



Ingo Paenke

Dynamics of Evolution and Learning



Ingo Paenke

Dynamics of Evolution and Learning

Dynamics of Evolution and Learning

by
Ingo Paenke



universitätsverlag karlsruhe

Dissertation, genehmigt von der Fakultät für Wirtschaftswissenschaften der
Universität Fridericiana zu Karlsruhe
Tag der mündlichen Prüfung: 29.02.2008
Referent: Prof. Dr. Hartmut Schmeck
Korreferent: Prof. Dr. Xin Yao

Impressum

Universitätsverlag Karlsruhe
c/o Universitätsbibliothek
Straße am Forum 2
D-76131 Karlsruhe
www.uvka.de



Dieses Werk ist unter folgender Creative Commons-Lizenz
lizenziiert: <http://creativecommons.org/licenses/by-nc-nd/2.0/de/>

Universitätsverlag Karlsruhe 2008
Print on Demand

ISBN: 978-3-86644-247-4

*To my mentor
Daisaku Ikeda*

Acknowledgements

Looking back, I feel immense gratitude to the many people around me that have supported me during the time in which I was working on this Ph.D. thesis. The past three and a half years have been exciting, challenging and truly valuable. The thesis is the result of a joint project between the Institute AIFB at the University of Karlsruhe and the Honda Research Institute Europe (HRI) in Offenbach a.M., Germany. In this project I have been working alternately at the HRI - living in Frankfurt in these times - and the AIFB. I am grateful to the HRI for providing the complete funding.

I wish to sincerely thank my doctoral advisor, Prof. Dr. Hartmut Schmeck. Despite his busy schedule in an increasingly large research group, he frequently took off some time for valuable discussions. I am grateful to you, Prof. Schmeck, for continuously pointing me to those aspects that were essential for successfully finishing my thesis in the planned time. I also thank you for letting me pursue my studies freely and for letting me decide whether and to what extent I get involved into teaching. At the AIFB, I was intensively advised by Dr. Jürgen Branke. I remember well when I walked into Jürgens office in 2003 where he immediatly offered me to write a Diploma thesis under his supervision and contacted the HRI to initiate a joint Diploma thesis project which finally led to this Ph.D. thesis. Thank you, Jürgen, for your great enthusiasm, your brilliant analyzes and your serious examination of every proposed idea. I am very grateful for having had such a caring advisor.

At the HRI, I was intensively advised by Dr. Yaochu Jin and Prof. Dr. Bernhard Sendhoff. Thank you, Yaochu and Bernhard for your immense support. Yaochu, you have been pushing me at the right pace while never being impatient. At any time you by yourself were giving me a great example of how to finalize a piece of work and take the next concrete steps. I will miss our fruitful and cheerful discussions. Bernhard, I remember our great meetings which ended with many, many notes in my hands. Not only did you help me to open up various new perspectives for the next steps of my scientific work - what I remember vividly is how much you treasured the work that I had done so far. I always left your office happily, looking forward to continue with my work. These experiences have shaped my working attitude and I am grateful for having met such a scientific mentor. I also thank HRI Europe President Prof. Dr. Edgar Körner and CFO Andreas Richter for supporting my research project. A special thanks goes to Claudia Schäfer for her caring support in the early phase of the project.

I also thank Bernhard Sendhoff and Hartmut Schmeck for initiating my two months research visit at Prof. Dr. Xin Yao's research group at the University of Birmingham. Xin, I thank you for giving me this opportunity, for your advice and for letting me get involved in the various research areas of your group. I greatly enjoyed my time in Birmingham.

A very special thanks goes to Prof. Dr. Tadeusz Kawecki from the University of Lausanne. My collaboration with Tad started after his talk at the HRI in the early stages of my thesis project. Tad, thank you so much for opening a multi-disciplinary perspective for my thesis and also for critically reviewing my work from the point of view of an excellent evolutionary biologist. I greatly benefit from our collaboration.

I am grateful to my many wonderful colleagues at the different institutes. Everyone has supported me in his or her own way, sometimes by commenting on my work, sometimes with technical assistance, and sometimes by simply listening to my problems. Each of the following persons deserves to be mentioned in length and I wish I could do this here. Among these people at the AIFB are Berndt Scheuermann, Christian Schmidt, Michael Stein, Sanaz Mostaghim, Holger Prothmann, Urban Richter, Peter Bungert, Matthias Bonn, Andreas Kamper, Lei Liu, Stefan Thanheiser, André Wiesner, Lukas König and our secretary Ingeborg Götz. Among these people at the HRI are the ELTec research group members Lars Gräning, Stephan Menzel, Markus Olhofer, Martina Hasenjäger, Thomas Bihrer and Till Steiner, the non-ELTec members Inna Mikhailova and Xavier Domont, and the ELTec affiliated Ph.D. students Neale Samways, Ben Jones, Dudy Lim, Aimin Zhou. A very special thanks goes to Miguel Vaz, who enormously supported me scientifically, technically and above all as a great friend. Thank you, Miguel! Thank you, Ramon Sagarna, Per Kristian Lehre, Andreas Soltioggo and Arjun Chandra for your kind support in Birmingham. A special thanks goes to Chrisantha Fernando for taking the initiative for our refreshing and productive collaboration.

I am grateful to Miguel, Yaochu, Bernhard, Stefan, Lars, Jürgen, Berndt, Sanaz and Holger for reviewing and proofreading my thesis. I thank Prof. Dr. Clemens Puppe and Prof. Dr. Andreas Geyer-Schulz for being the examiners of my defense and I thank Prof. Dr. Hartmut Schmeck and Prof. Dr. Xin Yao for additionally being referee of my thesis.

I would not have been successful in this endeavor without the tremendous background support from outside my workplace. Words can not appropriately reflect what I owe to my parents - *Meine lieben Eltern, Danke für alles!* I am indebted to my best friend Kentaro and his family for their continuous encouragement and support. *Danke, Kentaro und Sophia!* I am grateful to my flat mates in Frankfurt, Elke, Miguel, Daniela, Judith, Phillip for their flexibility and their understanding for my sudden arrivals and departures, and to my great friends in Karlsruhe, Moritz, Bernhard and Chrisoula.

As a member of the Soka Gakkai International, a lay Buddhist movement and United Nations NGO that fosters peace, education and culture around the globe, I am indebted to many of its members who so strongly supported me in my personal development. As representatives of the countless people that should be named here, I want to thank Kentaro, Sophia, Birgit, Moonja, Georg and Sonja.

Finally, I wish to express my gratitude to my mentor in life, Daisaku Ikeda, who is the President of the Soka Gakkai International, a Buddhist leader, peacebuilder and educator who has received the United Nations Peace award, more than 200 honorary doctorates, professorships and other academic honors for his contributions to peace. Through his life he encourages me to create value and to challenge myself for this purpose. I dedicate this work to my mentor, Daisaku Ikeda.

Ingo Paenke, Karlsruhe, 2008

Contents

Nomenclature and Symbols	IX
1 Introduction	1
2 Fundamentals	5
2.1 Evolution	5
2.1.1 Principles and Definitions	6
2.1.2 Evolutionary Computation - Transfer of Biological Principles to Computation	9
2.1.3 The Evolutionary Algorithm	10
2.1.4 Genotype-Phenotype Distinction in Evolutionary Computation	11
2.2 Learning	13
2.2.1 Principles and Definitions	13
2.2.2 Benefits of Learning	14
2.2.3 Cost of Learning	15
2.2.4 Types of Learning	17
2.3 Influence of Evolution on Learning	18
2.3.1 Biological Perspective	18
2.3.2 Computational Intelligence Perspective	18
2.4 Influence of Learning on Evolution	19
2.4.1 Definitions	19
2.4.2 Biological Perspective	19
2.4.3 Computational Intelligence Perspective	22
2.5 Summary and Conclusion	23
3 Lamarckian and Biological Inheritance	25
3.1 Related Work	25
3.2 Conditions for Lamarckism	27
3.3 A Simplified Model	28

3.3.1	Model Description	28
3.3.2	Simulation Experiments and Results	31
3.3.3	Discussion	35
3.4	Summary and Conclusion	35
4	Influence of Learning on Evolution - The Gain Function Framework	37
4.1	Related Work	37
4.2	Basic Idea	41
4.3	The Gain Function Framework	43
4.3.1	Formulation	43
4.3.2	Proof	44
4.4	Extended Gain Function Framework	47
4.4.1	Formulation	47
4.4.2	Proof	48
4.5	Summary and Conclusion	52
5	Conditions for Learning-Induced Acceleration and Deceleration of Evolution	55
5.1	A General Learning Function	55
5.1.1	Directional Learning	56
5.1.2	Learning Noise	59
5.1.3	Simulation	61
5.1.4	Conclusion	62
5.2	Separable Fitness Components	63
5.2.1	Positive, Decreasing $f_L(x)$	64
5.2.2	Constant $f_L(x)$	64
5.2.3	Positive, Increasing $f_L(x)$	64
5.2.4	Conclusion	65
5.3	Influence of Learning Curves on Evolution	65
5.3.1	Extension of the Fitness Landscape Model	65
5.3.2	Modeling Learning Curves	66
5.3.3	Genotype-Independent Learning Curves	67
5.3.4	Genotype-Dependent Learning Curves	69
5.3.5	Conclusion	71
5.4	A Non-Monotonic Gain Function	72
5.4.1	Fitness, Learning and Gain Functions	72
5.4.2	Simulation	73
5.4.3	Results	73
5.4.4	Conclusion	74
5.5	Summary and Conclusion	74
6	Gain Function Analysis of Other Models of Evolution and Learning	77
6.1	Hinton and Nowlan's In Silico Experiment	77
6.1.1	Original Model Formulation	77
6.1.2	Model Reformulation	80
6.1.3	Gain Function Analysis	80

6.1.4	Discussion	82
6.2	Papaj's In Silico Experiment of Insect Learning	82
6.2.1	Model Formulation	83
6.2.2	Gain Function Analysis	84
6.2.3	Extended Gain Function Analysis	86
6.2.4	Continual versus Posthumous Fitness Assessment	87
6.2.5	Discussion	87
6.3	Mathematical Models with Developmental Noise	88
6.3.1	Existing Models	88
6.3.2	Discussion	89
6.4	Biological Data - An Inverse Gain Function Application	89
6.4.1	In Vitro Evolution of Resource Preference	89
6.4.2	A Qualitative Gain Function Analysis	90
6.4.3	In Silico Evolution of Resource Preference	92
6.4.4	Discussion	94
6.5	Mathematical Models on the Fitness-Valley-Crossing Ability	94
6.5.1	Problem of Large State Spaces in Markov-Chain Analyses	96
6.5.2	Difficulty of Deriving the Transition Probabilities in Markov-Chain Analyses	96
6.5.3	Borenstein's Approach	97
6.5.4	The Role of the Gain Function	98
6.5.5	Discussion	98
6.6	Summary and Conclusion	98
7	Balancing Evolution and Learning	101
7.1	Computational and Biological Evolution/Learning Trade-Offs	102
7.2	Related Work	103
7.3	Simulation Model	105
7.3.1	Evolutionary Adaptation	106
7.3.2	Genotype-Phenotype-Mapping	107
7.3.3	Individual Learning	107
7.3.4	Remarks	107
7.4	Influence of Lifetime on Population Dynamics	108
7.4.1	Influence of Learning on Diversity	110
7.4.2	Influence of Learning on Exploration/Exploitation	113
7.5	Existence of an Optimal Evolution/Learning Balance	115
7.5.1	Optimality of Pure Evolution	115
7.5.2	Optimality of an Intermediate Degree of Learning	118
7.6	Summary and Conclusion	120
8	Self-Adaptation of the Evolution/Learning Balance	123
8.1	Related Work	123
8.2	Extension of the Analysis Model	124
8.3	An Initial Experiment of Lifetime Evolution	124
8.4	Lifetime Evolution with a Reproduction/Lifetime Trade-Off	126

8.4.1	Evolution of the Optimal Lifetime in Environment 4	127
8.4.2	Evolution of the Optimal Lifetime in Environment 3	127
8.5	Summary and Conclusion	127
9	Conclusion and Outlook	131
9.1	Conclusion	131
9.2	List of Contributions	133
9.3	Outlook	135
A	Geometric Explanation for the Fitness Valley in Exp. 1 of Chapter 3	137
B	Proof of Equation 5.16	141
C	Calculation of the Derivative of Equation 6.21	145
D	Basins of Attraction in Environments 2 and 4 of Chapter 7	147
E	Simulation Results for Deterministically Changing Env. 4 of Chapter 7	151
	Bibliography	155

Nomenclature and Symbols

Symbol, Domain	Description (cf. Chapter 2)
$x \in X$	Genotype x is an element of genotype space X , it contains all heritable information, including some (but not necessarily all) information needed to develop the innate phenotype, and possibly (but not necessarily) parameters that influence learning (cf. parameter a below).
$z \in Z$	Phenotype z is an element of phenotype space Z . It is an individual's physical state in an environment at a time, and is subject to selection.
t	Time (modelled in discrete or continuous units).
$a \in \mathbb{R}^n$	Learning parameter a is either a single value ($n = 1$) or a vector ($n > 1$). It defines parameters that influence individual learning behavior. a can either be part of genotype x or externally given. In the models of this thesis these include the number of learning trials, Examples are individual lifetime, learning stepsize, and others.
$e \in E$	Environment parameter e is either a single value or a vector that may influence development, learning and the adaptive value (see below), e.g. the location of the current optimum.
$\phi : (X, E) \mapsto Z$	Development function ϕ describes the mapping from genotype to innate phenotype. The innate phenotype is determined by the genotype and in some models by the environment.
$l : (Z, X, E) \mapsto Z$	Learning function l describes phenotypic changes. The outcome of learning may be influenced by the environment and/or genotype, e.g., when parameter a is part of the genotype.
$v : (Z, E) \mapsto \mathbb{R}^+$	Adaptive value function v determines the <i>adaptive value</i> or <i>viability</i> of a phenotype under environmental influence at a time. The adaptive value determines the probability to produce offspring.
$f : (X, E^L) \mapsto \mathbb{R}^+$	Absolute fitness f indicates the fitness of an individual with genotype x , e.g., the adaptive values v of its corresponding phenotype accumulated during lifetime L . Corresponding to v , f is influenced by all environmental states of its lifetime, i.e., a vector E^L .
$w :$ $(X, (X, Z, t)^n, E^L)$ or $(X, (X, Z, t)^n, E^L, a)$ $\mapsto \mathbb{R}^+$	Relative fitness w is a measure for the expected number of offspring of an individual with genotype x in its lifetime L . Besides the individual's own genotype x , w depends on the state of the population (n individuals) at its birth, i.e., the set of genotypes X^n and phenotypes Z^n , the environment during its life E^L and the external learning parameter(s) a (if not genetically encoded).

CHAPTER 1

Introduction

“Evolution, however, is change in the properties of groups of organisms over the course of generations. The development [...] of an individual organism is not considered evolution: individual organisms do not evolve. Groups of organisms, which we may call populations, undergo descent with modification.”

Douglas J. Futuyma, Evolution [47, page 2]

The seemingly unbounded growth of computational processing power, data storage capacity, and computer networks has led to digital processing systems with an unprecedented complexity. There is no indication that this trend will change in the future. From a software-engineering perspective it is evident that beyond a certain complexity the behavior of these systems is not fully predictable. Furthermore, the application scope of these systems is steadily growing and we observe an increasing interconnectivity of digital systems with the natural world. As a consequence operating conditions are no longer constant but are frequently changing.

