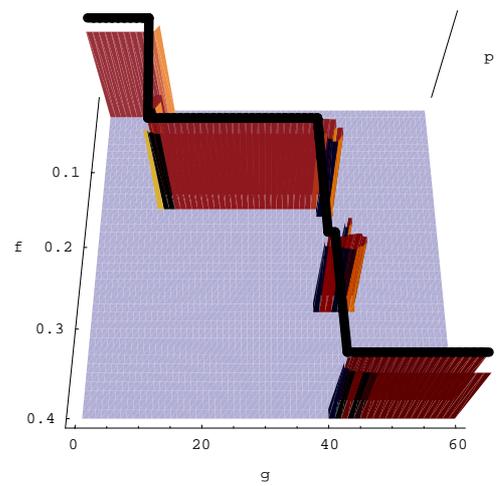
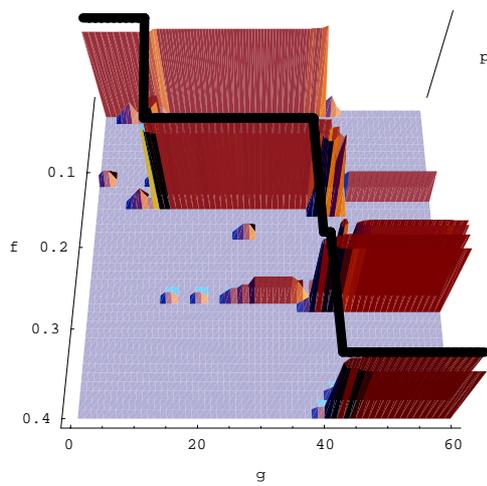
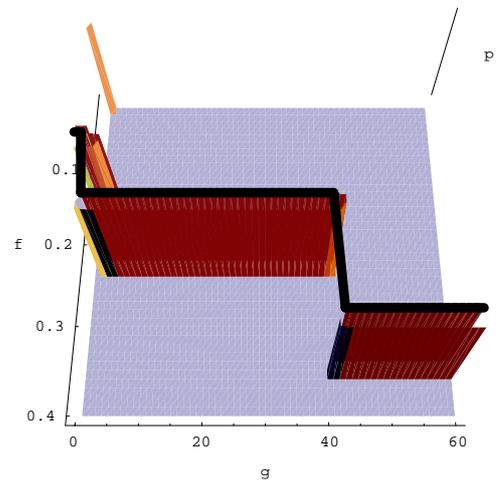
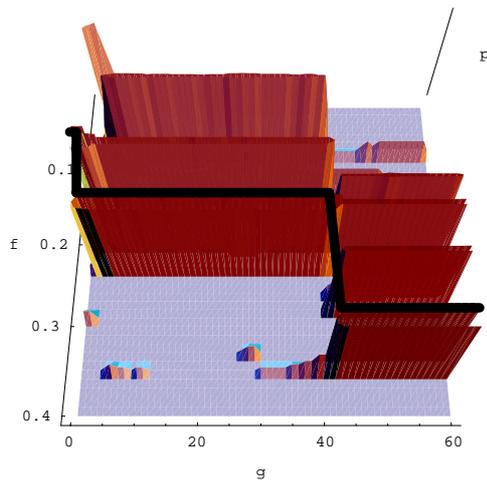
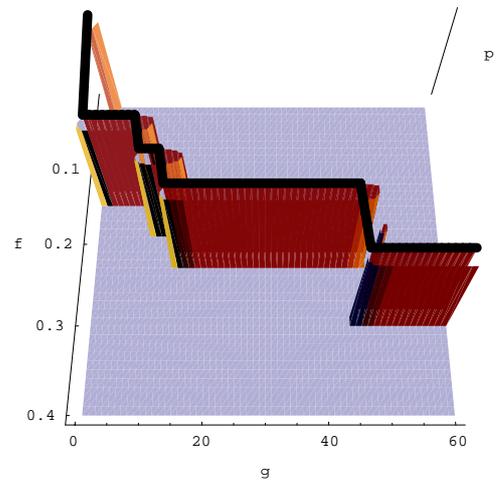
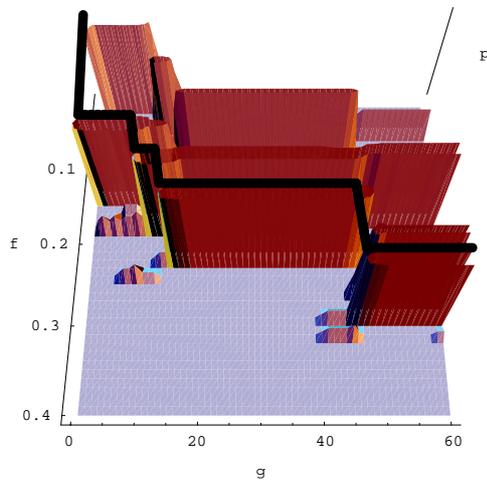
Higher-dimensional correlation

In the last chapter we observed that waiting time predictions based on one-dimensional (or pair) correlations alone are not sufficiently accurate under conditions of high neutrality. Pair-correlation-based theory assumes that all individuals of a given fitness class have identical fitness distributions of mutants. This assumption is likely to be inaccurate when extended neutral clusters exist. 'Interior nodes' of a given cluster have a very low probability for links to clusters of higher fitness, whereas for 'boundary nodes' such links are more probable, see Figure 3.1. Consequently, the different neighborhood structures, and therefore the different distributions of mutant fitness within a given fitness class need to be accounted for in order to describe the dynamics on a neutral network.

3.1 NKp fitness landscapes

3.1.1 Two-dimensional correlation

In Chapter 2 we have seen that for NKp landscapes with very high degree of neutrality a description of the fitness landscape based on one-dimensional correlation statistics is not sufficiently accurate. A comparison of actual and predicted evolutionary waiting times has shown that the time scale of evolution is underestimated by predictions based on this type of correlation information. This shows that it is not as easy for individuals to discover networks of higher fitness as the correlation matrix suggests. Individuals have to evolve from the entrance of a neutral network towards the exit to a cluster



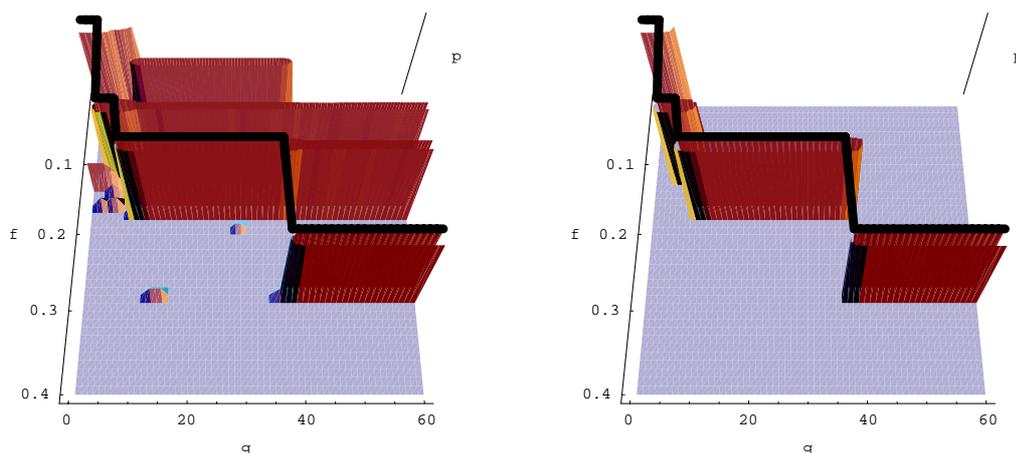


Figure 3.1: *Four different runs of the evolutionary algorithm on the NKp landscape with high neutrality. The left pictures show the actual distribution of fitness of all neighbors the individuals in a population have, and the current best fitness (black line), evolving over 60 generations. The right pictures show the corresponding fitness distribution of the individuals in the population, together with the best fitness. This is an confirmation of the hypothesis that a population evolving on a landscape with a high degree of neutrality searches the neutral cluster until the neighborhood structure changes, thus allowing jumps to a network of higher fitness (e.g. first run: generation 8, 10, and 40). Notice that the changes in neighborhood structure occur before the actual jump to the new cluster takes place. Possible connections to networks of higher fitness can also be lost (e.g. second run: generation 12, third run: generation 21, last run: generation 17) and found again later.*

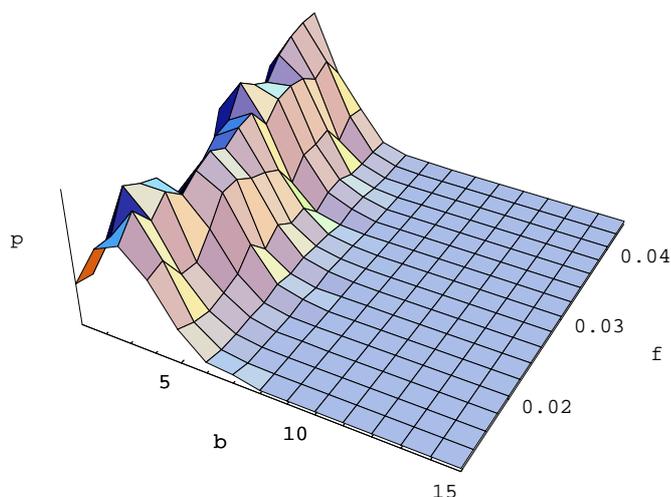


Figure 3.2: *Distribution of number of better neighbors b for genotypes with fitness f . The number of better neighbors gets gradually smaller for individuals of higher fitness.*

of higher fitness, see Figure 3.1. The number of generations necessary to discover higher fitness levels varies a lot from run to run of the evolutionary algorithm. Accordingly, individuals of a certain fitness may have very different neighborhood structures. In order to find sets of individuals for which this structure is sufficiently similar, a single fitness class has to be divided into several new classes. Only then a description of the fitness landscape based on neighborhood structures can be successful. For a division of individuals belonging to a certain fitness class into a set of new classes, crucial statistics about individuals additional to their fitness have to be introduced. Some candidate statistics are the average mutant fitness, the variance of mutant fitness, or the fraction of better neighbors (mutants that have higher fitness than the individual itself). I found that the number of better neighbors (see Figure 3.2) is the attribute best suited for this purpose and devised a refined approximation of the evolutionary algorithm, based on characterizing individuals in terms of the pair (fitness, number of better neighbors).