This development demands for computational systems capable to adapt to changing and unforeseen conditions. Biological information processing demonstrates great capabilities in this regard. The study of biology may inspire the design of computational systems that are highly adaptive as well. However, the laws of biological information processing are fundamentally different from the mechanisms of digital information processing. Rather than copying biology, promising biological principles need to be identified and tailored to the computational environment. The translation of biological principles to digital processing targeted toward problem solving has become a major subject of computer science. The various related approaches are often categorized under the umbrella of *Computational Intelligence*.

Evolution and learning are the two major mechanisms in natural adaptation. This thesis is devoted to an understanding of the dynamics of evolution and learning. Evolution is the change of the composition of heritable - genetic - information of a population of individuals over time. This change is driven by natural selection and by forces that introduce variation. Learning is the change of an individual's physical state - its phenotype - during its lifetime.

Compared to evolution, learning works on a much shorter time scale. The interplay between evolution and learning allows populations of organisms to adapt to various environmental changes.

The first transfer of principles of biological evolution to computer programs dates back half a century [43]. Since then, Evolutionary Computation has grown to an established discipline which develops algorithms that make use of principles of evolution for solving various complex optimization problems, the creation of art and music, process control, and a wide range of other applications. Two decades ago, Hinton and Nowlan [64] for the first time presented a computational model of evolution that incorporates individual learning. Their model demonstrates an adaptational advantage of the coupling of the two adaptation mechanisms.

Just before the work on this thesis started, a paper by Mery and Kawecki [107] was published that reports on a biological experiment on evolution and learning in fruit flies. This was the first time that biological experimental evidence for an influence of learning on evolution has been produced. The work reports on two experimental settings, one in which learning accelerates the rate of evolutionary change and one in which the opposite can be observed - decelerated evolutionary change in the presence of learning. A discussion with Tad Kawecki in the early stages of this work strongly influenced its research direction. It turned out that there was no satisfactory theory to explain his experimental results although a rich body of literature from different scientific fields is devoted to the study of advantages and disadvantages of learning in evolution. These studies approach the subject from various angles. Some of these studies come to the conclusion that learning accelerates evolution, while some conclude that learning decelerates evolution. There are also studies which describe experimental scenarios for both outcomes. So, the state of the art somehow confirms Mery's and Kawecki's results that learning does not generally accelerate or decelerate evolution. In each of these works, some explanation is provided as to what causes the respective results. However, the explanation derived from one study is not general enough to explain the results of others.

The aim of this thesis is to develop a unifying understanding of the dynamics of evolution and learning. It is based on the philosophy that simple models developed for the understanding of biological systems not only help to explain what we observe in nature but also to apply the biological principles in computation. The mathematical models of this thesis are proposed in this spirit. The simulation models developed in this thesis employ standard techniques from Evolutionary Computation aiming to ease the transfer of the studied biological phenomena to computational systems. Similarly, the simulation models are realized using standard techniques known from Evolutionary Computation. Therefore it is hoped that the presented analyses and studies contribute to an understanding of biological phenomena and serve as a basis for the transfer of principles of biological evolution and learning dynamics to computational systems.

Overview

In Chapter 2, the fundamentals of this thesis are introduced, not only as background but also to specify the working definitions and assumptions and to clarify the terminology used throughout the thesis.

In Chapter 3, adaptational effects of Lamarckian inheritance are studied. Lamarck [93] proposed that acquired properties are directly transferred from parent to offspring and that individual lifetime changes are the driving force of evolution. Although clearly rejected in evolutionary theory, Lamarckian inheritance is successfully employed in Evolutionary Computation. In Chapter 3, the conditions that favor Lamarckian and biological inheritance are studied thereby providing arguments why Lamarckian inheritance is often beneficial in evolutionary optimization even though it cannot be observed in nature. In the remainder of this thesis, biological (non-Lamarckian) inheritance is assumed. Chapter 3 also briefly discusses the reasons for this decision.

In Chapter 4, the Gain Function framework which represents the core of this thesis is introduced. The gain function is a mathematical framework that generally defines conditions under which learning accelerates or decelerates evolutionary change. It is formulated in terms of the influence of learning on the reproductive success of individuals (fitness) and considers how learning influences selection pressure. The central argument is the following: If genetically strong individuals benefit proportionally more from learning than genetically weak ones, learning accelerates evolution toward high fitness individuals. However, if weak individuals benefit more, evolution is decelerated. The gain function is formulated in a general way and can be applied to biological and computational models alike, although it naturally makes some simplifying assumptions.

In Chapter 5, several scenarios of coupled evolution and learning are studied using the gain function as an analysis tool. These scenarios are selected in order to cover a maximal range of typical environmental properties. In at least one setting, the gain function analysis yields a somewhat non-intuitive result.

In Chapter 6, several models from the evolutionary computation and biology literature are revisited and analyzed with gain function framework. The gain function perspective provides clear explanation of the results obtained from simulation studies. It also sheds some light on the result of Mery and Kawecki's evolutionary fruit fly experiment [107].

In Chapter 7, a further step toward the transfer of the dynamics of evolution and learning to computational paradigms is taken. There, the balance between the rate of evolutionary adaptation and the intensity of individual learning is studied. This issue is important in the presence of a trade-off between evolution and learning which arises from a computational resource conflict. The chapter concludes that in dynamic environments, the optimal balance depends on the type and rate of environmental change.

In Chapter 8, it is studied how the optimal balance between evolution and learning can emerge from a self-adaptation process. It is shown how the utilization of a biological principle - an individual energy trade-off between lifetime and reproduction - produces the appropriate conditions for successful self-adaptation.

Chapter 9 completes the thesis with conclusions and an outlook.

Major Contributions

In brief, the major contributions of this thesis are

- Explanation of the adaptational disadvantage of Lamarckism in rapidly changing environments,

- Formulation and proof of the gain function as a mathematical framework to predict the influence of learning on the rate of evolution,
- Identification of the conditions for learning-induced acceleration or deceleration for typical forms of learning,
- Theoretical underpinning of various studies of coupled evolution and learning,
- Discovery of a new type of adaptational advantage in presence of a resource-conflict between evolution and learning,
- Demonstration that biologically-plausible reproduction constraints allow successful self-adaptation of the evolution/learning balance.

See Chapter 9 for an extended review of the contributions of the thesis. The research approach of this work comprises simulation study and mathematical analysis. Subject of study are both computational and biological models of evolution and learning. The various perspectives constitute the multi-disciplinary nature of this thesis which allowed to make contributions to Mathematical Biology, Computational Biomodelling, Artificial Life, and Evolutionary Computation.

CHAPTER 2

Fundamentals

Evolution and *learning* are the two main adaptation processes that can be observed in nature and are also deployed in computational intelligence. *Evolution* is an adaptation process of the *genetic* composition of a *population*. In contrast, *learning* is an adaptation process of the *phenotype* of an *individual*. Considering a certain species that evolves and learns, the time scale on which the two adaptation processes operate is another point of distinction. Evolutionary adaptation takes place on a much larger time scale than individual learning (cf. Table 2.1).

This chapter provides an introduction to the principles of evolution in Section 2.1 and the principles of learning in Section 2.2. The coupling of evolution and learning produces a complex adaptive system. In the remainder of this chapter, the most well known aspects of the mutual influences in this system are reviewed, namely the influence of evolution on learning in Section 2.3 and the influence of learning on evolution in Section 2.4. Throughout this chapter the different aspects of evolution and learning are viewed from both the computational as well as from the biological perspective. This chapter is not intended to provide a complete review of theories and concepts of evolution and learning. Rather, it is tailored to the needs of this thesis. *Definitions* should therefore be understood as *working hypotheses* of this thesis.

2.1 Evolution

This section introduces concepts related to biological evolution and specifies the terminology of this thesis. The modeling approaches of this thesis are based on the standard Darwinian view of natural evolution that emphasizes the roles of variation and *natural selection* which is also often referred to as *survival of the fittest*¹.

¹ the term *Survival of the fittest* was actually coined by Herbert Spencer in his book *Principles of Biology* [162] after he had read Darwin's *The Origin of Species* [28]

Table 2.1: *Distinctive properties of evolution and learning*

Evolution	Learning
population-based	individual-based
genotype-level	phenotype-level
large time scale	short time scale

2.1.1 Principles and Definitions

Biological organisms are characterized by their physical manifestation which is defined as phenotype.

Definition 2.1 (Phenotype). *The phenotype of an individual is its physical state, including the physiology and behavior.*

The phenotype grows through a *development* phase, which is also known as *ontogenesis*. The resulting phenotype is strongly determined by its *genotype* which comprises the inherited characteristics of the organisms parents.

Definition 2.2 (Genotype). *The genotype of an individual is its set of heritable, also known as genetic, information.*

In biology, the genotype is often given in form of DNA (deoxyribonucleic acid). However, the development is not only influenced by genotypic information but also by the environment. The same genotype may result in a different phenotype depending on the environmental conditions under which development takes place. Abundant evidence is provided by several examples in West-Eberhard’s book [183]. There is a long-lasting debate in the biology literature about the relative influence of genes and the environment which is often referred to as the “Nature versus Nurture” debate [84, 88, 141]. A definition of development is given as follows.

Definition 2.3 (Development). *Development is the mapping from genotype to phenotype under the influence of the environment.*

Instinctively, the developed organisms struggle in order to produce offspring and thereby transferring their individual characteristics to progeny. As mentioned earlier, this process in which some individuals fail and some prevail is since Darwin known as *Natural Selection* [28] or simply *Selection*. Thus, the mechanism of selection determines which individuals reproduce offspring, i.e., pass on their genetic material to offspring. In nature, selection is an *intrinsic* mechanism of evolution which emerges from the struggle for survival and reproduction possibilities between individuals.

The concept of *fitness* is an attempt to capture the principles of natural selection in a theoretical framework. Haldane was the first to quantify fitness [53]. In agreement with similar definitions (e.g. [82]), the fitness of one particular genotype is quantified as the mean number of offspring of this genotype. A summarizing illustration of the relationship between genotype, phenotype and fitness is shown in Figure 2.1.

Haldane’s colleague Sewall Wright introduced the concept of the *fitness landscape*, or in Wright’s words the *adaptive landscape* [187] which visualizes the distribution of fitness values over a genotype *space*. Figure 2.2 illustrates such a fitness landscape for two dimensions




Figure 2.1: Relationship between genotype, phenotype and fitness. The phenotype is produced by the genotype under the influence of the environment and (natural) selection determines the fitness of a phenotype.




Figure 2.2: The Fitness Landscape, introduced by Sewall Wright [187]. Each point in genotype space maps to a fitness value. The mapping from genotype to fitness includes the more complex transformation with the phenotype as “stopover” (cf. Figure 2.1).

of the usually high-dimensional genotype space. The fitness landscape is often used as a means to picture the movement of a population of individuals on a landscape. In this image, individuals correspond to points in genotype space.

Since its introduction in 1932, the fitness landscape model has been subject to strong criticism in biology. See [87, 159] for recent surveys of the scientific discourse. The strongest argument of the criticism points to the concept of the population movement on a fitness landscape. Usually, fitness, in the sense of expected reproduction success, does not only depend on an individual's genetic configuration but is also dependent on other individuals. The same genotype (expressing a certain phenotype) in an otherwise identical environment may have different fitness values in different populations. Thus, a fitness landscape can only be drawn for one individual and under the assumption that the rest of the population is constant. This, however, contradicts the notion of population movement on the landscape. A similar argument has been described in [77] and [164]. For consistency, *two* concepts of fitness are employed in this thesis, namely *relative fitness* and *absolute fitness*.

Definition 2.4 (Relative Fitness). *Relative fitness is a measure for the expected reproductive success of an individual with a certain phenotype in a given population.*

Relative fitness refers to Haldane's fitness definition that measures the reproductive success. In the literature, this type of fitness has also been named *reproductive fitness* [164, 168].

Definition 2.5 (Absolute Fitness). *Absolute fitness is an individual lifetime measure for the survival and reproduction ability that can be evaluated independent of other individuals.*

The concept of *absolute fitness* allows to draw the picture of a population that moves on the (absolute) fitness landscape. The absolute fitness is similar to the concept of viability, which reflects an individual's strength and reproduction ability. A doubling of an individual's absolute fitness should therefore lead to an approximate doubling of its relative fitness, if other things are equal. Notice that this principle has become known in evolutionary computation as *fitness proportional selection* [36, page 59] which is discussed in more detail in Section 2.1.3. From a biologist's point of view, the concept of absolute fitness may be of little interest, because the difficulty of measurement makes it impracticable.

It is shown later that the concept of absolute fitness is indeed more useful in the realm of evolutionary computation, where the absolute fitness is usually *assigned* by a certain evaluation function. For convenience, the term *fitness landscape* is reserved for the mapping from genotype or phenotype to the *absolute* fitness value in this thesis.

The transfer of individual characteristics from parents to offspring through genes is known as *heredity*. Fit individuals produce offspring. This offspring is unlikely to have an identical genotype as its parents. First, in case of sexual reproduction, an offspring's genotype is a composition of the parent's genotypes which is also known as *recombination*. Secondly, when a parent's genetic information is replicated *copy errors* arise which is also known as *mutation*. It should be noted here that there are other sources of mutation. In summary, the whole reproduction process generates a *variation* of the genetic material. In case of sexual reproduction, there are at least two sources of genetic variation, namely recombination and mutation, where the asexual reproduction mutation is the major source of variation. However, the main point that should be emphasized here, is the *existence* of variation mechanisms.




Figure 2.3: *Abstract model of natural evolution. The population varies over time through the influence of selection and variation.*

There is no generally accepted definition of *evolution* but most definitions emphasize the *genetic changes in populations* or in Darwin’s words “descent with modification” [28]. As described above, there are three basic ingredients for evolution, namely *selection*, *heredity*, and *variation* that work to modify a population of individuals. In the following, a *definition* of evolution is proposed that more adequately includes all aspects relevant in the various models employed in this thesis.

Definition 2.6 (Evolution). *Evolution is the change in the composition of the genetic information of a population as a result of selection and variation, formally:*

$$\text{Population}(t + 1) = \text{Variation}(\text{Selection}(\text{Population}(t)))$$

This definition of evolution is illustrated in Figure 2.3. A population’s composition at time t may change within the next time step as a result of selection and variation. Usually the elements of evolution are modeled as having random components.² Notice that a discrete time model is employed in this formulation. Thus, essentially evolution is a repeated cycle of selection and variation. It should also be pointed out that natural evolution inherently requires a population. Without population no selection can take place.

2.1.2 Evolutionary Computation - Transfer of Biological Principles to Computation

The principles of natural evolution as described in the previous section can be implemented in computer programs which perform a “digital” evolution. The application of principles of natural evolution to computers is commonly summarized under the term *evolutionary computation*. There are at least two motivations for evolutionary computation.

First, the *principles* of natural evolution can be employed after appropriate modification for optimization, adaptation and creation in different domains such as, engineering, economics, medicine, and artificial life. The main tool for this is the *evolutionary algorithm*.

² The discussion on whether there is “real” randomness in nature (or randomness is just a model in order to deal with the enormous complexity of cause-and-effect relationships) is beyond the scope of this thesis.

Algorithm 2.1: Canonical Evolutionary Algorithm

input : Population size, Evaluation function, Specification of Mutation, Recombination and Selection
output: *Parents*

- 1 INITIALIZE(*Offspring*)
- 2 **repeat**
- 3 EVALUATE(*Offspring*)
- 4 *Parents* = SELECT(*Offspring*)
- 5 *Offspring* = RECOMBINE(*Parents*)
- 6 MUTATE(*Offspring*)
- 7 **until** *termination condition satisfied*

Secondly, biologists may use such programs as a tool to replicate natural evolution and thereby gain a deeper understanding of natural evolution itself. This type of simulated evolution is also known as *in silico* evolution.

2.1.3 The Evolutionary Algorithm

The evolutionary algorithm (EA) comprises the main elements of natural evolution described above. In the following, an EA as formulated in pseudo-code notation in Algorithm 1 is briefly described.

An EA usually starts with randomized initialization of a *set of solutions*, which is referred to as *population of individuals* in the biological metaphor. In the pseudo code notation, this population is called *Offspring*. The individuals of the *Offspring* population are evaluated with respect to a certain optimization, adaptation or creation task. This evaluation is commonly called *fitness evaluation*. Based on the evaluation result which is denoted f , individuals are selected as parents to produce individuals for the next generation. The higher the fitness of a solution, the more likely it is to be *selected*. In particular, this thesis employs the *fitness proportional selection* scheme, see e.g. [36, page 59]. Under the assumption of constant population size, the application of this selection scheme implies that the expected number of offspring of an individual is f/\bar{f} where f denotes the individual's quality according to the evaluation and \bar{f} denotes the average quality across all individuals of the population. One way to implement this scheme is Baker's stochastic universal sampling algorithm [6]. The individuals of the *Parents* population are combined. In analogy to biological evolution, this combination process is named *Recombination*. The resulting individuals are randomly altered which can be interpreted as *Mutation*. The resulting population represents the next generation of solutions. (Note that other versions of the EA allow to include parents from the current generation in the next generation.) This loop is repeated until a certain termination criterion is satisfied. Typical termination criteria are the excess of a predefined maximum computation time, the population convergence to a small region of the search space or a lack of fitness improvement in successive generations.

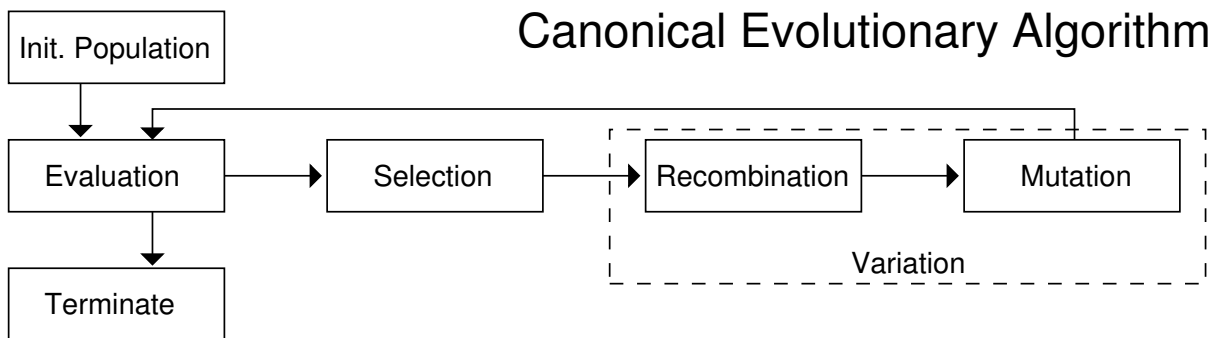


Figure 2.4: Flowchart of a canonical evolutionary algorithm. Similar to natural evolution, population search, and adaptation are realized through a cycle of selection and variation.

It is expected that through repeated application of this evolutionary loop of selection and variation, the population discovers points in the search space that correspond to high-quality solutions. The simplified canonical evolutionary algorithm is illustrated as a flowchart in Figure 2.4, highlighting that the EA basically processes repeatedly a cycle of selection and variation.

Note that the common usage of the term *fitness evaluation* in the EA corresponds to the concept of absolute fitness, cf. Definition 2.5. After the application of the selection operator, the relative fitness as defined in Definition 2.4 can be measured. So, the term fitness is commonly used in the sense of absolute fitness in the realm of evolutionary computation and in the sense of relative fitness in biology. Relative fitness can only be assessed by applying both the quality evaluation function and the selection function. So the *fitness proportional selection* scheme is more appropriately called *absolute fitness proportional selection*. Notice that in Evolutionary Computation various selection mechanisms have been proposed [11].

The relationship between the biological and the evolutionary computation perspective has also been pointed out by Colombetti and Dorigo [21, page 24]

“In the realm of artificial agents, the relationship between fitness and reproductive success is reversed: first, the fitness of an individual is computed as a function of its interaction with the environment; second, the fittest individuals are caused to reproduce. However, this pattern is not exclusive of artificial systems: it is applied by breeders (of cattle, horse, dogs, etc.) to produce breeds with predefined features. In fact, we contend that best metaphor of evolutionary computation is not biological evolution, but breeding.”

The evolutionary algorithms studied in this thesis have been implemented with the C++ open source programming library *EALib* which is part of the *SHARK* software package [155].

2.1.4 Genotype-Phenotype Distinction in Evolutionary Computation

In analogy to biological evolution, the distinction between *genotype* and *phenotype* is also often realized in EAs. In this case, a solution has two representations, the *genotype representation* which in evolutionary computation is often simply referred to as *representation* and the

phenotype representation which is often referred to as *solution* [14]. Evolution searches through the genotype space, hence, mutation and recombination are defined on the genotype representation. The quality of a solution is evaluated based on the phenotype representation. Thus, before an individual is evaluated, its genotype is transformed to the corresponding instance in the phenotype representation. This *genotype-phenotype-mapping (GPM)* is the analog to *development* in biological evolution. Complex genotype-phenotype mappings which resembles biological development in evolutionary computation seem to be of growing interest for evolutionary computation applications [92]. The genotype-phenotype distinction is often referred to as *indirect encoding* in evolutionary computation.

There are several reasons why an indirect encoding may have an advantage. Usually, the evaluation function expects a certain input format. This input format can be called the *natural representation* [110, 108]. There are at least two reasons why the natural representation might be inappropriate. These are described in the following.

Lack of Causality

The concept of *causality* originates from physics and means that small changes in the parameters of a system correspond to small changes in the system's performance. This concept can be transferred to the relationship between genotype and phenotype [153, 152] or between genotype and fitness [179, 139]: Here, *causality* or *strong causality* [153, 152] means that a small change in the genotype (the system parameters) produces only a small change in the phenotype or in the fitness (the system performance) of an individual. This means a single mutation in the genotype causes only a moderate fitness change. Causality is a prerequisite for evolutionary optimization [153]. Often, the natural representation makes it difficult to design mutation (or other variation) operators that provide causality in the mapping from genotype to fitness. In this case, the addition of a genotype level which maps to the phenotype space may help to construct the lacking causality.

Large Size of the Search Space

Often the natural representation specifies a solution in every detail. This produces a large search space which may be difficult to search successfully. By introducing the genotype as an additional representation layer, the space on which mutation and other variation mechanisms operate can be significantly reduced.

An example of such a case is the optimization of geometrical designs, such as turbines or wing airfoils. There, the fitness is usually assigned based on computational fluid dynamics (CFD) simulations. The simulation programs expect a certain input format of the geometrical design; usually a 2-d or 3-d grid representation that specifies all points of the grid. If the grid representation which can be considered as natural representation here is directly used for evolutionary search, the resulting search space becomes very large. An EA is unlikely to find a high quality solution under these conditions.

An alternative way to represent a geometry is to describe it by curves, whose shapes are controlled by a relatively small set of parameters, e.g., spline-curves [123, 81]. If the transformation function from the curves representation to the natural representation is known,

this small set of parameters fully describes a geometry. Evolutionary search on this (genotype) representation is often more successful than on the natural representation.

A second example is the evolution of artificial neural networks (ANNs). (It is here and in the remainder of this thesis assumed that the reader is familiar with the basic properties of ANNs. A comprehensive introduction is Haykin's book [62].) There are many examples in which a direct encoding, i.e., a complete specification of all properties of the ANN, leads to unsatisfactory evolutionary search. Various ways to indirectly encode an ANN with a relatively small set of parameters have been proposed such as a *parametric representation* [56] or *developmental rule representation* [83]. For a comprehensive survey see [188].

2.2 Learning

In the context of this thesis, learning should be understood as an *adaptation process of an individual agent*, i.e., either a natural organism or an individual in artificial evolution. Corresponding to Douglas J. Futuyma's quote [47, pg. 2] at the outset of this thesis, individual adaptation "is *not* considered evolution" - in contrast to the selection-variation loop of populations. *Similar* to evolution, however, learning should on average lead to some kind of improvement. In the following, learning and associated concepts are defined.

2.2.1 Principles and Definitions

In the context of evolution, learning should after all lead to an improvement of fitness, more precisely the absolute fitness, cf. Definition 2.5. However, since the concept of fitness *aggregates* over the whole lifetime of an individual, it is not appropriate to describe the lifetime *process* of learning.

More appropriate is the concept of *adaptive value*. In biology, the adaptive value refers to the degree to which a certain phenotypic characteristic helps an organism to survive and reproduce. During its lifetime, an individual may improve these characteristics, thereby increasing its adaptive value. In this thesis, the differentiation between components that contribute to fitness is of minor importance. More important is the temporal aspect of the adaptive value that is emphasized in the following definition.

Definition 2.7 (Adaptive Value). *The adaptive value is an individual's contribution to its absolute fitness at a time.*

Thus, if one draws the adaptive value of an individual against its age, learning curve is obtained. Based on this definition, a formal definition for learning can be specified.

Definition 2.8 (Learning). *Learning is an individual adaptation process that takes place on the phenotype level and is directed toward an increase in an individual's adaptive value.*

With Definition 2.7 it is self-evident that an increase in the adaptive value leads to an increase in the individual's absolute fitness. However, this may not necessarily lead to an increase in relative fitness, as it is shown in later chapters.

The same terminology can be applied to computational learning. For example, in a learning algorithm that is processed for a certain number of learning steps, the quality of a solution is

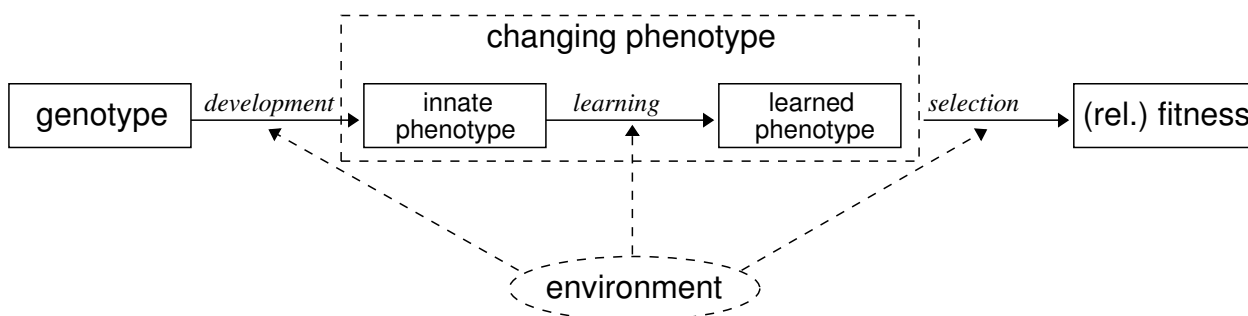


Figure 2.5: *Illustration of the relationship between genotype, phenotype and fitness that accounts for learning. The innate phenotype is produced by the genotype under the influence of the environment. The innate phenotype is modified through learning. Selection may act over time and not just on the learned phenotype.*

expected to increase with the number of learning steps. The *increase in solution quality over the processing time of the algorithm* can be rephrased as an *increase in adaptive value over lifetime*.

Figure 2.5 illustrates the relationship between genotype, phenotype, and fitness on a more fine-grained level that accounts for phenotype changes caused by learning, cf. Figure 2.1 for a coarse-grained illustration that does not account for learning. In biology, the transformation from genotype to phenotype is usually enormously complex. Development (ontogenesis) and learning (epigenesis) are parallel processes during the entire life time of individuals. There is no transition when one ceases and the other one starts. Nevertheless, in order to focus on the role of learning and for the sake of simple analysis, in this thesis the two processes are modeled as taking place in a sequential fashion: first development and then learning.

By definition, learning is directed toward an increase in the individual’s adaptive value and therefore an increase in the individual’s absolute fitness. However, there are various intermediate effects of learning that lead to this average increase in fitness. Some of these intermediate effects may actually be detrimental with respect to the individual’s adaptive value. The learning-induced increase in absolute fitness is therefore just the positive balance between benefits and cost of learning. Again, if this balance would be negative, there would be no point in learning. The benefits and cost of learning are discussed in the following.

2.2.2 Benefits of Learning

An obvious advantage of learning for an individual is that adaptation to its specific environmental conditions is possible [1, 102] which is not possible by evolutionary population-based adaptation. Furthermore, the genetic search mechanisms mutation, recombination (both variation operations), and selection may be inappropriate for a fine-grained adaptation. Thus, learning provides a clear benefit in this sense. This advantage applies to biology and computational intelligence alike.

Besides this, learning usually provides an adaptational advantage on the temporal scale. That is to say, learning allows to adapt quickly to changing environmental conditions [1, 102, 174]. In [173] and [167], it is concluded that learning can only provide an adaptational

advantage if the environmental dynamics are predictable to some degree. This advantage certainly applies to biology, and to those scenarios in computational intelligence, where population adaptation is applied to a quickly changing environment.

2.2.3 Cost of Learning

The cost of learning have received relatively little attention in the literature. However, in the few papers that study this issue a wide range of types of learning cost are discussed. From a purely biological point of view, Johnston [76] discusses six types of cost, namely

- Delayed reproductive effort and/or success
- Increased juvenile vulnerability
- Increased parental investment in each offspring
- Greater complexity of the central nervous system
- Greater complexity of the genome
- Developmental fallibility

and presents evidence for most of these. Mainly based on Johnston's work, Mayley [102] discusses several types of learning cost from an interdisciplinary point of view that considers both biology and computation. He groups these types of costs as follows

- Costs that are a function of the time spent for learning, e.g., time-wasting costs, delayed reproductive effort, energy costs
- Catastrophic costs, e.g., unreliability costs, damaging behavior
- Constant costs, e.g., increased ontogenetic costs
- Individual non-specific costs, e.g., parental investment, increased genotype length
- Non-evolutionary costs, e.g., program development/testing, CPU time

For details it is referred to the original papers [76] and [102].

In the following, the various types of learning cost are grouped with respect to whether they influence individual fitness or not. As it will be seen later, this categorization is tailored to the modeling approaches taken in this thesis.

Fitness Cost of Learning

Fitness cost of learning are those that can be modeled by a decrease in individual fitness. The two subcategories *energy consumption* and *exploration cost* cover most of the various cost aspects that arise on the individual level.