For the NK p landscape with high neutrality, $N = 30$, $K = 5$, and $p = 0.99$, a $(15 \times 10) \times (15 \times 10)$ correlation matrix is estimated (15 fitness classes, and 10 classes for the number of better neighbors; class 1: no bet-

ter neighbors, class 10: more than 8 better neighbors) by mutating random genotypes. The matrix entry at position $((f, b), (f', b'))$ describes the probability that an individual of fitness class f having b better neighbors results in an individual of fitness class f' and b' better neighbors after mutation.

In order to test this description of the NKp landscape we try to predict the waiting times of the actual evolutionary algorithm in a way similar to the method we applied in the previous sections: By reducing genotypes to the pair (fitness, number of better neighbors) we construct a simplified stochastic process, using the two-dimensional correlation information for mutating individuals, and employing the actual selection scheme, as this is only based on fitness values. Compared to the one-dimensional correlation description, the refined dynamics based on this two-dimensional correlation information give significantly more accurate predictions of waiting times.

Actual (first number) and predicted mean values of waiting times, based on two-dimensional correlation:

initial / final fitness	0.02	0.03	0.04
0.01	5.5 / 3.2	9.2 / 5.1	14.7 / 8.6
0.02	*	4.2 / 3.5	10.1 / 7.0
0.03	*	*	8.0 / 7.0

These results show that a more detailed description of neighborhood structures captures additional important information which determines the waiting times of evolution. But not only the mean value of waiting times is predicted more accurately, also their variances are increased towards the observed values. In some runs, the population happens to find exits to clusters of higher fitness much more quickly than in others, an observation which is true for the actual evolutionary algorithm as well as for the simplified process with two-dimensional correlation.

3.1.2 Three-dimensional correlation

Taking into account more and more features of individuals which influence their neighborhood structure, is expected to result in still more precise predictions. A bit surprising though is that the number of worse neighbors individuals have also plays an important role for the outcome of the evolutionary process.

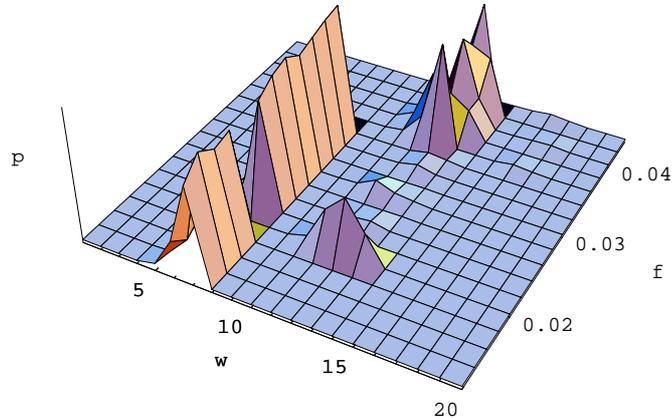


Figure 3.3: *Distribution of number of worse neighbors w for individuals with fitness f . The structure of this graph results from varying class frequencies. Larger classes have an increased probability for neutral mutation and thus a lower probability for worse neighbors.*

To describe the possible neighborhood structures of individuals in an even more accurate way, we now focus on the number of better and worse neighbors an individual of certain fitness has. While the probability for a random mutant to be of higher fitness is steadily decreasing for individuals of higher fitness classes, the probability for a mutant to be worse strongly depends on the probability for neutral mutation, p_{neutr} . The actual probability for neutral mutation in a certain fitness class depends on the fitness value and on the number of individuals that belong to this class. Lower fitness and greater number of individuals belonging to a fitness class results in higher probability for neutral mutation. In NKp landscapes with high neutrality the fraction of individuals which have fitness 0 is very large: $p^N = 0.74$ for our NKp landscape. For this reason the actual probability for a neutral mutation is higher than p_{neutr} for fitness 0. It falls below p_{neutr} with increasing fitness, but varies according to class frequencies. A higher probability for neutral mutation implies a smaller number of worse neighbors.

To predict the within-cluster dynamics in an even more precise way, NKp individuals in highly neutral landscapes are now described by the triple (fit-

ness, number of better neighbors, number of worse neighbors). For the current NKp landscape, this description results in a three-dimensional correlation information, i.e. a $(15 \times 10 \times 15) \times (15 \times 10 \times 15)$ matrix: 15 classes for fitness, 10 classes for the number of better neighbors and 15 classes for the number of worse neighbors. The matrix entry at position $((f, b, w), (f', b', w'))$ is the probability that an arbitrary individual with fitness f , b better and w worse neighbors will end up as an individual of fitness f' , having b' better and w' worse neighbors after a single mutation. This refinement of correlation description should be able to predict a population's drift along a neutral network towards its exit to a cluster of higher fitness in a still more accurate way than it is possible with two-dimensional correlation.

And indeed, the mean waiting times to reach a fitness threshold, predicted by a stochastic process based on the three-dimensional correlation matrix only, are closer to the actual mean values of waiting times, see Figure 3.4, than those predicted by two-dimensional correlation.

Actual (first number) and predicted mean values of waiting times for the distributions shown in Figure 3.4:

initial / final fitness	0.02	0.03	0.04
0.01	5.5 / 4.8	9.2 / 7.5	14.7 / 10.7
0.02	*	4.2 / 3.6	10.1 / 7.4
0.03	*	*	8.0 / 6.8

The number of worse neighbors thus gives additional important information of the evolutionary process. The conclusion drawn from these results is very encouraging: A population's drift along a neutral network can be understood very well, when more and more information is available about the neighborhood structure of individuals. Correlation statistics are not only appropriate for landscapes with low neutrality, but also for highly neutral landscapes. The only difference is that a simple description by means of fitness is not sufficient anymore. Clusters of equal fitness too large and too inhomogeneous for a reduction to fitness only. If we account for the varying neighborhood structure of individuals belonging to the same fitness class, we succeed in describing the drift from the entrance of a neutral network towards an exit that leads to a cluster of higher fitness. The correlation statistics introduced in this work seem to capture the salient features of a fitness landscape, like the degrees of ruggedness and neutrality, simultaneously.

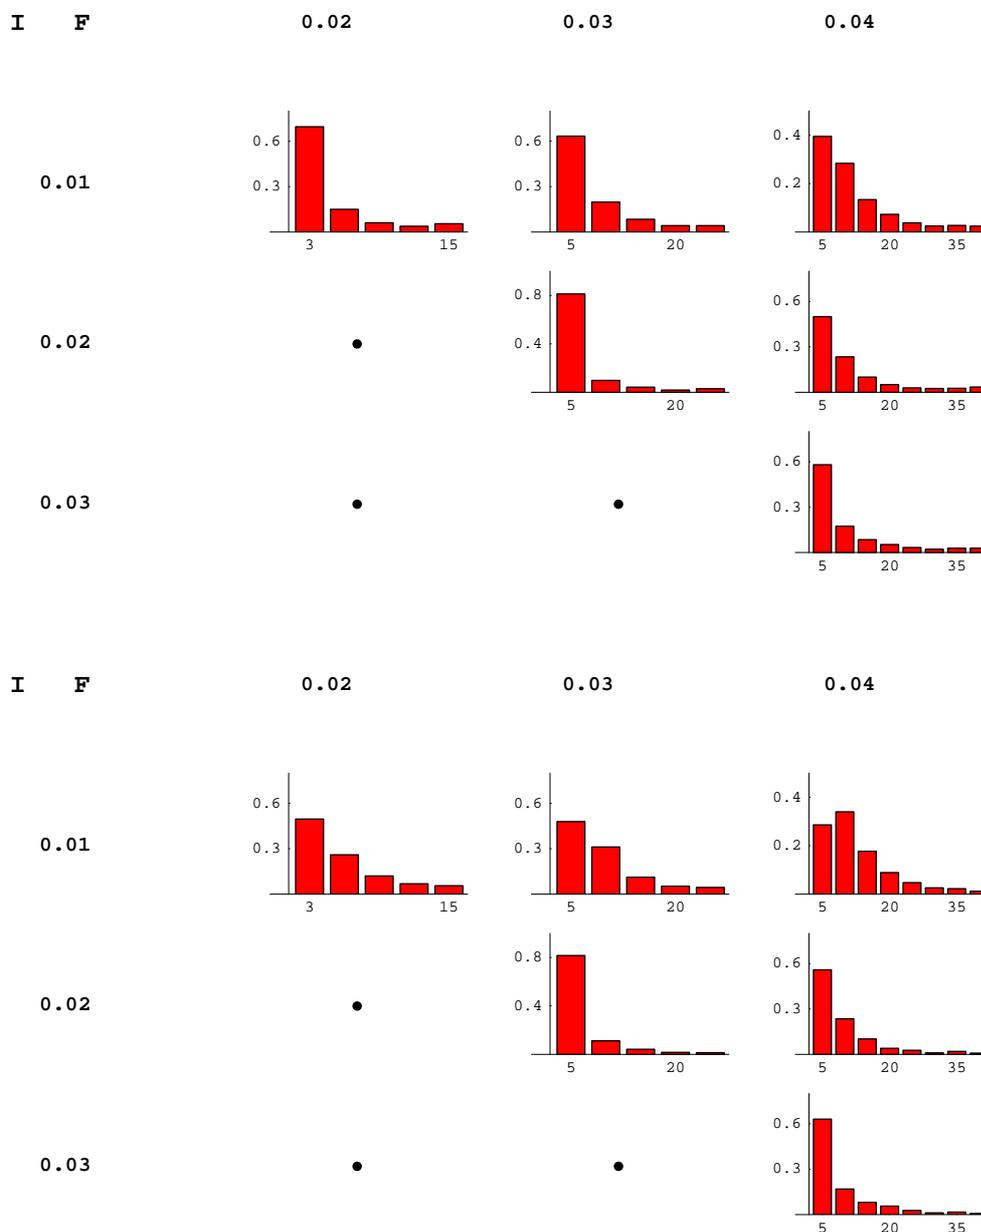


Figure 3.4: *Actual (top) and predicted (bottom) waiting times for evolution on a NKp landscape with a very high degree of neutrality ($p = 0.99$). Predictions are based on three-dimensional correlation statistics and give a very good approximation of evolution along the neutral networks. Compared to one-dimensional correlation statistics the number of better and worse neighbors of a genotype provides crucial additional information.*

3.2 The Royal Staircase Fitness function

3.2.1 Two-dimensional correlation

For highly neutral NKp landscapes the fact was incorporated, that genotypes of equal fitness may have different neighborhood structure, and a two-dimensional correlation matrix was constructed. The equivalent of this approach for RSF landscapes is to account for the fact that individuals of the same fitness have, according to the number of bits set correctly in their unaligned blocks, different probabilities to attain another fitness level under mutation.

As a first step I thus incorporate the number of bits set incorrectly in the first unaligned block, seen from the left. Genotypes are now described by the pair (f, b) where f is the fitness and b is the number of incorrect set bits, i.e. the number of zeros, in the first unaligned block, i.e. the f th block. b ranges from 1 to K , as genotypes of fitness f have at least one bit set incorrectly in their f th block.

A four-dimensional matrix representing this two-dimensional correlation information can now be constructed. The matrix entry at position $((f, b), (f', b'))$ describes the probability, that mutation of a random individual having fitness f with b bits set incorrectly in the f th block results in an individual of fitness f' and b' bits set incorrectly in the f' th block. In the case of the RSF landscape already analyzed in the last chapter ($N = 8$ and $K = 8$) this results in a $(9 \times 8) \times (9 \times 8)$ matrix whose entries are calculated analytically. The mutation operator for the evolutionary algorithm on this RSF landscape is also kept: every bit is independently changed from 0 to 1 and vice versa with probability q .

For an easier representation we introduce the following 3 functions. $P_0(b)$ denotes the probability that in a K -bit block the number of incorrectly set bits before and after mutation is b . I assume that this can either happen if none of the K bits is mutated or if an arbitrary 1 mutates to 0 and an arbitrary 0 mutates to 1, i.e. by a two-bits swap. I do not account for the possibility that more bits swap as for any reasonable mutation rate q the probability for such an event is practically zero. Thus:

$$P_0(b) = (1 - q)^K + b(K - b)q^2(1 - q)^{K-2}$$

The function $P_+(b, n)$ gives the probability that the number b of incorrectly set bits in a block increases by n under mutation. This either occurs

if n out of $K - b$ bits are mutated to 0 and the rest remains unchanged or if a swap of two bits is involved, i.e. one of b bits mutates to 1, $n + 1$ out of $K - b$ mutate to 0, and $K - n - 2$ bits are unchanged:

$$P_+(b, n) = \binom{K-b}{n} q^n (1-q)^{K-n} + b \binom{K-b}{n+1} q^{n+2} (1-q)^{K-n-2}$$

Similarly, the function $P_-(b, n)$ denotes the probability that the number b of incorrectly set bits in a block decreases by n under mutation, which happens if either n out of b bits mutate to 1 and the remaining bits are unchanged, or if one out of $K - b$ mutates to 0 and $n + 1$ out of b bits mutate to 1, and $K - n - 2$ bits are not mutated. Thus:

$$P_-(b, n) = \binom{b}{n} q^n (1-q)^{K-n} + (K-b) \binom{b}{n+1} q^{n+2} (1-q)^{K-n-2}$$

The matrix entry $c_{(f,b),(f',b')}$ at position $((f, b), (f', b'))$ can now be determined as follows:

- If $f' < f$, the first $f' - 1$ blocks have to remain aligned, in the next block $K - b'$ arbitrary ones have to remain unchanged and the remaining b' bits have to be mutated to zeros:

$$c_{(f,b),(f',b')} = (1-q)^{(f'-1)K} (1-q)^{K-b'} q^{b'} \binom{K}{b'} \quad \text{for } f' < f$$

- If $f' = f$ and $b' = b$, $f' - 1$ blocks remain unchanged and in the f' th block the number of incorrectly set bits before and after mutation is equal:

$$c_{(f,b),(f',b')} = (1-q)^{(f'-1)K} P_0(b') \quad \text{for } f' = f \text{ and } b' = b$$

- If $f' = f$ and $b' = b + n$ where $n \in \{1, \dots, K - b\}$, $f' - 1$ blocks remain unchanged. In the f' th block the number of incorrectly set bits increases from b to b' by n . Thus:

$$c_{(f,b),(f',b')} = (1-q)^{(f'-1)K} P_+(b, n) \quad \text{for } f' = f \text{ and } b' = b + n$$

- Similarly, if $f' = f$ and $b' = b - n$ for $n \in \{1, \dots, b - 1\}$, $c_{(f,b),(f',b')}$ is given by:

$$c_{(f,b),(f',b')} = (1-q)^{(f'-1)K} P_-(b, n) \quad \text{for } f' = f \text{ and } b' = b - n$$

- Finally, if $f' > f$, $f - 1$ blocks are unchanged, in the first unaligned the b incorrect bits are mutated whereas the others are not changed, the next $f' - f - 1$ blocks have to be of the aligned form, and in the f' th block b' bits are set to 0. Therefore:

$$c_{(f,b),(f',b')} = (1-q)^{(f-1)K} q^b (1-q)^{K-b} \left(\frac{1}{2^K}\right)^{f'-f-1} \binom{K}{b'} \frac{1}{2^K}$$

for $f' > f$

As selection operator, the stochastic scheme described in Chapter 2.3 is still applied: If P is the population size, we choose P times from the set of all individuals, where individual $i \in \{1, \dots, P\}$ is independently chosen with probability proportional to its fitness f_i , i.e. with probability

$$p_i = \frac{f_i}{\sum_{j=1}^P f_j}.$$

Refining the process by using this two-dimensional matrix results in a very good prediction of the behavior of the actual evolutionary algorithm. The following table compares actual (first number) and predicted mean values of waiting times:

initial / final fitness	5	7	9
1	618 / 587	1304 / 1230	2595 / 2398
5	*	545 / 508	1873 / 1680
7	*	*	960 / 886

Mean waiting times are predicted with an error of less than 7%. This shows that the assumption that genotypes have their unaligned blocks set randomly is not only incorrect, but also responsible for systematically inaccurate estimates of evolutionary waiting times. The number of bits set correctly strongly influences the outcome of mutation. To give even better results, the idea of correctly set bits is now tested by focusing on the first two unaligned blocks.

3.2.2 Three-dimensional correlation

In the following individuals are described by the triple $(f, b1, b2)$. f is the fitness of the genotype, $b1$ is the number of bits set incorrectly in the first unaligned block, the f th block, and $b2$ is the number of bits set incorrectly in the block after the first unaligned, the $f + 1$ st block. $b1$ ranges from 1 to K , $b2$ can range from 0 to K , as block $f + 1$ can be aligned without influencing the fitness of the genotype.

This three-dimensional distribution is calculated analytically in a similar way to the calculations of two-dimensional correlation above. The resulting $(9 \times 8 \times 9) \times (9 \times 8 \times 9)$ correlation matrix has at entry $((f, b1, b2), (f', b1', b2'))$ the probability $c_{(f,b1,b2),(f',b1',b2')}$ that an individual of fitness f having $b1$ bits set incorrectly in the first unaligned block and $b2$ in the second changes to an individual of fitness f' with $b1'$ and $b2'$ bits set incorrectly in its f' th and $f' + 1$ st block under a mutational step.

- If $f' < f - 1$, $f' - 1$ blocks remain unchanged under mutation, in the f' th block $b1'$ out of K bits are mutated to 0 and $K - b1'$ are not mutated, and in the $f' + 1$ st block the same happens for $b2'$ bits. Thus:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K+(K-b2')+(K-b1')} \binom{K}{b1'} q^{b1'+b2'} \binom{K}{b2'}$$

for $f' < f - 1$

- If $f' = f - 1$ and $b2' = b1$, $f' - 1$ blocks are unchanged. In the f' th block $b1'$ out of K ones are mutated to zeros, and in the following block $b1$ bits are set incorrectly after mutation:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K+(K-b1')} \binom{K}{b1'} q^{b1'} P_0(b1)$$

for $f' = f - 1$ and $b2' = b1$

- If $f' = f - 1$ and $b2' = b1 + n$ for $n \in \{1, \dots, K - b1\}$, $f' - 1$ blocks are not mutated, in the f' block $b1'$ bits are changed to zero, and in the $f' + 1$ st block the number of incorrectly set bits $b1$ is increased by n under mutation. Thus:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K+(K-b1')} \binom{K}{b1'} q^{b1'} P_+(b1, n)$$

for $f' = f - 1$ and $b2' = b1 + n$

- Similarly, if $f' = f - 1$ and $b2' = b1 - n$ where $n \in \{1, \dots, b1\}$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K + (K - b1')} \binom{K}{b1'} q^{b1'} P_-(b1, n)$$

for $f' = f - 1$ and $b2' = b1 - n$

- For the case that $f' = f$ and $b1' = b1$, $f' - 1$ blocks are not mutated and in the f' th block the number of incorrect bits stays the same under mutation. For the outcome of mutation in block $f' + 1$ we have three possibilities:

- If $b2' = b2$, the number of zeros in the $f' + 1$ st block is unchanged:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_0(b1') P_0(b2')$$

for $f' = f$, $b1' = b1$ and $b2' = b2$

- If $b2' = b2 + n$, $n \in \{1, \dots, K - b2\}$, the number of zeros in block $f' + 1$ increases by n :

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_0(b1') P_+(b2, n)$$

for $f' = f$, $b1' = b1$ and $b2' = b2 + n$

- If $b2' = b2 - n$, $n \in \{1, \dots, b2\}$, the number of zeros in block $f' + 1$ decreases by n :

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_0(b1') P_-(b2, n)$$

for $f' = f$, $b1' = b1$ and $b2' = b2 - n$

- For the case that $f' = f$ and $b1' = b1 + n$ for $n \in \{1, \dots, b1\}$, $f' - 1$ blocks are unchanged, and in the first unaligned block the number of incorrect bits increases by n . According to the outcome of mutation in block $f' + 1$, again three cases are possible:

- If $f' = f$, $b1' = b1 + n$, and $b2' = b2$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_+(b1, n) P_0(b2)$$

- If $f' = f$, $b1' = b1 + n$, and $b2' = b2 + m$, $m \in \{1, \dots, K - b2\}$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_+(b1, n) P_+(b2, m)$$

- If $f' = f$, $b1' = b1 + n$, and $b2' = b2 - m$, $m \in \{1, \dots, b2\}$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_+(b1, n) P_-(b2, m)$$

- For the case that $f' = f$ and $b1' = b1 - n$, $n \in \{1, \dots, b1 - 1\}$, we distinguish between the following three cases:

- If $f' = f$, $b1' = b1 - n$, and $b2' = b2$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_-(b1, n) P_0(b2')$$

- If $f' = f$, $b1' = b1 - n$, and $b2' = b2 + m$, $m \in \{1, \dots, K - b2\}$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_-(b1, n) P_+(b2, m)$$

- If $f' = f$, $b1' = b1 - n$, and $b2' = b2 - m$, $m \in \{1, \dots, b2\}$:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K} P_-(b1, n) P_-(b2, m)$$

- If $f' = f + 1$, $f - 1$ blocks stay aligned, and in the f th block only the $b1$ incorrect bits have to be mutated. In the $f' + 1$ st block, about which nothing is known, $b2'$ out of K possible bits have to be set to zero and the remaining $K - b2'$ to one. According to the change due to mutation in the f' th block, which has $b2$ zeros before and $b1'$ after mutation, three cases are possible:

- If $b1' = b2$, the number of incorrectly set bits in block f' before and after mutation is equal. Thus:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K+(K-b1)} q^{b1} P_0(b2) \binom{K}{b2'} \frac{1}{2^K}$$

for $f' = f + 1$, $b1' = b2$

- If $b1' = b2 + n$, $n \in \{1, \dots, K - b2\}$ the number of zeros in block f' increases by n :

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f'-1)K+(K-b1)} q^{b1} P_+(b2, n) \binom{K}{b2'} \frac{1}{2^K}$$

for $f' = f + 1$, $b1' = b2 + n$

– Finally, if $b1' = b2 - n$, $n \in \{1, \dots, b2\}$,

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f-1)K+(K-b1)} q^{b1} P_{-}(b2, n) \binom{K}{b2'} \frac{1}{2^K}$$

for $f' = f + 1$, $b1' = b2 - n$.