Energy consumption cost include cost for the development and for the maintenance of the learning system [76], e.g., brain or artificial neural network, as well as for the process

of learning itself. Organisms have a finite amount of energy available. An increase in the proportion of energy spent for learning implies a decrease in the proportion of energy spent for other activities. An example with obvious evolutionary consequences is the reduction of reproduction effort. In other words, *individuals reproduce less during learning*. Similarly, the survival probability may also suffer from increased learning effort. To some extent this type of learning cost also applies to computational intelligence because both digital replication for offspring production and learning demand computational resources, i.e., incur computational cost. It certainly applies to embodied computational intelligence, see, e.g., [174, 101].

Learning incurs various types of exploration cost. If not completely supervised (see Section 2.2.4), learning requires a certain degree of exploration in order to achieve improvement. Exploration bears the risk of trying out a worse solution than the current one. Therefore, an individual might experience setbacks or failures during the process of learning and the learning curve of a certain individual, i.e., its mapping from age to adaptive value may temporarily decrease, see [19, 174] for examples. However, in the end learning should yield an increase in the average adaptive value.

Non-Fitness Cost of Learning

Not all costs that arise from learning can be modeled as a decrease in individual fitness. The best example from an evolutionary computation point of view is CPU time which was classified by Mayley [102] as non-evolutionary cost.

Obviously, if available CPU time, or more generally computational resources, are limited, an increase in individual learning such as an increase in the number of iterations in a learning algorithm, implies that less computational resources are available for evolutionary adaptation.

The following reasoning holds for both, biological and computational scenarios. If the population size is more or less constant or has reached a certain limit, an increase in individual lifetime decreases the rate of evolutionary change, i.e., genetic change through variation and selection. The straightforward explanation is that with long lifetimes less individuals perish per time, hence less individuals can be born without breaking the population's size limit. In nature, the population size may be limited due to finite space, food resources etc.

However, the increased lifetime which can be interpreted as individual learning time does not reduce individual absolute fitness - unless such a reduction is externally assigned as possible in evolutionary computation. Instead, this type learning cost reduces the velocity of evolutionary adaptation for the population as a whole.

The conflict between resources allocated for evolutionary adaptation and resources allocated for individual learning has some relation to the above mentioned biological cost of energy consumption. In both cases, an increase in learning intensity implies a decrease in reproduction per time.

It should be noted here that the effect of non-fitness cost of learning on evolutionary dynamics cannot be explained by the standard model of evolution that is based on natural selection of *individuals*. However, there are models of evolution that consider larger units of selection, such as *group selection*, e.g., [104, 186] from the biology literature and [20, 95] in the computational intelligence literature, or even *species selection* [47, page 259].

2.2.4 Types of Learning

Learning theory assumes that the learning system has inputs and produces outputs. In both, biological and artificial systems typical inputs are sensor data, typical outputs are actions or decisions. In theory, it is often distinguished between three types of learning, namely *supervised learning*, *unsupervised learning* and *reinforcement learning*.

Supervised Learning

In supervised learning, a learning individual receives a teaching signal in addition to its (sensory) input. This teaching signal indicates the appropriate output, i.e., the teacher tells the learning individual what it should do. The goal of supervised learning is to construct an internal input-output mapping that can reproduce the input-output relationship that was provided by the teacher. This type of learning is widely used in machine learning applications, see e.g. [69]. The most well-known example is probably the back-propagation learning in artificial neural networks which is first described by Werbos [182] and further developed by Rumelhart et al. [147]. Supervised learning also appears in many animal species, e.g., where parents teach their offspring.

Unsupervised Learning

In unsupervised learning, the learning individual receives no information about the appropriateness of its output. Rather than “learning the right action”, the goal of unsupervised learning is “to extract an efficient internal representation of the statistical structure implicit in the inputs.” [65]. In machine learning, this for instance includes the discovery of statistical properties of (input) data, such as clusters. In biology, unsupervised learning takes place when developing animals learn to reduce visual input appropriately for further cognitive processing.

Reinforcement Learning

In reinforcement learning, a learning individual receives some information about the appropriateness of its output that was produced as a reply to an input stimulus. However, unlike the case of supervised learning, the individual is not taught *what* that appropriate action was. It only receives a feedback that indicates *how* appropriate the produced output was. In most of the cases in biology, there is no supervision available. Therefore, reinforcement learning is often observable in biology. In machine learning several reinforcement learning algorithms have been proposed and there is a wide range of applications [169].

Although there are several examples for each of the three types of learning in both machine learning and biology, it must be noted that the terminology is more common in the realm of machine learning [2]. However, an interesting hypothesis was proposed by Doya [32, 33] who argues that some brain regions are dominated by a certain operation “method” [32]. In particular, Doya claims that the network architecture of the *cerebellum* is specialized for supervised learning, the *basal ganglia* for reinforcement learning and the *cerebral cortex* for unsupervised learning.

2.3 Influence of Evolution on Learning

From a biological point of view, there is a simple answer to how evolution influences learning: Learning has evolved! This means the learning ability is the product of evolution. Research in evolutionary biology tries to answer, *how* learning has evolved, or in particular how learning in a certain species, a certain learning mechanism etc. has evolved. A prerequisite for learning is *phenotypic plasticity*. Thus, the evolution of phenotypic plasticity is another important research topic in this context. Ignoring genetic drift [94, 144], learning only evolves, if it provides a selective advantage. Hence, a condition for the evolution of learning is that its benefits outweigh the cost of learning. Recall Sections 2.2.2 and 2.2.3 where benefits and cost of learning have been discussed.

2.3.1 Biological Perspective

There are several studies that investigate the biological evolution of learning. One of the few examples that studies the biological evolution of learning with an *in vitro* experiment is Mery and Kawecki's work on the evolution of learning ability in fruit flies [106]. The experiments demonstrate that learning ability can indeed evolve. A similar evolution study with fruit flies [107] by the same authors is discussed in Section 6.4 of this thesis. A comprehensive review of the state of the art in biological evolution of learning is, however, beyond the scope of this thesis. The reader is referred to recent surveys [113, 134].

2.3.2 Computational Intelligence Perspective

The evolution of learning has also been investigated in artificial systems. The goal of such studies is either to gain new biological insights or to make progress in the design of adaptive technical systems. Artificial evolutionary systems naturally have constraints with regard to what *can* evolve. These constraints are either inherently present due to the limitation of computational resources, or they are set by the human designer - intentionally or accidentally. Correspondingly, different degrees of freedom with regard to how learning can evolve can be found in the literature.

The majority of the work employs artificial neural networks (ANNs) that perform a certain computational task during their lifetime. In its simple form, an ANN learns by adjusting its synapse weights and thereby modifies its input-output relation.

In the extreme case, the behavior of the ANN is fully genetically specified and only evolutionary adaptation is possible with no room for individual learning. This setting is known as *evolutionary learning*³. Several examples for evolutionary learning of synapse weights can be found in Yao's 1999 review [188, Section II.A-II.B]. To the knowledge of the author of this thesis, applications of the stand-alone evolution of ANN synapse weights, i.e., without any additional search mechanisms have rarely been published in recent years.

³Note that this usage of the term evolutionary learning differs from the terminology of this thesis where this type of adaptation is simply denoted as evolution and the term learning is reserved for individual level adaptation

The first step toward evolution *of* learning is to evolve some parameters of an ANN's learning algorithm. One example is the evolution of parameters of the backpropagation algorithm [55]. Others can be found in [188, Section IV].

More degrees of freedom for the evolution of learning are provided if not only parameters of a given learning algorithm but also the *learning rules* itself can evolve. Some examples for this category can be found in [188, Section IV]. In recent years, the evolution of learning rules has become an important issue in the field of *Evolutionary Robotics* [61, 119, 39, 180] where neural control systems are generated using experimental evolution. Several studies have demonstrated that under dynamic conditions, it is more appropriate to let the robot *learn* a good synapse weight configuration during lifetime using an evolved learning rule [40, 41, 42].

2.4 Influence of Learning on Evolution

In this section, the two mechanisms by which learning influences evolution, *Lamarckism* and the *Baldwin effect* named after the influential biologists Jean-Baptiste Lamarck (1744-1829) and James Mark Baldwin (1861-1934) are introduced. First, in Section 2.4.1 definitions for the two mechanisms are proposed. A detailed explanation follows in the remainder of this chapter. In Section 2.4.2, a biological perspective is presented based on a brief historical review of evolutionary theory and evidence related to the influence of learning on evolution. After that, in Section 2.4.3, a computational perspective is developed.

2.4.1 Definitions

The main effects, Lamarckism and the Baldwin effect, by which learning influences evolution are defined in the following.

Definition 2.9 (Lamarckism). *Lamarckism is the transfer of an individual's learned properties to its offspring.*

Some more precise definitions that distinguish between *weak Lamarckism*, *pure Lamarckism*, and *no Lamarckism* are developed in Chapter 3. In the literature, the Baldwin effect is mostly described qualitatively as a broad concept and there is no precise definition of it. Here, two definitions are suggested, one in the broader sense and one in the narrow sense.

Definition 2.10 (Baldwin Effect in the broader sense). *In the broader sense, the Baldwin effect is defined as a change in evolutionary pathways or the rate of evolution caused by individual learning in the absence of Lamarckism.*

Definition 2.11 (Baldwin Effect in the narrow sense). *In the narrow sense, the Baldwin effect is defined as a change in evolutionary pathways or the rate of evolution caused by individual learning, and the genetic fixation of previously learned properties through natural selection toward a reduction of the cost of learning (in the absence of Lamarckism).*

2.4.2 Biological Perspective

The biological perspective on the influence of learning on evolution has developed over the last 200 years with roughly one crucial finding every 50 years.

Lamarck (1809)

In his most influential work on evolutionary theory [93], Jean-Baptiste Lamarck (1744 - 1829) emphasized the strong role of the environment and the organism's capabilities to adapt to the local environmental conditions. Essentially, Lamarck argued that environmental change results in a behavior adaptation which causes a modification in the use of organs. The modified use of organs changes the organs "form" in the long run. Since at the time of Lamarck no theory of heredity had yet been developed, he assumed that this organic change is transmitted to offspring. He concluded that adaptive lifetime changes "[...] are preserved by reproduction to the new individuals [...]" [93]. Thus, in his view, individual lifetime adaptation to environmental conditions are the driving force for evolutionary change.

Darwin (1859)

A half century later Charles Darwin (1809-1882) extended Lamarck's theory [28]. As it is introduced earlier in this chapter, Darwin proposed that the driving force for evolutionary adaptation is the interplay of variation and *natural selection*. He assumed that from a population with variations of traits those which have an adaptational advantage are more frequently reproduced. Notice that Darwin had no valid theory of heredity available either. Although not explicitly stated, Darwin saw no significant influence of individual's lifetime adaptation or "learning" on evolutionary change.

Baldwin (1896)

A synthesizing view was developed another half century later by James Mark Baldwin (1861-1934). His main proposal [7] was that individual learning can change the evolutionary pathways of a species, even in the absence of Lamarckism, because learning influences fitness. Furthermore, he argued that learning involves cost. Selection acts to reduce the learning cost and the previously learned behavior eventually becomes instinctive. This mechanism is another half century later named "The Baldwin effect" by George Gaylord Simpson [158]. The reader is referred to the review by Depew [29] who argues that it could as well have been named after other biologists of the late nineteenth century.

Crick (1958)

Although Nobel Prize winner Francis Crick was not directly involved with the question how learning influences evolution, his publication of the *central dogma* of molecular biology (first in 1958 [23] and reformulated in 1970 [24]) provided important evidence for this question. The central dogma states that the information flow from DNA to Protein is uni-directional. Figure 2.6 illustrates this relationship (Figures 2.6(a) and (b) are redrawn from Crick's original article [24]). DNA influences protein (after translation to RNA and in rare case directly), but protein has no influence on DNA. The arrows in Figure 2.6(a) show all theoretically possible transfers between the three families of polymers (DNA, RNA and Protein). Figure 2.6(b) shows the actually observed information flow in nature, where the solid lines are based on clear evidence and the dashed lines represent either special cases or are based on uncertain evidence. Despite extensive research efforts, this figure remained almost unchanged since its




Figure 2.6: *Illustration of the Central Dogma of molecular biology formulated by Crick, Figures a) and b) adapted from [24]. a) shows the theoretically possible information flow if all three families of polymers would influence each other; b) shows the actually observed information flow in nature (solid lines are based on clear evidence, dashed lines represent special cases or are based on uncertain evidence); c) is a simplification of b), omitting the intermediate state (RNA); d) shows the simplified conclusion of c) with regard to the relationship of genotype and phenotype.*

publication in 1970. In Figure 2.6(c), the intermediate state (RNA) in the transition from DNA to Protein has been omitted, thereby emphasizing the uni-directional information flow from DNA to Protein. Since DNA is the carrier of genetic information and Protein represents the elementary unit from which cells (which make up the phenotype) are built, the central dogma can also be visualized as in Figure 2.6(d). *The genotype influences the phenotype, but the phenotype does not influence the genotype.* Despite the lack of evidence for “backward translation” from protein to genotype, Crick provides theoretical arguments why this cannot be observed in nature. He points out that the forward translation from DNA to RNA to Protein involves a very complex machinery and that it is unlikely that this machinery works backwards. An alternative could be the existence of an “entirely separate set of complicated machinery” [24] for back translation. However, there is no trace that such a machinery exists. With Crick’s findings, Lamarckism is to be rejected.

Recent Biological Perspectives (1986-2007)

Although Crick’s and other findings clearly reject Lamarckism, there has recently been discovered a range of “Lamarckian-like mechanisms” that occur without breaking the central dogma. These can roughly be categorized as *sustaining heritable epigenetic variation* [73], *phenotypic memory* [74] and so called *neo-Lamarckian inheritance* [140]. Examples include, mutational hotspots and adaptive mutations occurring during bacterial stress [45], chromatin marks that control differentiation in multi-cellular organisms [68], RNA silencing allowing potential influence by somatic RNA on germ line gene expression [97], inheritance of immune system states by antibody transfer in breast milk [160], and behavioral and symbolic inheritance systems such as food preference, niche construction traditions and all information transmission dependent on language [73].

Also recently, the first evidence for the Baldwin effect from a biological experiment has been produced. Mery and Kawecki [107] demonstrate the Baldwin effect in the *in vitro evolution* of fruit flies. Particularly, they show that *learning* of a resource preference can speed up the

Algorithm 2.2: Canonical Memetic Algorithm

input : Population size, Evaluation function, Specification of Mutation, Recombination, Selection, Learning
output: *Parents*

- 1 INITIALIZE(*Offspring*)
- 2 **repeat**
- 3 INDIVIDUAL LEARNING(*Offspring*)
- 4 EVALUATE(*Offspring*)
- 5 *Parents* = SELECT(*Offspring*)
- 6 *Offspring* = RECOMBINE(*Parents*)
- 7 MUTATE(*Offspring*)
- 8 **until** *termination condition satisfied*

evolution of the innate resource preference. Their study is revisited in Section 6.4 of this thesis.

2.4.3 Computational Intelligence Perspective

The first computational model that demonstrates the Baldwin effect has been published by Hinton and Nowlan already 20 years ago [64]. Since then, several other simulation studies demonstrating the Baldwin effect have been published, see Belew and Mitchel's collection [10], Bruce and Weber's collection [181] and the special issue on the Baldwin effect in the *Evolutionary Computation* journal [175] for a few examples. Several other examples are described in the course of this thesis.

Evolutionary computation allows to break the central dogma by designing a backward machinery that translates phenotypic changes to the genome and thereby allows to investigate Lamarckism.