- At last, if $f' > f + 1$, $f - 1$ blocks are not mutated and in the f th and $f + 1$ st block only the $b1$, respectively $b2$ zeros have to be mutated to ones. After the $f + 1$ st block, $f' - f - 2$ blocks have to be in the aligned form, in the f' th and $f' + 1$ st block exactly $b1'$ respectively $b2'$ bits out of K possible are zeros. Thus:

$$c_{(f,b1,b2),(f',b1',b2')} = (1 - q)^{(f-1)K+(K-b1)+(K-b2)} q^{b1+b2}$$

$$\left(\frac{1}{2^K}\right)^{f'-f-2} \binom{K}{b1'} \frac{1}{2^K} \binom{K}{b2'} \frac{1}{2^K}$$

for $f' > f + 1$.

Based on the three-dimensional correlation information, a simplified stochastic process, approximating the actual evolutionary algorithm, can be constructed. The waiting times predicted by this process are practically exact, see Figure 3.5.

The following table compares actual (first number) and predicted mean values of waiting times for the distributions shown in Figure 3.5:

initial / final fitness	5	7	9
1	618 / 619	1304 / 1260	2595 / 2531
5	*	545 / 528	1873 / 1793
7	*	*	960 / 946

The difference between predicted and actual mean values of waiting times is of less than 3%. This shows that incorporating more and more correlation information about the individuals and their mutants can capture all necessary features of the evolutionary process.

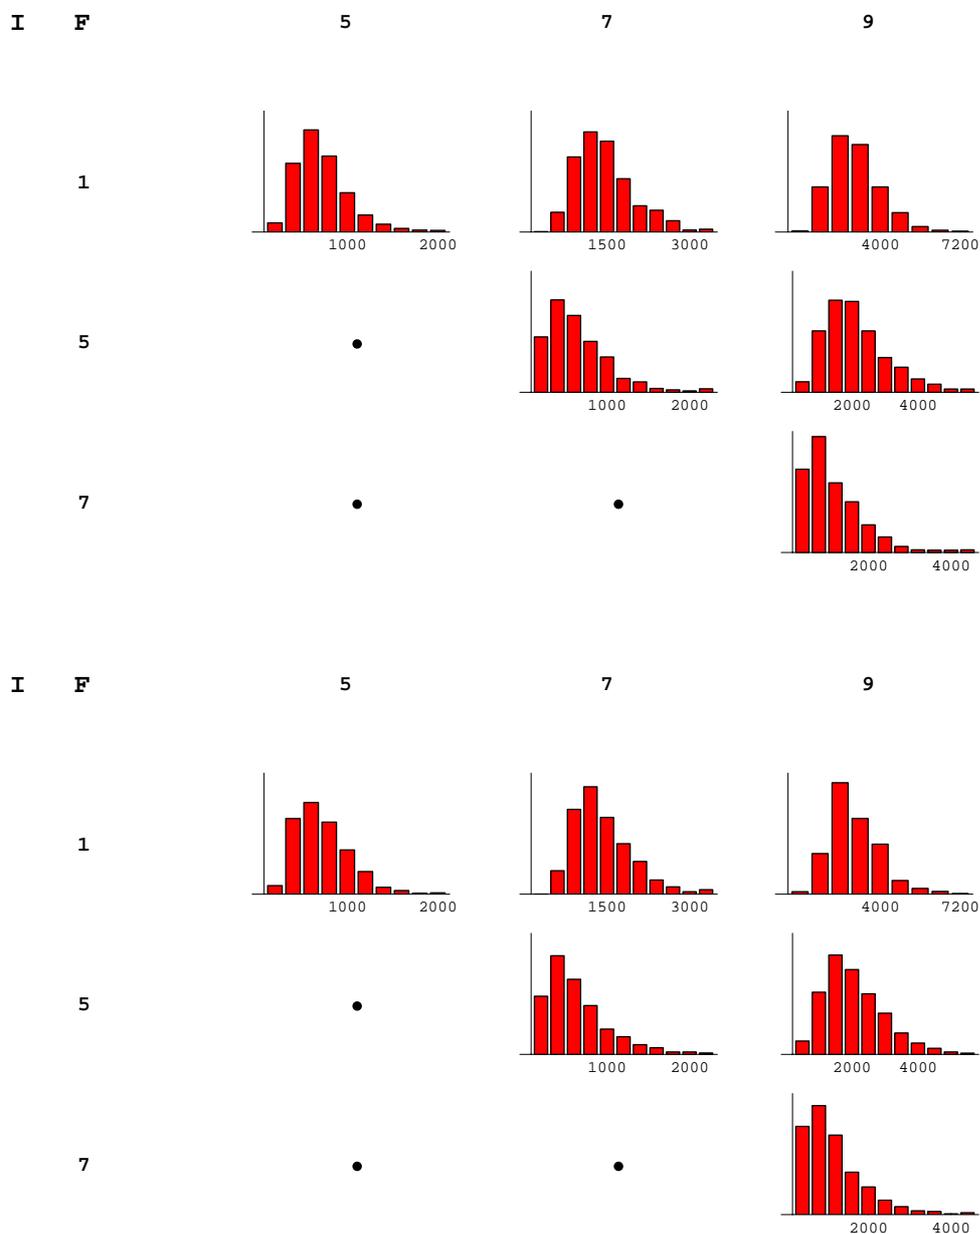


Figure 3.5: *Actual (top) and predicted (bottom) waiting time distributions derived from 1000 runs each, based on three-dimensional correlation. Including more information about the individuals evolving on an RSF landscape allows for still better predictions of evolutionary waiting times than it was possible with two-dimensional correlation.*

In the last two chapters I have demonstrated that correlation information provides the information necessary to predict evolutionary waiting times on fitness landscapes with great accuracy. As the analyzed evolutionary algorithms involve mutation and selection, I tried to find a way of predicting the outcome of a mutational step. The selection scheme is based on the fitness values of individuals only and therefore is relatively easy to describe. Once the current fitness distribution in the population is known, the fitness distribution of the selected new generation can be determined. By contrast, the working of the mutation operator is much more difficult to understand, and depends on the structure and topology of the fitness landscape. As a mutant of an individual is always a neighbor of this individual in genotype space, the outcome of mutation depends on the neighborhood structure of individuals.

I described a fitness landscape by the correlation under single mutational steps between classes of individuals which have a similar neighborhood structure. This new type of correlation information proved to be successful in predicting the performance of an evolutionary algorithm at the level of fitness. For fitness landscapes involving low degrees of neutrality, a one-dimensional correlation description, based on characterizing individuals by their fitness only, succeeded in predicting the distribution of waiting times between two fitness levels. This method was applied successfully to TSP and NKp landscapes with low neutrality. In fitness landscapes with high neutrality, extended neutral networks of individuals with the same or very similar fitness values were encountered. Those networks are too inhomogeneous to be described by one-dimensional correlation information, and I thus extended this description to two- and three-dimensional correlation. Based on this type of information, almost accurate predictions of waiting times on NKp landscapes with high neutrality as well as on RSF landscapes were derived. I conclude that correlation information has the potential for capturing all the features of fitness landscapes that determine the fitness-level dynamics of evolutionary processes.

Chapter 4

Towards analytic descriptions

In the last chapters we have studied a stochastic process for predicting evolutionary waiting times on several specific fitness landscapes. This process describes the state of an evolving population in terms of the distribution over classes of individuals. These classes of individuals are determined by stating crucial properties of genotypes, which influence the fitness of their neighbors, i.e. their mutants. The neighborhood structure is then assumed to be the same for all individuals belonging to the same class. In particular, the correlation between all pairs of classes is estimated, which is defined as the probability, that a random mutant of an individual belonging to the first class is in the second class. With this correlation information, the outcome of mutating a genotype concerning the fitness is approximated. As selection depends on fitness values only, it can relatively easily be described. By reducing genotypes to the class they belong, and by applying the simplified mutation and the actual selection operator, the evolutionary algorithm is approximated by a simpler stochastic process based on correlation statistics only.

Given that S is the current distribution of individuals over the classes, the dynamics are in principle determined by the transition probabilities $P(S'|S)$ where S' is the distribution over classes of the next generation. Unfortunately, for any reasonable population size and for a fine-grained classification of genotypes, the number of possible distributions of a population over classes is enormous. This makes it almost impossible to construct the transition matrix of a Markov-chain for an analytic prediction of mean and variance of waiting times.

The outcome of many realizations of the stochastic process itself was

therefore used to predict the distribution of waiting times. Computation time is much shorter for such a simplified process than for the actual evolutionary algorithms, which require large numbers of fitness function evaluations. An analytic description of evolutionary dynamics would not only allow still quicker predictions, but also more insight into how mutation and selection operators influence the scale of evolutionary waiting times. For the case of fitness landscapes with low neutrality and one-dimensional correlation statistics showing normally distributed mutant fitness, an analytic description of the average evolutionary dynamics on the level of fitness in large populations is feasible.

As shown in Chapter 2, fitness landscapes of low neutrality can be described very well by one-dimensional correlation statistics. Moreover we have observed, that in the case of the analyzed TSP landscapes and the NKp landscape of low neutrality, these one-dimensional correlation statistics, and thus the distribution of mutant fitness, can be well approximated by normal distributions. Accordingly, when starting with a population in generation 0 with all individuals having the same genotype, one step of mutation alone leads to a normal distribution of fitness in infinite populations. A stochastic selection scheme based on an exponential selection function is expected to shift the mean of this normal distribution. Fitness in generation 1 will therefore still be normally distributed. Based on these ideas I now try to develop a deterministic description for the evolution of the normally distributed fitness in large populations. As a normal distribution is determined by the two parameters mean and variance, formulas for their change under mutation and selection are aspired.