In fact, there are several examples for coupled evolution and learning, with and without Lamarckism in the evolutionary computation literature. Most of these evolutionary computation algorithms have been developed for the purpose of evolutionary optimization. These algorithms have been named *Memetic Algorithms (MAs)* [115, 58].

The Memetic Algorithm

In the context of evolutionary optimization algorithms, learning is added through a local search method. The entire optimization algorithm is called *Memetic Algorithm (MA)* [115, 58]. A canonical memetic algorithm is described in pseudo code notation in Algorithm 2. In comparison to the canonical evolutionary algorithm, a memetic algorithm is characterized by an extended evaluation scheme. In particular, the final evaluation of an individual is preceded by an individual search in its local neighborhood in the search space. See also the flowchart in Figure 2.7. Memetic algorithms are typically used to solve stationary optimization problems.

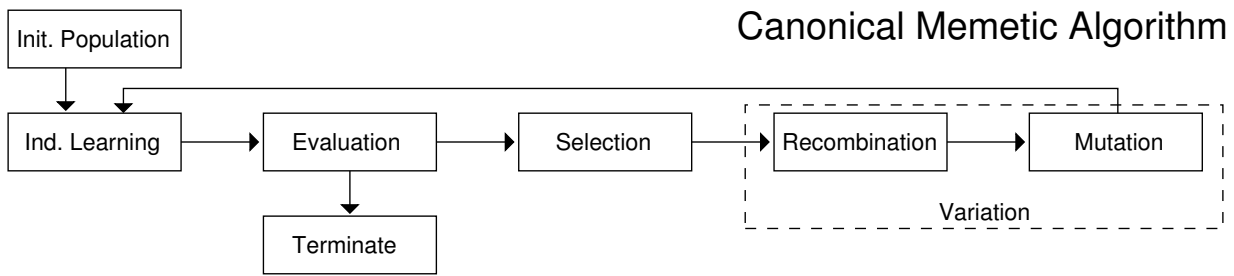


Figure 2.7: Flowchart of a canonical memetic algorithm. In comparison to the canonical evolutionary algorithm, a memetic algorithm is characterized by an extended evaluation scheme. In particular, the final evaluation of an individual is preceded by a local search of the algorithm which aims to improve fitness.

Terminology

In memetic algorithms and related evolutionary computation fields, the terms *Lamarckian learning* and *Darwinian learning* have become standard terminology. Lamarckian learning refers to the case when evolution and learning are coupled and some form of Lamarckism is employed. Darwinian learning refers to the case when evolution and learning is coupled without Lamarckism. Sometimes the term *Baldwinian learning* is used as a synonym for Darwinian learning.

In other words, learning is called *Lamarckian* if the result of an individual’s learning is transferred to its offspring, and it is called *Darwinian* if this is not the case and the result of learning is “thrown away”. However, in both cases learning influences the fitness of individuals. In the absence of Lamarckism, this potentially causes the Baldwin effect. In evolutionary computation, the terms “Lamarckian learning” and “Darwinian learning” (or “Baldwinian learning”) are somewhat misleading because usually the difference between these evolutionary systems does not lie in the learning procedure but solely in the inheritance mechanisms. Therefore, the terms *Lamarckian inheritance* (respectively *Lamarckism*) and *Darwinian inheritance* seem to be more appropriate to distinguish the two cases.

Interestingly, Lamarck’s and Darwin’s findings are often presented as to oppose each other (“Darwinian versus Lamarckian inheritance”), in particular in the realm of evolutionary computation. As briefly outlined earlier, their works had a very different focus, and Darwin’s theory can rather be seen as an extension of Lamarck’s. Therefore, the terms *Lamarckian inheritance* and *biological inheritance* are chosen in this thesis to refer to coupled evolution and learning with and without Lamarckism, respectively.

2.5 Summary and Conclusion

This chapter has presented the fundamentals upon which this thesis is built. Several definitions have been proposed that are valid for the remainder of this thesis. Corresponding to the definitions of this chapter, symbols and domains are specified and a short description of the

the symbol's meanings are provided in the *Nomenclature* on page IX in the preface of this thesis.

In this chapter, the main principles of evolution and learning have been introduced briefly. More examples that further explain these principles are presented in the corresponding related work sections of this thesis, namely in Section 3.1, Section 4.1, Section 7.2, and Section 8.1.

Lamarckian and Biological Inheritance

As reviewed in Section 2.4, Lamarckian inheritance is not biologically plausible, but it can be developed for artificial systems of evolution and learning. Furthermore, Lamarckian-like mechanisms exist in nature (cf. Section 2.4.2). Therefore, one might ask whether there is an advantage of Lamarckism from a purely adaptational point of view. In other words, if an evolutionary system can be endowed with either Lamarckism or biological inheritance, what is the preferred choice?

Ignoring the cost of Lamarckian inheritance, e.g., development and maintenance of a backward-machinery that transfers phenotype information to genotype, this chapter compares Lamarckian and biological inheritance from an adaptational point of view. It turns out that Lamarckism produces an adaptational disadvantage in rapidly changing environments. However, in slowly changing environments a population endowed with Lamarckian inheritance shows a better adaptation behavior than a population without Lamarckian inheritance.

Apart from this chapter, in this thesis biological inheritance is assumed. This chapter also aims to highlight the main arguments for this decision. Large parts of this chapter are based on [132].

This chapter begins with a review of the related work (Section 3.1). Then, the conditions that need to be satisfied in order to observe Lamarckism are formally derived (Section 3.2). A simplified model is suggested and evaluated in Section 3.3 based on a simulation study. The chapter closes with a summary and conclusions (Section 3.4).

3.1 Related Work

The majority of computational studies that couple evolution and learning can be found in the field of memetic algorithms (cf. Section 2.4.3) which are also sometimes called *hybrid genetic algorithms* [115, 58]. Recall that this class of optimization algorithms is typically used to solve stationary optimization problems.

With regard to the question, whether Lamarckism provides an adaptational advantage, indirect evidence may be provided by the fact that in most applications to stationary optimization problems the Lamarckian inheritance mechanism is employed (cf. comment in [58, p.15]).

Unfortunately, only a small fraction of the published work in this research field focuses on a direct comparison of Lamarckian and biological inheritance.

In one of these studies, Gruau and Whitley [51] compare coupled evolution and learning of (artificial) Boolean neural networks with Lamarckian inheritance to the case of biological inheritance. The evolutionary goal is to find a configuration of Boolean neural network that produces a certain target function. It turns out that evolution finds the target function earlier under Lamarckian inheritance than under biological inheritance.

This result is consistent with the findings of Julstrom [78] who compares Lamarckian and biological inheritance for optimization of a modified traveling sales person problem, where the optimization goal is to find a collection of sub-routes with 4 cities that yield a short overall tour. After a fixed number of generations it turns out that the population that employs Lamarckian inheritance has on average found a better solution than the population that employs the biological inheritance mechanism.

Ku and Mak [89] apply Lamarckian and biological inheritance in coupled evolution and learning to a recurrent neural network which should learn a temporal relationship of inputs. In their experiments Lamarckian inheritance is clearly superior over biological inheritance with regard to the rate at which evolution finds a good neural network.

In [148] and [149], Sasaki and Tokoro compare Lamarckian and biological inheritance for stationary as well as dynamic environments in an artificial life framework in which neural network agents have to discriminate food from poison. Their model does not only allow either pure Lamarckian or pure biological inheritance but also allows intermediate levels. The simulation results correspond to Gruau and Whitley [51], i.e., a high degree of Lamarckian inheritance solves the optimization task much better than biological inheritance in a stationary environment. However, in simulations with a dynamic environment, populations with a high degree of biological inheritance show a better adaptation ability over time.

In [142], Rocha and Cortez come to a similar result: Lamarckian inheritance is preferable in stationary settings, while biological inheritance “reveals a greater robustness in dynamic ones” [142, page 382].

Houck et al. [70] find for a range of stationary optimization problems that some form of partial Lamarckism is superior to both pure Lamarckism and biological inheritance .

In another study by Whitley et al. [185] with binary encoded genotype and phenotype applied to a set of common benchmark fitness functions, Lamarckian inheritance finds much quicker a good solution for every tested fitness function. In the long run, however, biological inheritance finds better solutions in some of the fitness functions.

The latter paper suggests that Lamarckian inheritance is preferable with respect to efficiency, and biological inheritance seems to be preferable with respect to effectiveness. This may also explain that in practice most algorithms employ a Lamarckian inheritance mechanism. So, the explanation for this could indeed be the large number of fitness evaluations required by biological inheritance to find a good solution. Unfortunately most of the experiments reviewed here, have a fixed and relatively short runtime, and in the light of the results of

Whitley et al. [185] it would have been interesting to see how the results compare to each other when the evolutionary optimization is run for a very long time.

More consistent across several experiments is the result that Lamarckian inheritance is superior over biological inheritance in stationary environments (such as stationary optimization problems) and on the contrary biological inheritance has an advantage in dynamic environments. In the reviewed papers, several qualitative explanations are provided to explain these results. However, a fine-grained analysis is not presented, possibly because the simulation models are too complex for such an analysis.

In this chapter, a simplified model of evolution and learning is presented which produces similar results as e.g. by Sasaki and Tokoro [149] but which allows a fine-grained analysis leading to a clear understanding of the results. Before this model is presented, the conditions for the observation of Lamarckism are formally derived.

3.2 Conditions for Lamarckism

As mentioned earlier, Lamarckian inheritance can be implemented in artificial evolutionary systems. This requires to design a reverse genotype-phenotype mapping (the “backward machinery”) in addition to the artificial forward genotype-phenotype mapping. In the following, some conditions for the construction of the forward and backward mapping are derived that need to be satisfied in order to observe Lamarckism.

Recall that in the time of Lamarck, no valid theory of heredity had been developed, yet. Thus, Lamarck was not aware of the distinction between genotype and phenotype. This distinction is assumed in the following, because it is crucial for the formal analysis of the forward and backward genotype-phenotype mappings. In particular it is assumed that phenotypic changes of an individual can only be transferred to its offspring after modification of its own genotype. Thus, under Lamarckism, learning does not only change the individual’s phenotype, but also its genotype.

The function ϕ represents the mapping from a genotype x to a phenotype z with respect to the environmental state e , thus $z = \phi(x, e)$. Which phenotype is expressed by genotype x depends on the current environmental state e . The function γ represents the change of a genotype x to a genotype x' under the influence of e , i.e., $x' = \gamma(x, e)$. Since there might be random influences, the expected outcome of the mappings is considered which is denoted as $E(\gamma(x, e))$ and $E(\phi(x, e))$. In the following three cases are distinguished: *No Lamarckism*, *Weak Lamarckism*, and *Pure Lamarckism*.

Definition 3.1 (No Lamarckism). *No Lamarckism is present if the environment e has no influence on the (expected) genotype, i.e.,*

$$\forall (x, e', e'') : E(\gamma(x, e')) = E(\gamma(x, e'')) ,$$

where e' and e'' are environmental states with $e' \neq e''$ and x is the initial genotype.

Definition 3.2 (Weak Lamarckism). *Weak Lamarckism is present if learning or environment has an influence on the (expected) genotype, i.e.,*

$$\exists (x, e', e'') : E(\gamma(x, e')) \neq E(\gamma(x, e''))$$

where e' and e'' are environmental states with $e' \neq e''$ and x is the initial genotype.

Definition 3.3 (Pure Lamarckism). *Pure Lamarckism is present if the genotype-phenotype-mapping ϕ and the change of the genotype γ are related in the following way: Under an arbitrary learning influence e' , genotype x produces the same (expected) phenotype as the resulting genotype would produce in the absence of learning. Denoting, the absence of learning as e_0 pure Lamarckism is formally defined as*

$$\forall (x, e') : E(\phi(x, e')) = E(\phi(\gamma(x, e'), e_0)) .$$

Alternatively, definitions 3.1-3.3 can be formulated with a genotype-phenotype mapping that is influenced by learning parameter a , i.e., $\phi(x, a)$ respectively $\gamma(x, a)$. In this case, e needs to be replaced by a in definitions 3.1-3.3. Learning parameter a can be interpreted as learning time, life time or learning intensity, or it can simply mean the absence of learning (if $a = 0$) respectively the presence of learning ($a = 1$).

The conditions for pure Lamarckism (Definition 3.3) highlight an interesting conceptual difficulty of Lamarckism. In order to observe the inheritance of the parent's learned characteristics it must be avoided that the offspring overwrites these with newly learned characteristics. Thus, this form of Lamarckism can only be observed unequivocally if there is a "neutral" environment. Correspondingly, in the case that the genotype-phenotype mapping is influenced by a learning parameter, pure Lamarckism can only be observed if the offspring's innate phenotype can clearly be identified, e.g., by disabling offspring learning.

In most cases where Lamarckian inheritance is employed in evolutionary computation, the two search mechanisms, evolution and learning, work on one representation. In this case, a *direct transfer* of the phenotype to the genotype after learning is the straightforward backward-machinery which satisfies the conditions for pure Lamarckism. In fact, there is no need to distinguish between genotype and phenotype any longer in this case (cf. Section 2.1.4).

3.3 A Simplified Model

In this section, first the simplified model and its simulation set-up are described (Subsection 3.3.1). The results of the simulation are presented in Subsection 3.3.2 and discussed in Subsection 3.3.3.

3.3.1 Model Description

Inspired by the model of Jablonka et al [74], the simplified model of evolution and learning, allows two environmental states $e \in \{E_0, E_1\}$. Two phenotypes $z \in \{P_0, P_1\}$ are possible, where P_0 is better adapted to E_0 , and P_1 is better adapted to E_1 , i.e.,

$$\begin{aligned} f(P_0|E_0) &> f(P_1|E_0) \quad , \\ f(P_0|E_1) &< f(P_1|E_1) \quad , \end{aligned} \tag{3.1}$$

where f denotes the absolute fitness score. In the simulations of Section 3.3.2 fitness scores are set such that

$$\forall (i, j), i \neq j : \frac{f(P_i|E_i)}{f(P_i|E_j)} = 2 \quad , \tag{3.2}$$

i.e., the fitter phenotype reproduces twice as much as the unfit. This ratio defines the selection pressure. The real-valued genotype $x \in [0, 1]$ represents the predisposition toward phenotypes P_0 and P_1 . A low x value corresponds to a genetic predisposition toward P_0 , and a high x value corresponds to a genetic predisposition toward P_1 . The probability to realize a certain phenotype also depends on a learning parameter $a \in [0, 1]$ (the larger a , the higher the learning rate) and the environmental state $e \in \{E_0, E_1\}$, in particular

$$\begin{aligned} p(z = P_0|x, E_i, a) &= \begin{cases} \phi(1 - x, a), & \text{if } i = 0 \\ 1 - \phi(x, a), & \text{if } i = 1 \end{cases}, \\ p(z = P_1|x, E_i, a) &= \begin{cases} \phi(x, a), & \text{if } i = 0 \\ 1 - \phi(1 - x, a), & \text{if } i = 1 \end{cases}, \end{aligned} \quad (3.3)$$

where

$$\phi(x, a) = \begin{cases} x^{1/(1-a)}, & \text{if } 0 \leq L < 1 \\ 1, & \text{if } L = 1 \end{cases}. \quad (3.4)$$

Equation 3.3 indicates that in both environments, the probability to produce the high-fitness phenotype (P_0 in E_0 , P_1 in E_1) increases with a , i.e., learning is adaptive. Notice that the probability to express phenotype P_0 is always the counter-probability of realizing P_1 . Figure 3.1 illustrates the relationship as formulated in Equation 3.3 for different values of the learning parameters L .