This analysis is unfortunately not as straightforward as it might seem at first sight. Once the population has passed generation 0 and is normally distributed with respect to fitness, the mutation process amounts to a combination of two normal distributions whose parameters are interdependent. First, the fitness in the population is normally distributed. Secondly, every individual of the population has mutants whose fitness values are again normally distributed, but the mean and variance of this distribution depend on the fitness of the parent individual. In the case of the TSP and NKp landscapes with low neutrality, one-dimensional correlation statistics can be approximated by normal distributions, whose mean value depend linearly on the parent fitness. When further assuming that the variance of mutant fitness is the same for all individuals in a certain generation, this combination of two normal distributions results again in a normal distribution after mutation.

Under these conditions, formulas for the evolution of normally distributed fitness under mutation and selection can be derived.

4.1 TSP fitness landscapes

The analytic approximation of an evolutionary algorithm on a fitness landscape presented in the following is based on several assumptions. First, it is required that a one-dimensional correlation description of the landscape captures all the essential information for predicting evolutionary change on the level of fitness. Secondly, the distribution of mutant fitness for fixed parent fitness, as provided by this correlation information, needs to be normal in a good approximation. Thirdly, the mean and variance of the normally distributed mutant fitness must depend linearly on the parent fitness, and the function for variance has to be nearly constant. Moreover, the selection scheme must be based on an exponential selection function, and the population has to be large enough so that its fitness distribution can be approximated by normal distributions.

In the following, these prerequisites for the analytic description will be clarified, and I will show that the assumptions that have to be made are justified for the TSP landscape already studied in the last chapters. Thereafter, these assumptions are used to develop recursion equations for the change of the normal distribution of fitness under mutation and selection, and a consequential approximation of the best fitness in only finite populations is presented. Finally, the analytic predictions are compared with the dynamics of an actual evolutionary algorithm. In Section 4.2 I show that these deterministic approximations of evolutionary dynamics are also successful for the NKp landscape of low neutrality already studied in the last chapters.

4.1.1 Prerequisites

The coordinates of the cities in our Travelling Salesman Problems have been chosen randomly. Consequently, the distances between pairs of cities are random numbers. The length of a tour visiting all cities is the sum of N of these random distances, if N is the number of cities. According to the central limit theorem, tour lengths are approximately normally distributed. If genotypes, i.e. tours of the TSP, are mutated, n distances between cities are replaced by other random values. Under point mutation, $n = 4$, for

remove-and-reinsert mutation $n = 3$, and for reverse mutation, $n = 2$ (see Section 2.1.2). Even for reverse mutation, the tour lengths of the mutants of an individual can therefore be approximated by a normal distribution. The fitness function of the TSP studied in Section 2.1.4 is a linear function of the length. The corresponding distributions for fitness instead of length can thus also be approximated by normal distributions, with different mean and variance.

The one-dimensional correlation matrix not only confirms that the neighbors of individuals belonging to a certain fitness class are normally distributed with respect to fitness, see Figure 2.24. We also observed that the mean value and the variance of this distribution depend linearly on the parent fitness value in a very good approximation, see Figure 2.25. Let $f_m(x)$ be the function for the mean value of mutant fitness for parent fitness x , and let $f_v(x)$ be the corresponding function for the variance of fitness. As both functions are linear functions of parent fitness in a close fit, they are well approximated by:

$$\begin{aligned} f_m(x) &= k_m x + d_m \\ f_v(x) &= k_v x + d_v \end{aligned}$$

Together with the assumption of normal distribution, the one-dimensional correlation matrix is approximated by specifying these two functions. This representation of correlation statistics was successfully tested as substitute for the correlation matrix, see Figure 2.26. Consequently, the assumption of normal distribution of mutant fitness, and linear dependence of the parameters on parent fitness is justified.

4.1.2 Approximation of fitness distribution dynamics

Let $p_g(x)$ be the probability density of fitness of the population on the TSP landscape in generation g , and let $c(x, y)$ be the conditional density of mutant fitness y given that the parent fitness is x . The function $c(x, y)$ is the continuous equivalent to the one-dimensional correlation matrix. According to the total probability theorem, the joint density function of mutant fitness and parent fitness, $F(x, y)$, is given by $F(x, y) = c(x, y)p_g(x)$. In the following it will be shown how the density of fitness after mutation, $p'_g(y)$, and the density after selection, $p''_g(y)$, can be approximated. By combining both processes, recursion relations are derived for the change of mean and variance of normally distributed fitness over generations.

Mutation

At first, it will be demonstrated how to calculate the probability density of fitness after all individuals have been mutated, $p'_g(y)$. This density is determined by

$$p'_g(y) = \int_{-\infty}^{\infty} F(x, y) dx = \int_{-\infty}^{\infty} c(x, y) p_g(x) dx.$$

Let $N_{m,v}(x)$ be the density function of a normal distribution with mean m and variance v . Using the functions $f_m(x)$ and $f_v(x)$ for mean and variance of mutant fitness depending on parent fitness x , as stated above, the conditional density $c(x, y)$ can be written as

$$c(x, y) = N_{f_m(x), f_v(x)}(y).$$

As $p_g(x)$ is assumed to be the density function of a normal distribution with mean m_g and variance v_g , the joint density $F(x, y)$ can now be written as

$$F(x, y) = N_{f_m(x), f_v(x)}(y) N_{m_g, v_g}(x).$$

Without any additional simplification, $p'_g(y)$ is given by

$$\begin{aligned} p'_g(y) &= \int_{-\infty}^{\infty} F(x, y) dx \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi f_v(x)}} e^{-\frac{(y-f_m(x))^2}{2f_v(x)}} \frac{1}{\sqrt{2\pi v_g}} e^{-\frac{(x-m_g)^2}{2v_g}} dx \\ &= \frac{1}{2\pi\sqrt{v_g}} \int_{-\infty}^{\infty} \frac{1}{\sqrt{k_v x + d_v}} e^{-\frac{(y-k_m x - d_m)^2}{2(k_v x + d_v)} - \frac{(x-m_g)^2}{2v_g}} dx. \end{aligned}$$

For this integral no analytic solution was found.

However, when assuming that the function $f_v(x)$ is not depending on x but constant in each generation, the function $p'_g(y)$ can be calculated analytically. If k_v is close to zero and therefore $f_v(x)$ is almost constant, which is true for this TSP landscape, we can approximate the function $f_v(x)$ in every

generation by its value at the current mean fitness in the population $f_v(m_g)$ and get

$$\begin{aligned}
F(x, y) &= N_{f_m(x), f_v(m_g)}(y) N_{m_g, v_g}(x) \\
&= \frac{1}{\sqrt{2\pi f_v(m_g)}} e^{-\frac{(y-f_m(x))^2}{2f_v(m_g)}} \frac{1}{\sqrt{2\pi v_g}} e^{-\frac{(x-m_g)^2}{2v_g}} \\
&= \frac{1}{2\pi \sqrt{f_v(m_g) v_g}} e^{-\frac{1}{2f_v(m_g)v_g} [v_g(y-(k_m x+d_m))^2 + f_v(m_g)(x-m_g)^2]} \\
&= \frac{1}{2\pi \sqrt{f_v(m_g) v_g}} e^{-A} e^{-B}
\end{aligned}$$

$$\begin{aligned}
\text{with } A &= \frac{1}{2f_v(m_g)v_g} [x^2 (f_v(m_g) + k_m^2 v_g) + 2x (d_m k_m v_g - f_v(m_g)m_g) \\
&\quad - 2xy(k_m v_g) - 2y(d_m v_g) + y^2(v_g)] \\
\text{and } B &= \frac{1}{2f_v(m_g)v_g} (v_g d_m^2 + f_v(m_g)m_g^2).
\end{aligned}$$

If we introduce the 5 variables

$$\begin{aligned}
n_1 &= m_g \\
n_2 &= k_m n_1 + d_m \\
s_1 &= \sqrt{v_g} \\
s_2 &= \sqrt{f_v(m_g) + v_g k_m^2} \\
r &= \frac{k_m s_1}{s_2}
\end{aligned}$$

we obtain

$$\begin{aligned}
A &= \frac{1}{2f_v(m_g)v_g} (x^2 s_2^2 + 2x(n_2 r s_1 s_2 - n_1 s_2^2) - \\
&\quad - 2xy(r s_1 s_2) - 2y(n_2 s_1^2 - n_1 r s_1 s_2) + y^2(s_1^2)) \\
&= \frac{1}{2(1-r^2)} \left[\frac{(x-n_1)^2}{s_1^2} - \frac{2r(x-n_1)(y-n_2)}{s_1 s_2} + \frac{(y-n_2)^2}{s_2^2} \right] - C
\end{aligned}$$

$$\text{with } C = \frac{1}{2(1-r^2)} \left[\frac{n_1^2}{s_1^2} - \frac{2r n_1 n_2}{s_1 s_2} + \frac{n_2^2}{s_2^2} \right]$$

$$\begin{aligned}
&= \frac{s_2^2}{2(s_2^2 - k_m^2 s_1^2)} \left[\frac{n_1^2 s_2^2 - 2k_m n_1 n_2 s_1^2 + n_2^2 s_1^2}{s_1^2 s_2^2} \right] \\
&= \frac{f_v(m_g) m_g^2 + d_m^2 v_g}{2f_v(m_g) v_g}.
\end{aligned}$$

Since $C = B$, $F(x, y)$ can be written in the form

$$F(x, y) = \frac{1}{2\pi \sqrt{f_v(m_g) v_g}} e^{-\frac{1}{2(1-r^2)} \left[\frac{(x-n_1)^2}{s_1^2} - \frac{2r(x-n_1)(y-n_2)}{s_1 s_2} + \frac{(y-n_2)^2}{s_2^2} \right]}.$$

The marginal density $p_{g_{mut}}(x)$ is now given by:

$$\begin{aligned}
p'_g(y) &= \int_{-\infty}^{\infty} F(x, y) dx \\
&= \frac{1}{2\pi \sqrt{f_v(m_g) v_g}} \int_{-\infty}^{\infty} e^{-\frac{1}{2(1-r^2)} \left[\frac{(x-n_1)^2}{s_1^2} - \frac{2r(x-n_1)(y-n_2)}{s_1 s_2} + \frac{(y-n_2)^2}{s_2^2} \right]} dx.
\end{aligned}$$