In each generation, each of 100 individuals reproduces asexually an expected number of $w = f/f$ offsprings. f is the individual's absolute fitness, \bar{f} the mean absolute fitness of the population and w is the relative fitness of an individual. This implies a constant population size over time. The selection scheme, known as *linear-fitness-proportional selection*, is implemented by the *stochastic universal sampling* algorithm [6] which implements sampling (with replacement) of n offspring from n parents, where the probability of an individual being sampled is proportional to its absolute fitness.

Lamarckian inheritance is implemented as follows. The offspring's genotype x' depends on the parent's genotype x , its learning-induced increase in the probability of realizing the high-fitness phenotype p , and a Lamarckian parameter λ , in particular

$$x' = \lambda p + (1 - \lambda)x. \quad (3.5)$$

Pure Lamarckism is given for $\lambda = 1$ and no Lamarckism is present for $\lambda = 0$. Figure 3.2 illustrates this implementation of Lamarckian inheritance. Mutation is modeled by adding a random number drawn from Gaussian probability distribution with mean $\mu = 0$ and standard deviation $\sigma = 10^{-4}$, cut off at the genotype space boundaries. Lamarckism and mutation are the two forces that modify the genotype. The mutation strength is chosen rather low in order to emphasize the effect of Lamarckism.

In some of the experiments, the Lamarckian parameter λ and/or the learning parameter a evolves as well. In these cases each individual's genotype is extended by an additional gene that stores its λ and a , respectively. The average time between the two environmental changes is specified by an environmental parameter T . Notice that T is an environment parameter (cf. Nomenclature table, the length of the change interval is denoted T instead of e for the

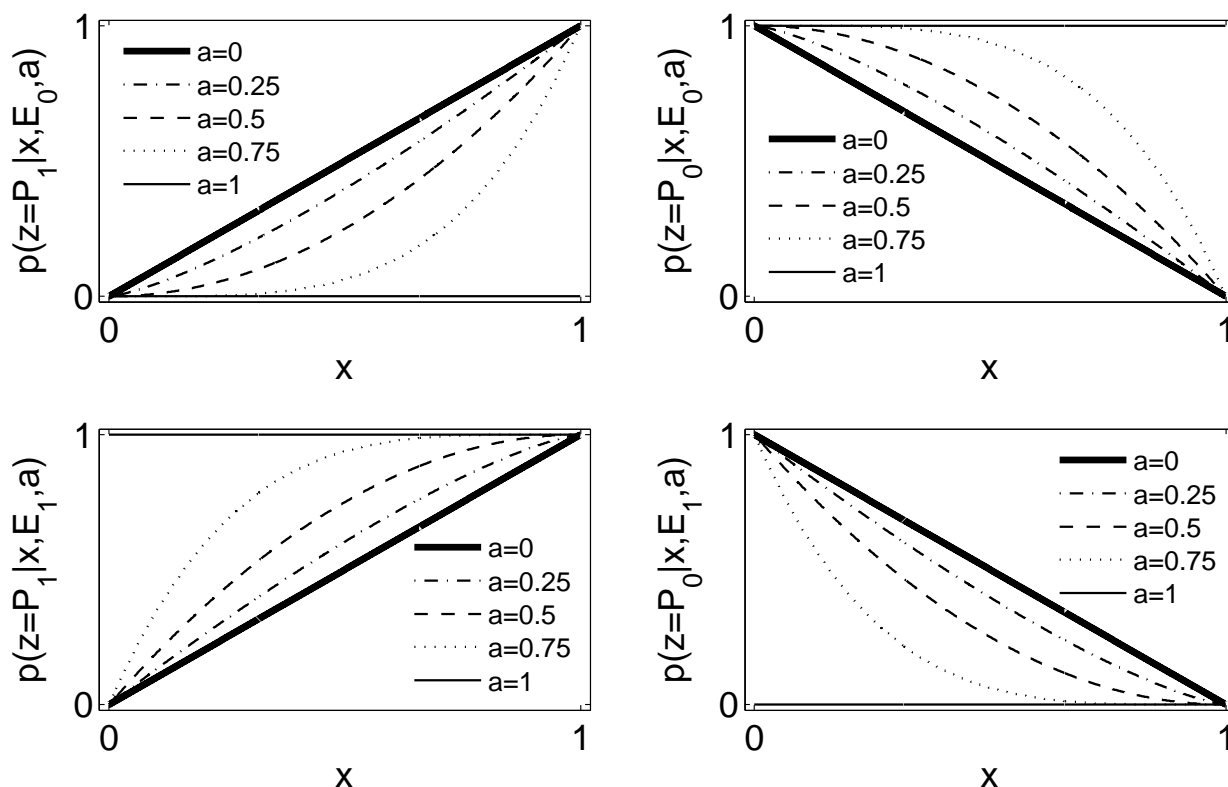


Figure 3.1: Illustration of the probabilistic genotype-phenotype-mapping: Influence of the learning parameter a on the probability to express phenotype P_0 and P_1 for genotype value x , in Environments E_0 and E_1 , as formulated in Equation 3.3. The probability to realize the optimal phenotype (P_0 in E_0 and P_1 in E_1) increases with a .




Figure 3.2: Implementation of Lamarckian inheritance: Learning increases the probability of realizing the optimal phenotype from genetic predisposition x to p (cf. Equation 3.1 and Figure 3.1). Depending on the Lamarckian parameter λ the offspring benefits from this increase directly because it inherits a value x' , with $x \leq x' \leq p$, where λ determines the closeness of x' to x and p .

sake of readability). The actual change periods are either deterministic (cyclic changes) or stochastic. The *population fitness* is defined as follows.

Definition 3.4 (Population fitness). *The population fitness is the average of all absolute fitness values in the population, formally $\frac{1}{n} \sum_{i=1}^n f_i$, where f_i is the realized absolute fitness value of individual i and n is the population size.*

The quality of the adaptation of the population is measured as the population fitness over time. To avoid an initialization bias, only the absolute fitness values from generation 1000 to 2000 are sampled. Three experiments have been carried out which are described in the next section.

3.3.2 Simulation Experiments and Results

In this subsection, three experiments (Experiments 1 to 3) are described and their results discussed.

Experiment 1 - Evolution with Constant Evolutionary Parameters

In this experiment, all evolutionary parameters are held constant during an evolutionary run. Evolution is simulated for Lamarckian parameters $\lambda \in \{0, 0.05, \dots, 0.95, 1.0\}$ combined with environmental change intervals $T \in \{1, 5, 10, \dots, 95, 100, 200\}$. The whole set of parameter combinations is evaluated for constant learning parameters $a = 0.5$ and $a = 0.75$, and for the cases of probabilistic and deterministic environmental changes. The results are shown in Figure 3.3. The figure shows for each combination of T and λ the population fitness, averaged over time and over 25 independent evolution runs. The following findings are qualitatively consistent over all settings: In rapidly changing environments (small T), the population fitness over time is maximal for $\lambda = 0$, i.e., without Lamarckism (see the thick gray line). On the contrary, in slowly changing environments (large T) the population fitness over time (thick gray line) is maximal for $\lambda = 1$, i.e., pure Lamarckism.

The minimum of the population fitness over time (thick black line) is produced with pure or high levels of Lamarckism (λ close to 1) in rapidly changing environments, and without or with low level of Lamarckism (λ close to 0) in slowly changing environments. Interestingly, for intermediate T , the lowest adaptation success is found for intermediate λ . For example, in the top-left panel, for $T = 20$ the minimum population fitness over time is produced with $\lambda = 0.4$. The peculiar (population) fitness valley disappears for very low or high T . A geometric explanation for this fitness valley is given in Appendix A.

In summary: Compared to biological inheritance, Lamarckism produces a better adaptation behavior in slowly changing environments and worse behavior in rapidly changing environments. For an intermediate rate of environmental change, the worst adaptation behavior is produced by an intermediate degree of Lamarckism. The slower the environment changes, the lower is the degree of Lamarckism that produces the worst overall adaptation behavior. The same set of experiments has been carried out for higher mutation probabilities. The results of these experiments are qualitatively consistent with the effects described in this paragraph, even though quantitatively weaker than with low mutation probabilities. Although qualitatively consistent, the observed effects are weaker with higher mutation rates (not shown).

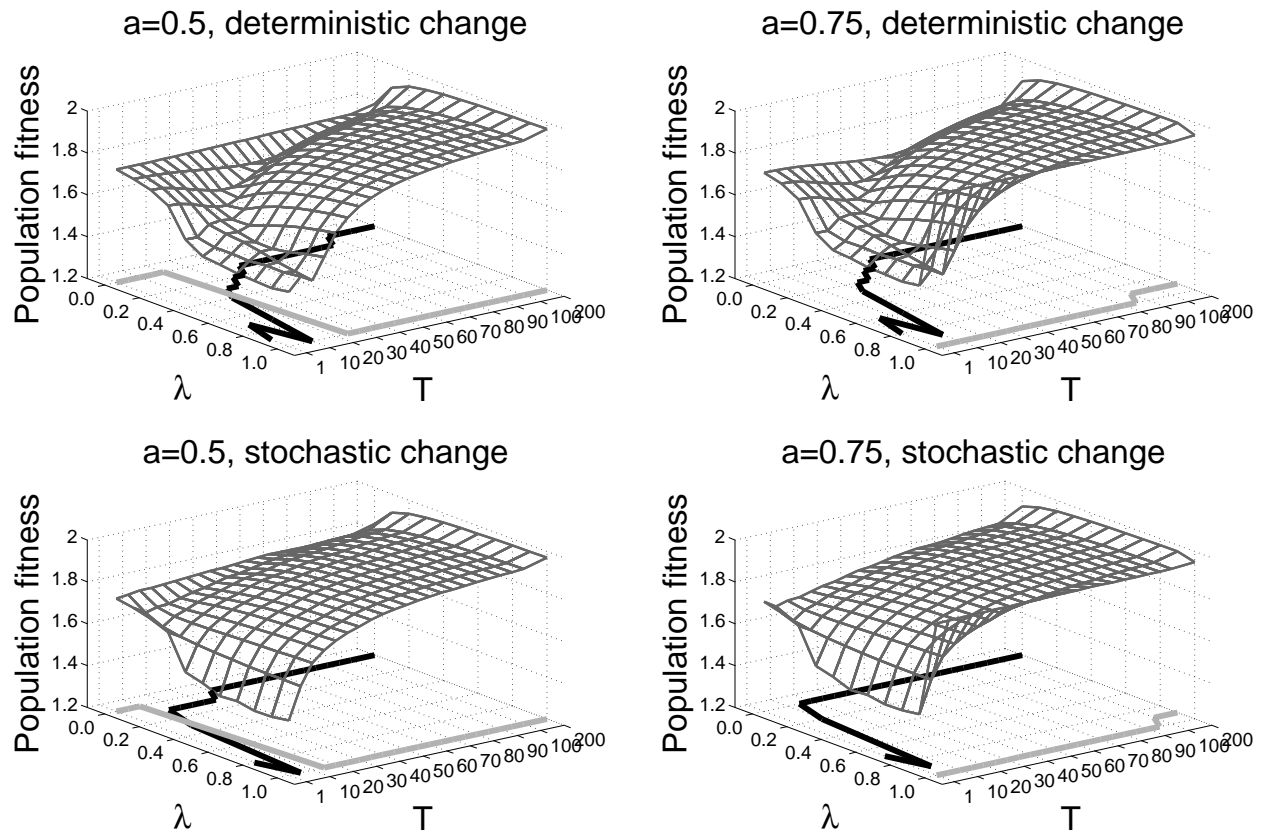


Figure 3.3: *Experiment 1: Measuring population fitness over evolution time for different Lamarckian parameters λ and environmental change intervals T . Environmental changes are deterministic or stochastic, combined with a learning parameter a that is set to either 0.5 or 0.75. A high population fitness over time indicates good, a low population fitness indicates poor population adaptation. The thick gray line shows for which λ the population fitness over time is maximal for a given T . The thick black line shows the corresponding minimum. For rapidly changing environment (small T), the worst adaptation is given for a low degree of Lamarckism (λ). On the contrary, for slowly changing environments (large T), a high degree of Lamarckism is worst. Interestingly, for intermediate T an intermediate λ produces the worst adaptation.*

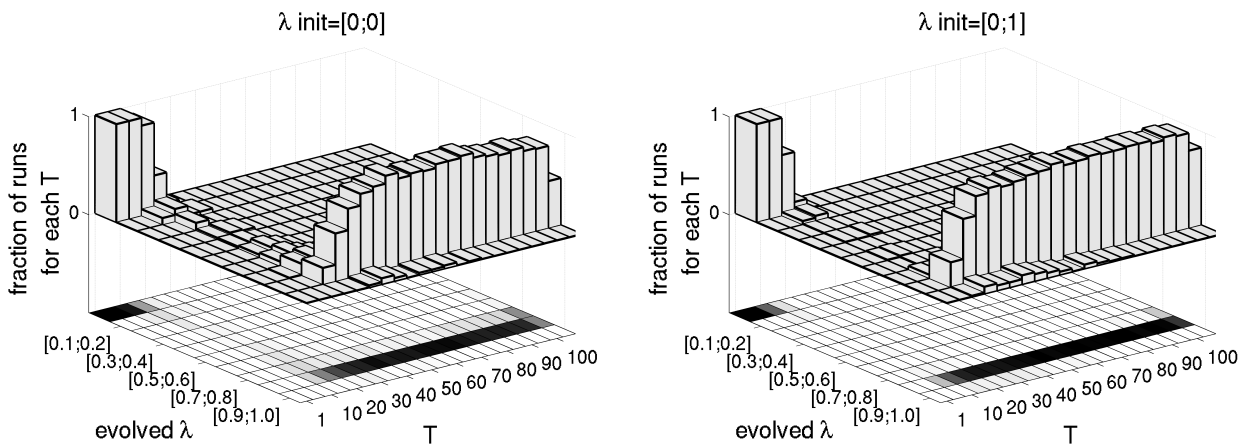


Figure 3.4: *Experiment 2: Evolving the Lamarckian parameter λ , initialized uniformly on $[0, 1]$ (left panel), and starting without Lamarckism, i.e., $\lambda = 0$ for all individuals (right panel), in case of deterministic environmental changes and with $a = 0.5$. Bar-heights indicate the relative number of evolutionary runs that evolved a λ in the corresponding interval. A near-optimal λ according to Figure 3.3 evolves in most evolutionary runs but the initialization of λ has a visible influence.*

Experiment 2 - Evolution of Lamarckism

This experiment aims to test whether the optimal level of Lamarckism λ (cf. thick gray line in the top-left panel of Figure 3.3) evolves when each individual has its λ encoded in the genotype. Note that a second-order adaptation process is necessary for this. Figure 3.4 presents the results of a set of evolutionary runs. For each $T \in \{1, 5, 10, \dots, 95, 100, 200\}$, evolution is run 100 times with learning parameter $a = 0.5$. The mutation strength is set to $\sigma = 10^{-4}$ for x (as in Experiment 1) and to the same value for mutation of λ . The length of the bars in Figure 3.4 represents the fraction of the runs that resulted in a mean λ in the interval $[0, 0.1], [0.1, 0.2] \dots [0.9, 1.0]$. The left panel of Figure 3.4 shows the case in which the initial population was distributed uniformly on the entire λ -range.