Since

$$\begin{aligned}
&\frac{(x-n_1)^2}{s_1^2} - \frac{2r(x-n_1)(y-n_2)}{s_1 s_2} + \frac{(y-n_2)^2}{s_2^2} = \\
&\left[\frac{x-n_1}{s_1} - \frac{r(y-n_2)}{s_2} \right]^2 + \frac{(1-r^2)(y-n_2)^2}{s_2^2}
\end{aligned}$$

we obtain

$$p'_g(y) = \frac{1}{2\pi \sqrt{f_v(m_g) v_g}} e^{-\frac{1}{2} \left(\frac{y-n_2}{s_2} \right)^2} \int_{-\infty}^{\infty} e^{-\frac{1}{2(1-r^2)} \left[\frac{x-n_1}{s_1} - \frac{r(y-n_2)}{s_2} \right]^2} dx.$$

With

$$\frac{1}{\sqrt{1-r^2}} \left(\frac{x-n_1}{s_1} - \frac{r(y-n_2)}{s_2} \right) = z$$

it is easy to see that

$$\begin{aligned}
 p'_g(y) &= \frac{1}{2\pi\sqrt{f_v(m_g)v_g}} e^{-\frac{1}{2}\left(\frac{y-n_2}{s_2}\right)^2} s_1\sqrt{1-r^2} \int_{-\infty}^{\infty} e^{\frac{z^2}{2}} dz \\
 &= \sqrt{\frac{1-r^2}{2\pi f_v(m_g)}} e^{-\frac{1}{2}\left(\frac{y-n_2}{s_2}\right)^2} \\
 &= \frac{1}{s_2\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{y-n_2}{s_2}\right)^2}.
 \end{aligned}$$

Thus $p'_g(y)$ is normal with mean $n_2 = k_m n_1 + d_m = f_m(m_g)$ and variance $s_2^2 = f_v(m_g) + v_g k_m^2$.

In summary, the mutation operator changes the mean and variance of the current fitness in the following way:

$$\begin{aligned}
 m_g &\longrightarrow f_m(m_g) \\
 v_g &\longrightarrow f_v(m_g) + v_g k_m^2.
 \end{aligned}$$

Selection

In Section 2.1.4 we have chosen a stochastic selection scheme which selects individual j with probability proportional to the exponential function $e^{Sf(j)}$, where $f(j)$ is the fitness of j and S is a positive selection coefficient. If $p(x) = N_{m,v}(x)$ is the normal density function for the current fitness with mean m and variance v , the probability density after selection $p''_g(x)$ is proportional to the product of the two functions. Thus:

$$\begin{aligned}
 p''_g(x) &= \frac{1}{c} p(x) e^{Sx} \quad \text{with } c = \int_{-\infty}^{\infty} p(x) e^{Sx} dx \\
 &= \frac{1}{c} \frac{1}{\sqrt{2\pi v}} e^{-\frac{(x-m)^2}{2v}} e^{Sx} \\
 &= \frac{1}{c} \frac{1}{\sqrt{2\pi v}} e^{-\frac{1}{2v}((x-m)^2 - 2vSx)} = \frac{1}{c} \frac{1}{\sqrt{2\pi v}} e^{-\frac{1}{2v}(x^2 - 2x(m+vS) + m^2)} \\
 &= \frac{1}{c} \frac{1}{\sqrt{2\pi v}} e^{-\frac{1}{2v}((x-(m+vS))^2 - 2mvS - v^2S^2)}
 \end{aligned}$$

$$= \frac{1}{c} \frac{1}{\sqrt{2\pi v}} e^{-\frac{(x-(m+vS))^2}{2v}} e^{-\frac{2mS+vS^2}{2}}.$$

Since

$$\begin{aligned} c &= \int_{-\infty}^{\infty} p(x) e^{Sx} dx = \frac{1}{\sqrt{2\pi v}} \int_{-\infty}^{\infty} e^{-\frac{(x-m)^2+2vSx}{2v}} dx \\ &= \frac{1}{\sqrt{2\pi v}} \int_{-\infty}^{\infty} e^{-\frac{(x-(m+vS))^2}{2v}} e^{-\frac{2mvS+v^2S^2}{2v}} dx = e^{-\frac{2mS+vS^2}{2}}, \end{aligned}$$

we obtain

$$p_g''(x) = \frac{1}{\sqrt{2\pi v}} e^{-\frac{(x-(m+vS))^2}{2v}} = N_{m+vS,v}(x),$$

Starting with normally distributed fitness in a very large population, the exponential selection function leads again to a normal distribution of fitness after selection with the mean value shifted towards better fitness. The variance stays the same. If m'_g is the current mean fitness and v'_g the variance of fitness, the selection scheme changes those parameters in the following way:

$$m'_g \longrightarrow m'_g + v'_g S$$

$$v'_g \longrightarrow v'_g$$

Recursion relations

Starting with a normal distribution of fitness in generation g with mean m_g and variance v_g , the processes of mutation and selection ensure that the fitness distribution in generation $g + 1$ is still normal. For the change of mean and variance under mutation and selection of fitness, we now have the following expressions:

$$m_g \longrightarrow f_m(m_g)$$

Mutation:

$$v_g \longrightarrow f_v(m_g) + v_g k_m^2$$

$$f_m(m_g) \longrightarrow f_m(m_g) + S(f_v(m_g) + v_g k_m^2)$$

Selection:

$$f_v(m_g) + v_g k_m^2 \longrightarrow f_v(m_g) + v_g k_m^2$$

The development of mean and variance of fitness over generations is now given by the recursion relations

$$m_{g+1} = f_m(m_g) + S [f_v(m_g) + v_g k_m^2]$$

$$v_{g+1} = f_v(m_g) + v_g k_m^2.$$

4.1.3 Approximation of maximal fitness

For a comparison of actual waiting times of the evolutionary algorithm on the TSP landscape and predictions based on these equations, we have to approximate the maximal fitness f_{max} in a finite population of individuals whose fitness values are approximately normally distributed with mean m and variance v . Then we can predict the number of generations necessary, until one individual's fitness exceeds the final fitness threshold. In a population of N individuals, $N - 1$ individuals have fitness lower or equal to f_{max} . We therefore determine f_{max} from the equation

$$\frac{1}{\sqrt{2\pi v}} \int_{-\infty}^{f_{max}} e^{-\frac{(x-m)^2}{2v}} dx = \text{erf}_{m,v}(f_{max}) = \frac{N-1}{N}$$

and obtain

$$f_{max} = \text{erf}_{m,v}^{-1}\left(1 - \frac{1}{N}\right).$$

4.1.4 Testing the predictions

Predictions are expected to be most accurate for a very large population size and a rather small selection coefficient. If the population size is too small, the actual fitness distribution cannot be approximated by a normal distribution. Similarly, a big selection coefficient in a finite population leads to an

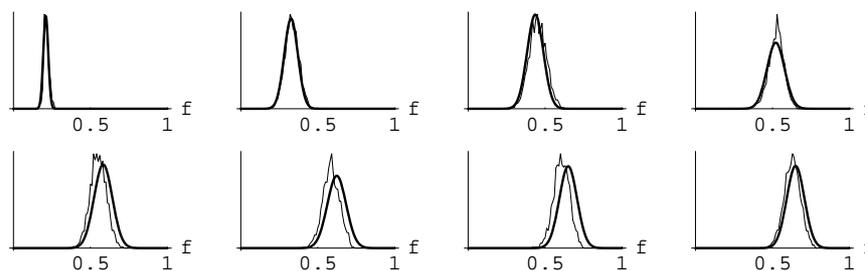


Figure 4.1: *Actual (thin line) and predicted (thick line) fitness distribution in a population of 5000 individuals in generation 0, 10, 20, 30 in the upper row and generation 40, 50, 100 and 200 in the lower row. The actual fitness distributions are derived from a single run of the evolutionary algorithm.*

actual distribution of fitness after selection which is qualitatively different from a normal distribution, because in a finite population, the probability for an individual to have fitness above a certain threshold is 0. In an approximation based on normal distributions, however, all fitness values have a positive probability to occur. The outcome of selection in a finite population, therefore, is not simply a shift of the mean of the fitness distribution, but depends on the actually occurring fitness values.

In order to test the analytic predictions, the parameters of the evolutionary algorithm are set as follows: The population size is chosen at $P = 5000$ and the selection coefficient at $S = 5$. The predictions are based on the two recursion relations for mean and variance of fitness and thus only on the functions $f_m(x)$ and $f_v(x)$ for mean and variance of mutant fitness for parent fitness x . Actually, the best predictions are obtained when using not only one linear function for $f_m(x)$ and $f_v(x)$ but instead for each a set of 10 linear functions for different fitness ranges. These functions are derived from the correlation matrix by linear interpolation between the points of mean and variance of mutant fitness for all different classes. The 10 linear functions per parameter are nearly equal to each other, but nevertheless slight differences are important for the accuracy of predictions.

Figure 4.1 shows a comparison of the development of actual and predicted fitness distribution over generations. Actual and predicted mean and best fitness in the population evolving over 200 generations are depicted in Figure 4.2. The statistics of actual behavior of the evolutionary algorithm are de-

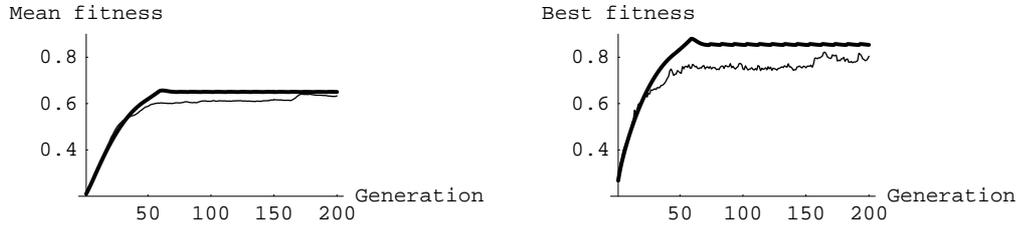


Figure 4.2: *Actual (thin line) and predicted (thick line) mean and best fitness in a population of individuals evolving on a TSP fitness landscape. The predictions are based on an analytic description of the evolutionary dynamics.*

rived from a single run. For high population size, different runs of the actual process do not show important differences.

The accuracy of the predictions is remarkable. This shows not only that the TSP fitness landscape can be described by one-dimensional correlation statistics. It also demonstrates that the correlation matrix itself can be approximated in a very simple way. We reduced the high-dimensional fitness landscape to two sets of linear functions, which allow for an analytic prediction of the evolutionary process. In the next section this method of analytic description is applied to the NK p fitness landscapes of low neutrality.