In rapidly changing environments ($T \leq 10$), the majority of the runs produces a small λ , and for slower changing environment ($T \geq 15$) a large λ . Comparing this to the results of Experiment 1 (top-left panel of Figure 3.3), we see that the optimal λ indeed evolves in a second-order process. In another experiment (Figure 3.4, right panel) evolution starts without Lamarckism ($\lambda = 0$) for all individuals. In this case, a large λ is only evolved for $T \geq 25$. The likely reason for this difference is the observed fitness valley for intermediate λ in case of intermediate levels of environmental change. Apparently, the population can not cross the fitness minimum for T around 20. In an additional experiment (results not shown) the learning rate a is evolved as well. In the absence of learning cost, a high a quickly evolved and suppressed the evolution of the Lamarckian parameter λ in slowly changing environments: With very high learning ability, there is only weak selection pressure for a large a in slowly changing environments, which leads to the evolution of only intermediate levels of λ .

In summary, in most cases, a near-optimal level of Lamarckism evolves in a second order process. However, in cases where there is a population fitness minimum for intermediate

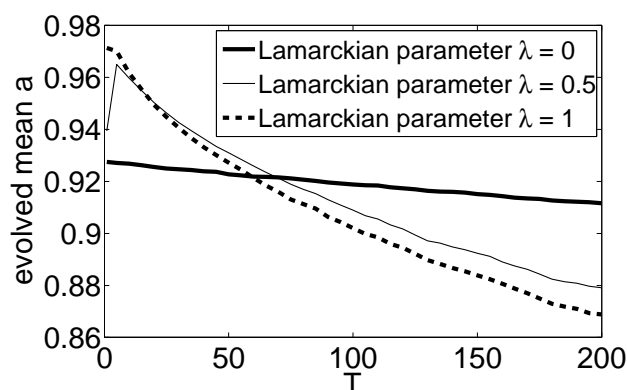


Figure 3.5: *Experiment 3: Evolving the learning parameter a while the level of Lamarckism λ is constant. The figure shows the evolved mean a , in the cases pure Lamarckism ($\lambda = 1$), intermediate level of Lamarckism ($\lambda = 0.5$) and no Lamarckism ($\lambda = 0$).*

levels of Lamarckism (see Experiment 1), the globally optimal level of Lamarckism does not always evolve.

Experiment 3 - Evolution of Learning Ability

The aim of this experiment is to test if Lamarckism influences the evolution of learning ability a . Holding the level of Lamarckism λ and the environmental change interval constant during the evolution, a is evolved. In particular the cases “no Lamarckism” ($\lambda = 0$), pure Lamarckism ($\lambda = 1$) and an intermediate level of Lamarckism ($\lambda = 0.5$) are investigated. The simulation results are presented in Figure 3.5.

Comparing the two extreme cases no ($\lambda = 0$) and pure ($\lambda = 1$) Lamarckism, we see that in quickly changing environments ($T < 60$) a larger mean a evolves with pure Lamarckism, and in slower changing environments a lower mean a evolves with pure Lamarckism than without Lamarckism. The case of intermediate level of Lamarckism ($\lambda = 0.5$) lies between the two extremes cases, but is closer to the case of $\lambda = 1$. So, Lamarckism suppresses the evolution of learning ability in slowly changing environments and facilitates the evolution of learning ability in quickly changing environments. An explanation for this is that for large T , there is a relatively low selection pressure for high a in case of Lamarckism, because a high λ alone allows good adaptation. For small T , however, it has been shown that Lamarckism is detrimental, and there is a relatively high selection pressure to evolve a high a that can compensate for the Lamarckian disadvantage.

In summary, when Lamarckism provides an adaptational advantage (slowly changing environments) a lower learning ability is evolved because there is less selection pressure for it, but when Lamarckism provides an adaptational disadvantage (rapidly changing environments) a higher learning ability is evolved because there is stronger selection pressure for it, i.e., learning compensates the disadvantage of Lamarckism here.

3.3.3 Discussion

The results of the experiments suggest that Lamarckian inheritance has an adaptational disadvantage in rapidly oscillating environments, compared to stationary environments. This disadvantage in rapidly changing environments is explained by the movement of the mean genotype. With Lamarckian inheritance, genotype movement is faster than with genetic mutation alone. In rapidly oscillating environments, Lamarckism increases the integral of genotype distance from the optimum. The advantage of Lamarckian inheritance in slowly changing environments is because the genotype converges to the optimum more rapidly than by random mutation alone. A peculiar finding at intermediate levels of environmental oscillation is that a minimum value of population fitness is associated with a particular value of Lamarckian inheritance. This is in contrast to the monotonic changes in population fitness observed at the very high and very low rates of environmental change.

The near-optimal degree of Lamarckism with respect to the rate of environmental change can be produced by evolutionary self-adaptation. However, the afore mentioned fitness valley may prevent the evolution of Lamarckism from scratch even though high levels of Lamarckian inheritance are a global optimum.

A follow-up experiment in which learning rate is evolvable, showed that the introduction of Lamarckian inheritance in rapidly oscillating environments increases selective pressure for better learning mechanisms, whilst introduction of Lamarckian inheritance in slowly oscillating environments decreases the selective pressure for learning mechanisms.

Note that these findings are limited to instances where environmental changes occur cyclically such that the genotype is able to establish itself in an area where a high fitness under several environmental conditions can be attained through learning. In nature, simple binary oscillating environments involve geophysical rhythms such as diurnal and seasonal cycles. If however, the environment rapidly changes in a non-oscillating path, Lamarckism may be beneficial even in this rapidly changing environments. On the other hand, if there are slow global environmental trends with superimposed rapid cyclic changes, then the conclusions of this chapter are likely to hold as well.

3.4 Summary and Conclusion

The aim of this chapter was to review and study some arguments for the decision to investigate biological (non-Lamarckian, Darwinian) inheritance in the remainder of this thesis.

First of all, Lamarckian inheritance is biologically implausible. For the sake of its interdisciplinary character, it is therefore more appropriate to concentrate on biological inheritance in this thesis.

Furthermore, the construction of Lamarckian inheritance is only straight-forward if both genotype and phenotype have the same representation. Lamarckian inheritance would limit our focus to evolutionary systems where a backward-machinery for the translation of phenotype information to the genotype is possible. Turney [174] points out that apart from the trivial case without genotype-phenotype distinction the inverse mapping incurs high computational cost which may quickly become intractable.

Recall, also the conceptual difficulty of observing Lamarckian inheritance that was pointed out earlier.

Finally, this chapter has shown from a purely adaptational point of view that the optimal degree of Lamarckian inheritance depends on the rate of environmental change, in particular that Lamarckian inheritance can actually produce a disadvantage in dynamic environments.

It should be noted here that the rate of environmental change is only one of the factors influencing the optimal degree of Lamarckian inheritance. However, the complexity and diversity of the environmental dynamics are other factors that are likely to have a strong influence on the optimal degree of Lamarckian inheritance. These issues should be considered in future research.

Apart from this chapter, biological inheritance is assumed in this thesis.

Influence of Learning on Evolution - The Gain Function Framework

The influence of learning on evolution has been intensively studied in the last decades. Contributions from both perspectives biology and evolutionary computation highlight various aspects on how learning influences evolution. In this chapter, it is investigated how learning influences the rate (or velocity) of evolution. As the study of related work will show, there is no consensus on whether learning accelerates or decelerates evolution.

In this chapter, a general mathematical framework for the analysis of the influence of learning on the rate of evolution, the *gain function*, is proposed. This framework defines the conditions under which learning accelerates respectively decelerates evolution. Section 4.3 is based on [131], Section 4.4 is based on [130]. As mentioned earlier biological inheritance is assumed.

4.1 Related Work

A rich body of literature is devoted to the influence of learning on evolution. Here, selected representative examples are reviewed and categorized.

Baldwin's original paper [7] has been introduced in Section 2.4.2 concluding that learning accelerates evolution. The probably most influential paper on the interaction between evolution and learning in the field of computational intelligence has been published by Hinton and Nowlan [64] in 1987. There, they present a simulation study with a population of sexually reproducing individuals in an environment where a single phenotype produces a high fitness. In this rather simplistic model, Hinton and Nowlan demonstrate the Baldwin effect (cf. Definitions 2.10 and 2.11) and show how learning can accelerate evolution. Their model is revisited and described in detail in Section 6.1.

After reviewing some statistical properties of Hinton and Nowlan's model, Maynard Smith [105] points out that with a population size much larger than the one of the original paper (1000), and with asexual reproduction instead of sexual reproduction, evolution would quickly find the optimum even in the absence of learning. Belew [9] revisits Hinton and Nowlan's

model with an analytical treatment of the evolutionary dynamics that confirm the original simulation results. In addition, Belew includes “culture”, a type of learning in which weak individuals learn directly from strong individuals. Belew shows that “cultural learning” leads to an even stronger emphasized Baldwin acceleration effect. Fontanari and Meir [44] also revisit the model of Hinton and Nowlan. They employ a quantitative genetics framework with infinite population size. With a dynamical system analysis Fontanari and Meir confirm Hinton and Nowlan’s conclusion that learning speeds up evolution.

Behera and Nanjundiah [8] extend the model of Hinton and Nowlan [64] by a gene-regulation mechanism which provides phenotypic plasticity. In their simulations, it turns out that learning or phenotypic plasticity accelerates evolution. Note that Behera and Nanjundiah’s aim was to replicate and further understand the results of a famous biological study with fruit flies by Waddington [177, 178].

In the artificial evolution of neural networks coupled with supervised learning for pattern recognition, Kessing and Stork [80] show that evolution is only accelerated through learning in the case of an intermediate degree of learning. With too much or too little individual learning, evolution is decelerated. French and Messinger [46] who study the evolution and learning in simple interacting agents in a 2-d grid-world come to similar conclusions as Kessing and Stork [80], namely that the strongest acceleration effect can be observed for an intermediate degree of individual learning. Despite this, French and Messinger conclude that the “reproduction mode” (sexual or asexual) play an important role in the influence of learning on evolution.

In his biology review, Gordon [49] mentions that learning can decelerate evolution. In support of this claim, Papaj’s [133] simulated evolution of insect learning shows that individual learning slows down the evolution of genetic configurations with high fitness. Papaj’s model is revisited and described in detail in Section 6.2. Further evidence of a deceleration effect of learning is presented by Mayley [103]. Mayley’s simulation study with Kauffman’s NK fitness landscapes [79] shows that learning may work to hide genetic differences between individuals and thereby decelerate evolution. The main factors that constitute this “Hiding effect” [103] are identified as the cost of learning and the degree of interaction between genes (epistasis). Mayley also mentions that a similar effect has already been described in [76] and [49]. In another simulation study with a similar set-up, Mayley [102] presents examples for both learning-induced acceleration and deceleration. Here, the cost of learning play an important role as well. However, another factor that influences the impact of learning is the correlation between genotype and phenotype neighborhood. A high correlation is given when a small distance between individuals in genotype space should corresponds to a small distance in phenotype space. This condition is similar to the concept of causality as reviewed in Section 2.1.4. A number of studies in the field of evolutionary robotics [119, 180] have elaborated on the importance of correlation with differing results [135, 120, 60, 118].

Bull [16] studies the coupling of evolution and a simple trial-and-error learning mechanism on NK fitness landscapes. Contrary to Mayley [103], Bull concludes that learning accelerates evolution. Bull identifies the rate of learning as a crucial parameter that impacts the influence of learning on evolutionary change.

Ku et al. [90] investigate the influence of learning on evolution for the optimization of recurrent neural networks [99]. In their study, they combine a cellular genetic algorithm [184] with different local hill climbing methods for the optimization of the synapse weights of the recurrent neural network. The optimization runs show that learning decelerates evolution.

Despite the deceleration effect of learning in case of biological inheritance, they also show that evolution is accelerated with Lamarckian inheritance. The latter result is in support of the finding of Chapter 3. Accounting for the computational cost of learning the deceleration effect is even more evident. The same authors confirm these results in a similar study in [91].

Ancel [3] argues that phenotypic plasticity does not universally accelerate evolution. She provides an example of a Gaussian fitness function in which the addition of a noise component in the mapping from genotype to phenotype decelerates evolution. This example and related work [18, 5] are revisited and analyzed in detail in Section 6.3. Noteworthy, all the three papers [18, 5, 3] employ a quantitative genetics approach and are therefore among the few examples that provide a mathematical analysis on the influence of learning on evolution.

Dopazo et al. [30] study an extended version of Hinton and Nowlan’s model in which the fitness landscape is relatively smooth in comparison to the original model in [64]. The simulation results suggest that the greater the amount of learning, the stronger is evolution of genetically strong individuals decelerated. Dopazo et al. employ both analytical tools and simulations.

The first biological evolutionary experiment intended to demonstrate the Baldwin effect is proposed by Mery and Kawecki [107]. In an *in vitro* experiment, they study the effect of a simple form of learning on the evolution of resource (food) preference in fruit flies. In one experimental set-up (evolution of resource preference A), learning accelerates the evolution of the genetic predisposition toward the target preference (A). However, in another set-up (evolution of resource preference B) learning decelerates the evolution of the genetic predisposition toward the target preference (B). Mery and Kawecki’s study is revisited in Section 6.4 in detail.

Borenstein et al. [13] develop a mathematical model to explain the influence of learning on a population’s ability to cross valleys in fitness landscape. In particular, a stochastic model approach is employed that aggregates the population movement to one moving point of the fitness landscape. Although the meaning of fitness in the fitness landscape model of Borenstein et al. is not explicitly specified, it is probably used in the sense of absolute fitness, cf. Definition 2.5. This model is reviewed in more detail in Section 6.5.3. It turns out that the degree to which learning accelerates evolution is positively correlated with the learning-induced reduction in the so called “drawdown” of the fitness landscape (cf. Section 6.5.3). The drawdown reduction in turn is influenced by the type and the amount of learning.

Similar conclusions have been derived in the study of Mills and Watson [109]. They confirm that learning accelerates evolution by easing the crossing of a fitness valley.

The above review is (with necessary simplifications) summarized in Table 4.1. Indeed, no general conclusion on whether learning accelerates or decelerates evolution can be drawn. The analysis framework developed in the remainder of this chapter attempts to derive general conditions for learning-induced acceleration and deceleration of evolution. The framework allows to predict the results of several (but not all) models that have been reviewed in this section.