4.2 NK p fitness landscapes

In an NK p fitness landscape, the fitness of a genotype is on average a sum of $N(1 - p)$ random numbers. For not very high values of p , the fitness distribution therefore tends to be normal. Under mutation, on average $(K + 1)(1 - p^2)$ bits of the genotype vector change their contribution and mutant fitness is likewise approximately normally distributed, if p is not very high. In the NK p landscape of low neutrality ($N = 40$, $K = 3$, and $p = 0.3$) that was investigated in Section 2.2, $N(1 - p) = 28$ and $(K + 1)(1 - p^2) = 3.64$. Consequently, the fitness distribution, as well as mutant fitness for most fitness classes, are approximated quite well by normal distributions (see Section 2.2 and Figure 2.30).

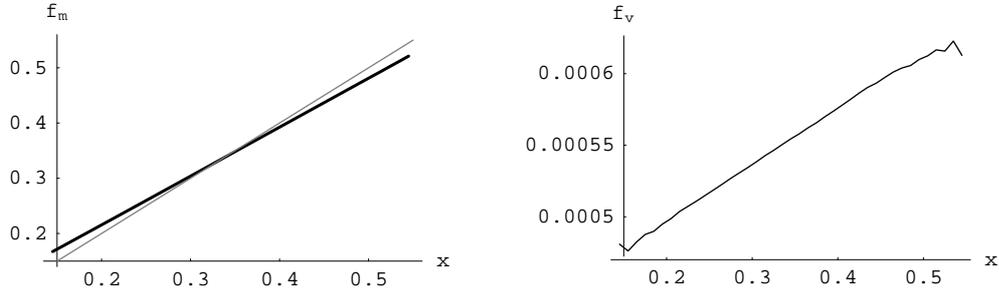


Figure 4.3: *The functions $f_m(x)$ and $f_v(x)$ for mean and variance of neighbor fitness depending on parent fitness x in the NKp landscape of low neutrality are linear in a good approximation. In the first picture the identity function is depicted in gray.*

The analytic predictions for the evolution of fitness distribution that we have developed for TSP landscapes can therefore, in principle, also be applied to the NKp landscape of low neutrality investigated in Section 2.2, as for this landscape a one-dimensional correlation description has been successful. The population size is set to $P = 5000$ and we apply the stochastic selection scheme already used for the evolutionary algorithm of the TSP, see last section. The selection coefficient is chosen at $S = 10$.

4.2.1 Testing the predictions

For each of the functions of mean and variance of mutant fitness for parent fitness x , $f_m(x)$ and $f_v(x)$, not only one linear function is taken. Like in the case of the TSP of the last section predictive accuracy is improved, when for both $f_m(x)$ and $f_v(x)$, 10 linear functions for different fitness ranges are taken. They are derived from the correlation matrix by linear interpolation between mean values and variances of mutant fitness for different classes, see Figure 4.3. With the assumption of normally distributed mutant fitness, the correlation matrix now has a continuous representation by means of the functions $f_m(x)$ and $f_v(x)$.

A comparison of actual and predicted evolution of fitness distribution is presented in Figure 4.4. The actual distribution is derived from a single run. For a population size of $P = 5000$, different runs of the evolutionary

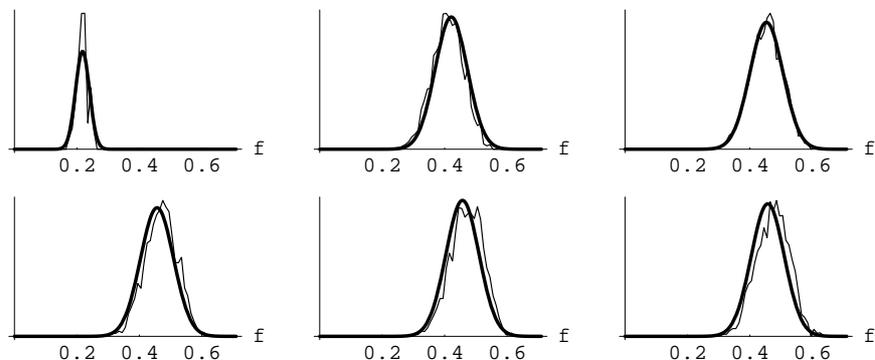


Figure 4.4: *Actual (thin line) and predicted (thick line) fitness distribution in a population of 5000 individuals in generation 0, 20, 40 in the upper row and generation 60, 80 and 100 in the lower row.*



Figure 4.5: *Actual (thin line) and predicted (thick line) mean and best fitness for evolution in an NK_p landscape. The actual statistics are derived from a single run, whereas the predictions are based on an analytic approximation of the average dynamics on the level of fitness.*

algorithm do not show important differences. A comparison of best and mean fitness per generation in one run of the actual algorithm and the analytic predictions is presented in Figure 4.5.

The predictions, which are only based on the two sets of linear functions for the mean and variance of mutant fitness, are very close to the actual dynamics. An analytic approximation of evolutionary dynamics based on normal distributions was successful for the TSP landscape and the NKp landscape of low neutrality. In the next subsection, however, some ideas for improving predictive accuracy and for more general approximations are presented.

4.2.2 Possible improvements and extensions

When we calculated the probability density function of fitness after all individuals have been mutated, p'_g , the linear function $f_v(x)$ for the variance of neighbor fitness was approximated by its value at the current mean fitness in the population $f_v(m_g)$:

$$\begin{aligned} p'_g(y) &= \int_{-\infty}^{\infty} F(x, y) dx \\ &= \int_{-\infty}^{\infty} N_{f_m(x), f_v(x)}(y) N_{m_g, v_g}(x) dx \\ &\approx \int_{-\infty}^{\infty} N_{f_m(x), f_v(m_g)}(y) N_{m_g, v_g}(x) dx. \end{aligned}$$

Without this simplification, p'_g could not be calculated analytically. Numerical evaluations show that this density function is in general not normal, but can be approximated very well by the density function of a normal distribution, if the slope of the function $f_v(x)$ is not too high. In the case of the analyzed TSP and NKp landscape of low neutrality, the slope of $f_v(x)$ is very small. For that reason, the predictions derived when making this simplification are very close to the actual dynamics.

An substitution of $F(x, y)$ by its Taylor-polynomial of first and second order yields exactly the same mean and variance of the so approximated

density function p'_g , as is achieved when using the above simplification. Consequently, the mean and variance of the actual distribution of fitness after mutation is certainly very close to the moments we could calculate analytically when allowing for this simplification. However, predictive accuracy of the analytic description could possibly be improved when accounting for the more accurate distribution of fitness after mutation.

In principle, it should be feasible to base predictions on the moments of distributions different from normal distributions. Such a description would no longer be restricted to fitness landscapes for which mutant fitness is normally distributed. Predictions for the fitness distribution evolving over generations in landscapes with large degree of neutrality, like NKp landscapes for very high p or RSF landscapes, would then be possible. To develop this approach systematically is an interesting option for future research.

Nevertheless, the analytic description of the evolutionary algorithm on a TSP landscape and an NKp landscape with low neutrality, based on normally distributed fitness, provided surprisingly good estimates of the actual process. One-dimensional correlation statistics have proved to be capable of describing TSP and NKp landscapes of low neutrality. Moreover, the correlation statistics for these landscapes have themselves shown such a simple structure, that they can be reduced to a set of normal distributions. This simplicity of the correlation structure of these high-dimensional fitness landscapes was unexpected and allows for considerably reducing the complexity of such fitness landscapes.

Chapter 5

Summary

The structure of fitness landscapes

Although the metaphor of a fitness landscape on which a population is evolving towards adaptive peaks has served as an important basis for understanding evolutionary processes in different areas of science, the question which landscape statistics are critical for predicting evolutionary change on these landscapes so far has not been resolved. In this study, we have suggested a potential answer to this question, based on correlation approximations of evolutionary algorithms.

Evolutionary algorithms, most commonly used for solving problems of optimization, can be regarded as evolutionary systems in which the dynamics at the level of fitness are resulting from the dynamics at a genotype level. Whereas the specification of evolutionary algorithms is generally straightforward, their actual dynamics are often difficult to understand. The picture of a fitness landscape in which individuals of an evolutionary algorithm are engaged in hill-climbing can, in principle, help to improve this understanding. Actual fitness landscapes of specific evolutionary algorithms, however, are mostly high-dimensional, and identifying features which are critical for the evolutionary outcome has turned out to be a challenging problem.

Until now, mainly two different correlation functions have been used to describe fitness landscapes; the *direct correlation function* and the *auto-correlation function*, see Section 1.4. Based on these functions rough approximations of the behavior of the evolutionary process are possible (Stadler 1995, Manderick 1997). The direct correlation function is based on considering the correlation of fitness values between pairs of genotypes at varying mutational

distances. The auto-correlation function is based on time series of fitness values that result from random walks on a fitness landscape. The dynamics of such random walks are only determined by mutation, selective pressures are not considered. Auto-correlations are then averaged over all possible initial conditions for the random walk. It can be shown that these two types of correlation functions carry equivalent information (Stadler 1995).

Fitness landscapes without neutrality

In this study I presented a new correlation measure which, relative to the previous approaches, at the same time reduces and enhances the statistical information provided: the focus of this new measure is on single mutational steps (reduced information) but the initial fitness of considered individuals is maintained in the measure (enhanced information). Unless the fitness function is degenerate, it is very probable that an individual's fitness is to some degree correlated with the fitness values of its mutants or neighbors. The underlying idea of this correlation approximation is to describe a fitness landscape by means of its neighborhood structure.

In Chapter 2, I therefore investigated one-dimensional correlation measures, which are based on the assumption that the neighborhood structure in terms of fitness values is similar for individuals of the same fitness range. The description of a fitness landscape then amounts to stating the distributions of mutant fitness for different fitness classes.

One of the essential features of the dynamics of an evolutionary algorithm concerns the probability for a population of a certain initial fitness to reach a fitness threshold within g generations. This information is summarized in the distribution of waiting times. Comparing the actual waiting times of evolutionary algorithms for several specific problems to predictions based on various candidate statistics allowed for assessing the relative merits of these simplified landscape descriptions.

In Section 2.2 the focus was on the Travelling Salesman Problem and on its solution by an evolutionary algorithm that utilizes a point-mutation operator. Within this setting, the complexity of landscape statistics has been increased in a sequence of several steps. I have shown that a landscape description neglecting the correlation between neighboring fitness values (percolation) is not successful in predicting the evolutionary dynamics; subsequently, I investigated one-dimensional correlation statistics:

1. *Percolation*. All individuals on the landscape are considered to have

fitness values above or below a given threshold with probability p and $1 - p$, respectively (Gavrilets and Gravner 1997). The results of an approximate evolutionary process based on this simplification showed that introducing more fitness classes and accounting for their specific distribution of neighboring fitness values are vital steps for overcoming the poor predictive accuracy of the percolation approximation.

2. *Coarse correlation with monomorphic population.* In a next step we have used correlation statistics of the landscape. This has enabled us to take into account that genotypes in different fitness classes are surrounded by different neighborhood structures. Although based on a coarse classification of fitness values into just a small number of fitness classes, while treating the population as being monomorphic, predictions improved relative to the percolation approximation. Yet, the simplified processes still resulted in too long waiting times.

3. *Fine correlation with monomorphic population.* A fined-grained correlation matrix served as the basis for the next step and was supposed to better represent the possibility of small changes in fitness values, which can be critical for describing the evolutionary process. Predicted time scales of evolution lay quite close to those of the actual evolutionary algorithm. However, the fact that the waiting times were systematically underestimated led us to conjecture that the simultaneous presence of different fitness values within a population is important for predicting evolutionary change.

4. *Fine correlation with polymorphic population.* Allowing the population to be polymorphic, so that individuals in one generation can belong to different fitness classes, resulted in remarkably accurate predictions of evolutionary waiting times. All the information needed for this very close approximation of the evolutionary algorithm is provided by a fine-grained correlation matrix.

In the following subsections I have then shown that descriptions based on one-dimensional correlation are equally successful for different mutation operators, and also for another Travelling Salesman Problem with increased number of cities and an evolutionary algorithm using a stochastic selection scheme.

Fitness landscapes with neutrality

On TSP landscapes, neighboring genotypes hardly ever have the same fitness, i.e. are neutral. Therefore, we turned our attention to fitness land-

scapes which allow for tuning the degree of neutrality in order to test if one-dimensional correlation approximations are equally successful under such different conditions.

In Section 2.2 we observed that for NKp fitness landscapes with low neutrality predictions based on one-dimensional correlation statistics are very close to actual evolutionary waiting times. By contrast, investigating an NKp landscape with high degree of neutrality has demonstrated that the fitness landscape decomposes into a set of large and extended networks of neutral genotypes that are too inhomogeneous for being described by mean neighborhood structures at different fitness levels. For the class of Royal Staircase Functions, which allow to construct fitness landscapes of high neutrality with well-structured networks of equal fitness, we also observed that one-dimensional correlation can only provide rough predictions of evolutionary waiting times.

In order to describe classes of individuals with similar mutant fitness, additional properties of individuals have to be introduced. In Chapter 3, I investigated highly neutral fitness landscapes and showed that higher-dimensional correlation statistics are appropriate for predicting a population's drift along neutral networks. For NKp landscapes with large-scale neutrality, I accounted not only for the fitness of individuals but also for their number of better, and, finally, also for the number of worse neighbors. Genotypes that are similar with respect to those three parameters also show a similar neighborhood structure. In the case of RSF landscapes I tried an equivalent approach and characterized individuals by their fitness and by the number of incorrectly set bits on the genotype vector for attaining higher fitness. For both types of landscapes with high degree of neutrality, an approximation of the evolutionary process based on three-dimensional correlation statistics proved to be successful and nearly accurate in predicting evolutionary waiting times. Describing fitness landscapes by these correlation statistics is conceptually very simple, and allows to considerably reduce the complexity of even highly neutral fitness landscapes.

Analytic descriptions

All predictions of evolutionary waiting times described so far were based on many realizations of stochastic processes using correlation information. I therefore developed an analytic description of the fitness dynamics on TSP and NKp landscapes with low neutrality, presented in Chapter 4. The one-

dimensional correlation matrices have, in both cases, shown that mutant fitness, as a very good approximation, follows a normal distribution. Moreover, mean and variance of mutant fitness depend linearly on parent fitness in a very close fit. On this basis, the correlation matrices were represented by two linear functions describing the mean and variance of mutant fitness. Due to this special structure of the correlation matrices, the fitness distributions in large populations can be approximated by normal distributions. By choosing an exponential selection function, I derived formulas for the change of mean and variance of normally distributed fitness values present in a large evolving population. These analytic predictions yielded very good approximation of the actual evolution of fitness distributions. By approximating the maximal fitness in a finite population whose fitness distribution is represented by a normal distribution, it was also possible to predict the evolution of the best fitness in a population.

Directions for future research

For landscapes of higher neutrality and landscapes with mutant fitness distributions that are not normal, I have, so far, not been able to derive analytic descriptions. As the correlation statistics provided nearly accurate predictions of waiting times for all different landscapes, it should in principle be feasible to analytically derive predictions for evolving fitness distributions, based on their moments. This requires an investigation of the structure of higher-dimensional correlation information of fitness landscapes with high degree of neutrality, and thus additional study.

I conclude that the crucial features of a fitness landscape are captured in a surprisingly accurate way by correlation statistics that describe the neighborhood structures of individuals. The success of these landscape statistics, allowing for almost perfect prediction of evolutionary waiting times, shows that this approach offers a new pathway for improving our understanding of complex fitness landscapes.

References

- Barnett L. (1997). Tangled Webs - Evolutionary Dynamics on Fitness Landscapes with Neutrality, MSc. diss., School of Cognitive and Computing Sciences, Sussex Univ. UK.
- Barnett L. (1998). Ruggedness and neutrality - the NKp family of fitness landscapes, *Artificial Life VI, Proceedings of the Sixth International Conference on Artificial Life*, MIT press
- Beasley D. (1997). Possible applications of evolutionary computation, In: "Handbook of Evolutionary Computation" (A1.2) Bäck T., Fogel D.B., and Michalewicz Z. (eds.) *Institute of Physics Publishing and Oxford University Press*
- Eigen M., McCaskill J., and Schuster P. (1989). The molecular Quasispecies, *Adv. Chem. Phys.*, **75**: 149 - 263.
- Fontana W., Stadler P.F., Bornberg-Bauer E.G, Griesmacher T., Hofhacker I.L., Tacker M., Tarazona P., Weinberger E.D., and Schuster P. (1993). RNA folding and combinatorial landscapes, *Phys. Rev. E*, **47**(3): 2083-2099.
- Gavrillets S. (1997). Evolution and Speciation on Holey Adaptive Landscapes, *Trends Ecol. Evol.*, **12**(8): 307-312.
- Gavrillets S. and Gravner J. (1997). Percolation on the fitness hypercube and the evolution of reproductive isolation, *J. Theor. Biol.* **184**(1): 51-64.
- Grimmett G. (1989). Percolation, *Springer*
- Happel R. and Stadler P.F. (1996). Canonical Approximation of Fitness Landscapes, *Complexity* **2**: 53-58.

- Huynen M.A., Stadler P.F., and Fontana W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation, *Proc. Natl. Acad. Sci. (USA)* **93**(1): 397-401.
- Kauffman S.A. and Levin S. (1987). Towards a general theory of adaptive walks on rugged landscapes, *J. Theor. Biol.* **128**: 11-45.
- Kauffman S.A. (1993). The Origins of Order - Self-Organization and Selection in Evolution, *Oxford University Press*.
- Kemeny J.G. and Snell J.L. (1960). Finite Markov Chains, *D. Van Nostrand Company*.
- Kesten H. (1982). Percolation theory for mathematicians, *Birkhaeuser*.
- Manderick B. (1997). Correlation analysis, In: "Handbook of Evolutionary Computation" (B2.7.3) Bäck T., Fogel D.B., and Michalewicz Z. (eds.) *Institute of Physics Publishing and Oxford University Press*.
- Michalewicz Z. (1992). Genetic algorithms + data structures = evolution programs, *Springer*
- Pasemann F., Steinmetz U., and Dieckmann U. (1999). Evolving structure and function of neurocontrollers, *Proceedings of the 1999 Congress on Evolutionary Computation* July 6th to 9th 1999, Madison, Washington DC, USA (1999).
- Reidys C.M. and Stadler P.F. (2001). Neutrality in Fitness Landscapes, *Appl. Math. & Comput.* **117**: 321-350.
- Rudolph G. (1997). Stochastic processes, In: "Handbook of Evolutionary Computation" (B2.2) Bäck T., Fogel D.B., and Michalewicz Z. (eds.) *Institute of Physics Publishing and Oxford University Press*.
- Sahimi M. (1994). Applications of percolation theory, *Taylor&Francis*.
- Schuster P., Fontana W., Stadler P.F., and Hofacker I.L. (1994). From sequences to shapes and back: A case study in RNA secondary structures, *Proc. Roy. Soc. Lond. b Bio.* **255**(1344): 279-284.
- Schuster P. (1996). How does complexity arise in evolution? *Complexity* **2**(1): 22-30.

- Schuster P. (1997). Landscapes and Molecular Evolution, *Physica D* **107**(2-4): 351-365.
- Stadler P.F. (1992). Correlation in landscapes of combinatorial optimization problems, *Europhys. Lett.* **20**(6): 479-482.
- Stadler P.F. and Schnabl W. (1992). The Landscape of the Travelling Salesman Problem, *Phys. Lett. A* **161**(4): 337-344 .
- Stadler P.F. (1995). Towards a Theory of Landscapes, In: "Complex Systems and Binary Networks" (Proc. of the Guanajuato Lectures 1995) López-Peña R., Capovilla R., García-Pelayo R., Waelbroeck H., and Zertuche F. (eds.), *Springer*.
- Stadler, P.F. (1996). Landscapes and their correlation functions, *J.Math.Chem.* **20**(1-2): 1-45.
- Stauffer D. and Aharony A. (1994). Introduction to percolation theory, *Taylor&Francis*.
- van Nimwegen E., Crutchfield J.P., and Mitchell M.(1999). Statistical Dynamics of the Royal Road Genetic Algorithm, *Theoretical Computer Science* **229**: 41-102.
- van Nimwegen E. (1999). The Statistical Dynamics of Epochal Evolution, PhD Thesis, www.santafe.edu/~erik/.
- van Nimwegen E. and Crutchfield J.P. (1999). Optimizing Epochal Evolutionary Search: Population-Size Dependent Theory, *Machine Learning Journal* in press.
- van Nimwegen E. and Crutchfield J.P. (2000). Optimizing Epochal Evolutionary Search: Population-Size Independent Theory, *Computer Methods in Applied Mechanics and Engineering* **186** (2-4).
- Weinberger E. (1990). Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference, *Biol.Cybern.* **63**(5): 325-336.