Table 4.1: *Related work overview: Conclusions from studies on the influence of learning on evolution. The conclusions of the papers are assigned with respect to the following categories (with necessary simplifications). Has learning-induced acceleration or deceleration of evolution (in terms of the rate of evolution) or both been observed? Is the aim of the study evolutionary computation or biology oriented, or both? What was the analysis approach - simulation, mathematical analysis, in vitro experiment or theoretical considerations? Have there any particular factors been identified that determine the influence of learning on evolution - population size, recombination model (sexual or asexual), epistasis (degree of interaction between genes in the genotype-phenotype mapping), cost of learning (cf. Section 2.2.3), G-P-correlation (correlation between genotype and phenotype space)? Some papers derive factors that are less general. Such factors are omitted here. The papers marked with a “*” (star) are revisited in detail in Chapter 6.*

	Acceleration	Deceleration	Evolutionary Computation	Biology	Simulation	Mathematical	In Vitro Experiment	Theoretical Considerations	Population Size	Recombination Mode	Epistasis	Amount of Learning	Cost of Learning	G-P-Correlation
Baldwin (1896) [7]	×			×				×						
Johnston (1982) [76]	×	×		×				×						
Hinton & Nowlan (1987)* [64]	×		×		×									
Maynard Smith (1987) [105]	×			×				×	×	×				
Belew (1990) [9]	×		×			×								
Fontanari & Meir (1990) [44]	×		×	×		×				×				
Keesing & Stork (1991) [80]	×	×	×		×							×		
Gordon (1992) [49]		×						×						
French & Messinger (1994) [46]	×	×	×		×					×		×		
Papaj (1994)* [133]		×		×	×									
Andersson (1995)* [5]		×		×		×								
Mayley (1996) [102]	×	×	×		×								×	×
Mayley (1997) [103]		×	×		×						×		×	
Bull (1999) [16]	×		×		×							×		
Ku et al. (1999) [90]		×	×		×							×		
Ku et al. (2003) [91]		×	×		×							×		
Ancel (2000)* [3]		×		×		×								
Dopazo et al. (2001) [30]		×	×		×	×						×		
Behera & Nanjundiah (2004) [8]	×			×	×									
Mery and Kawecky (2004)* [107]	×	×		×			×							
Borenstein et al (2006)* [13]	×		×		×	×						×		
Mills & Watson (2006) [109]	×		×		×									

4.2 Basic Idea

As outlined in Chapter 2 development (ontogenesis) and learning (epigenesis) are treated in a sequential fashion. It is assumed that genotypic information alone is sufficient to produce an innate phenotype (development) that can be assigned a fitness. The innate phenotype is modified through learning resulting in the learned phenotype, cf. Figure 4.1.

The rate of evolutionary (genotypic) change increases with the relative differences in fitness among different individuals. Learning influences the fitness (cf. the Baldwin effect) of individuals that have a certain genetic pre-disposition and may thereby influence fitness differences between individuals with “strong” and “weak” genetic pre-dispositions. “Weak” and “strong” genetic predisposition refer to a low, respectively high expected fitness of a genotype. Learning may, for example, amplify relative fitness differences between individuals with “strong” and “weak” genetic pre-dispositions. In this case, genetically strong individuals benefit more from learning than their genetically weak rivals, and evolution is accelerated. The opposite case may occur as well. Learning may reduce relative fitness differences between individuals with “strong” and “weak” genetic pre-dispositions.

Figure 4.2 visualizes the claim. The figure shows the mapping from genotype to fitness for two cases of learning. Recall that the fitness f of a genotype x is given by $f(\phi(x))$ where $\phi(x)$ is the mapping from genotype to phenotype. The dashed curve shows the innate fitness for a given genotype, which is assumed to be linearly increasing with the genotype value. A “weak” individual with genotype $x = 0.5$ has a fitness of $f(\phi(0.5)) = 0.5$ and a “strong” individual with genotype $x = 0.75$ has a fitness of $f(\phi(0.75)) = 0.75$. In relative terms, the strong individual’s fitness is $0.75/0.5 = 1.5$ times the weak individual’s fitness. Thus, it is expected that the strong individual produces 50 percent more offspring than the weak individual.

Learning influences this ratio. In case 1, where the learning-induced change in the mapping from genotype to expected fitness results in the gray curve, the strong individual’s fitness is $1.07/0.56 \approx 1.9$ times the weak individual’s fitness. Hence, now it is expected that the strong individual produces 90 percent more offspring than the weak one.

In case 2, where the learning-induced change in the mapping from genotype to expected fitness results in the solid black curve, the ratio of strong to weak individual’s fitness becomes $1.75/1.44 \approx 1.2$, i.e., the strong individual is expected to produce only 20 percent more offspring than the weak individual.

As a consequence, in case 1 learning accelerates genetic evolution toward a high fitness region and in case 2 learning decelerates genetic evolution toward a high fitness region, compared to the case where no learning is present. Formally, $f(\phi(x))$ denotes the innate fitness of a genotype x and $f(l(\phi(x)))$ the fitness after learning. For convenience $f(\phi(x))$ is substituted by $f_\phi(x)$, and $f(l(\phi(x)))$ is substituted by $f_{\phi_l}(x)$.

Generally, learning-induced acceleration of evolution is expected if

$$\frac{f_{\phi_l}(x_{\text{strong}})}{f_{\phi_l}(x_{\text{weak}})} > \frac{f_\phi(x_{\text{strong}})}{f_\phi(x_{\text{weak}})}. \quad (4.1)$$

Rewriting Equation 4.1,

$$\frac{f_{\phi_l}(x_{\text{strong}})}{f_\phi(x_{\text{strong}})} > \frac{f_{\phi_l}(x_{\text{weak}})}{f_\phi(x_{\text{weak}})}, \quad (4.2)$$

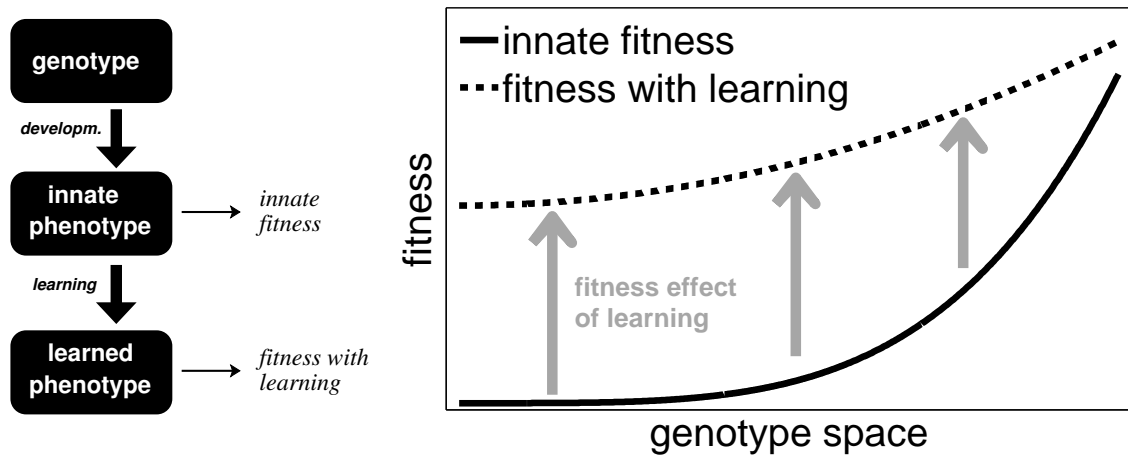


Figure 4.1: The basic model to analyze the influence of learning on evolution. By changing the phenotype (left), learning also changes the mapping from genotype to fitness (right).

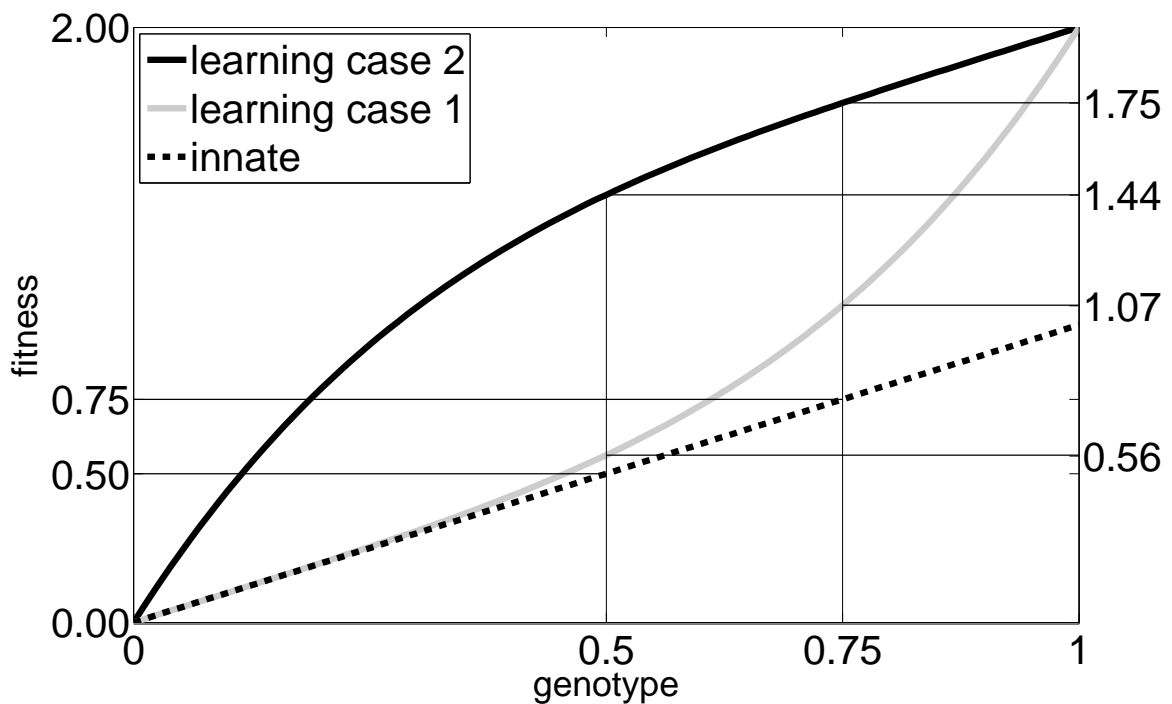


Figure 4.2: Illustration of the basic idea of the gain function. In the absence of learning the fitness ratio of a strong individual ($x = 0.75$) and a weak individual ($x = 0.5$) is $0.75/0.5 = 1.5$. After learning the ratio is $1.07/0.56 \approx 1.9$ in case 1 and $1.75/1.44 \approx 1.2$ in case 2. Hence, in case 1, genetically strong individuals reproduce more frequently in the presence than in the absence of learning, and in case 2, genetically strong individuals reproduce less frequently in the presence than in the absence of learning.

reveals the basic idea of the gain function. If the relative fitness gain of learning, $f_{\phi_l}(x)/f_{\phi}(x)$, increases toward higher fitness, learning is predicted to accelerate evolution.

In the following, a mathematical framework, called the *gain function*, is developed.

4.3 The Gain Function Framework

The gain function framework builds upon the above presented idea that the increase respectively the decrease in relative fitness gain toward a higher fitness region determines the influence of learning on the rate of evolution.

4.3.1 Formulation

In the following, an individual is characterized by a real-valued genotypic variable x and a real-valued phenotypic variable z and the mapping from genotype to innate phenotype (development) is

$$z = \phi(x) . \quad (4.3)$$

An individual changes its innate phenotype via a learning function $l(z)$. This means that a genotype x produces phenotype $\phi(x)$ in the absence of learning, and phenotype $l(\phi(x))$ in case of learning. The absolute fitness of an individual is assigned using a fitness function $f(z)$, defined on the phenotype space. Thus, fitness is given by $f(l(\phi(x)))$ in case of learning and by $f(\phi(x))$ in the absence of learning. As mentioned earlier, $f(\phi(x))$ is denoted as $f_{\phi}(x)$, and $f(l(\phi(x)))$ is denoted as $f_{\phi_l}(x)$. When $l(x)$ is a stochastic function, $f_{\phi_l}(x)$ needs to be replaced by the expected fitness of the learned phenotype, denoted $\bar{f}_{\phi_l}(x)$. It is assumed that the fitness function $f_{\phi}(x)$, respectively $f_{\phi_l}(x)$, is positive and monotonic within the range of population variability.

We now consider a finite population of n individuals, whose genotype values are labeled x_i , $i = 1 \dots n$. The rate of evolution is measured as the distance that the population's mean genotype

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad (4.4)$$

moves toward the optimum in one generation. An individual's reproduction probability is assumed to be proportional to its absolute fitness value f . With regard to the biological concept of *fitness*, where fitness corresponds to the number of offspring produced by an individual, this is the most reasonable selection model. Notice that in the field of evolutionary computation, this selection method is known as *fitness proportional selection*, cf. Section 2.1.3. With this assumption, the expected mean genotype after selection \bar{x}^* can be calculated as

$$\bar{x}^* = \frac{\sum_{i=1}^n x_i f_{\phi}(x_i)}{\sum_{i=1}^n f_{\phi}(x_i)} . \quad (4.5)$$

Assuming an unbiased, symmetric mutation this is equal to the mean genotype of the next generation. The expected change of the mean genotype, S_x in one generation is given by

$$S_x = \frac{\sum_{i=1}^n x_i f_{\phi}(x_i)}{\sum_{i=1}^n f_{\phi}(x_i)} - \frac{1}{n} \sum_{i=1}^n x_i . \quad (4.6)$$

In quantitative biology, S_x is also known as *selection differential*. The mean genotype change in case of learning S_{xl} is derived analogously by replacing f_ϕ with f_{ϕ_l} in Equation 4.6. Thus, learning accelerates (decelerates) evolution if

$$\text{sign}(S_{xl} - S_x) = \text{sign} \left(\frac{\sum_{i=1}^n x_i f_{\phi_l}(x_i)}{\sum_{i=1}^n f_{\phi_l}(x_i)} - \frac{\sum_{i=1}^n x_i f_\phi(x_i)}{\sum_{i=1}^n f_\phi(x_i)} \right) \quad (4.7)$$

is positive (negative). The *gain function* is now defined as the quotient between the genotype-to-fitness function with learning and the genotype-to-fitness function without learning, i.e.,

$$\boxed{g(x) = \frac{f_{\phi_l}(x)}{f_\phi(x)}}. \quad (4.8)$$

Under the assumption that g is monotonic over the range of population variation, it is shown that

$$g'(x) \begin{cases} > 0 \Leftrightarrow S_{xl} - S_x > 0 \text{ (case A)} \\ < 0 \Leftrightarrow S_{xl} - S_x < 0 \text{ (case B)} \\ = 0 \Leftrightarrow S_{xl} - S_x = 0 \text{ (case C)}. \end{cases} \quad (4.9)$$

Equation 4.9 shows that whether learning accelerates or decelerates evolution is determined by the sign of the derivative of the gain function. A positive derivative implies acceleration, a negative implies deceleration and a constant gain function implies that learning has no effect on evolution. Conversely, if we find that learning has accelerated (decelerated) evolution we know that the gain function derivative is positive (negative), under the assumptions given above.

4.3.2 Proof

Given that there is genetic variation, f_{ϕ_l} and f_ϕ are increasing in x and that the sign of $g'(x)$ is constant within the range present in the population ($[x_{\min} \leq x \leq x_{\max}]$) Equation 4.9 is proved by induction. In the following, the proof for case A of Equation 4.9 ($g'(x) > 0$) is outlined. The other cases are omitted because the respective proofs are analogous and the transfer from the first case is straightforward. Recalling Equation 4.5, Statement **S(n)** is defined as

$$\mathbf{S(n)} := \frac{\sum_{i=1}^n x_i f_{\phi_l}(x_i)}{\sum_{i=1}^n f_{\phi_l}(x_i)} - \frac{\sum_{i=1}^n x_i f_\phi(x_i)}{\sum_{i=1}^n f_\phi(x_i)} = \bar{x}_l^* - \bar{x}^* = S_{xl} - S_x > 0. \quad (4.10)$$

Recalling the gain function definition $g(x) = f_{\phi_l}(x)/f_\phi(x)$, we obtain

$$\forall x, x_i, x_j \in [x_{\min}, x_{\max}], x_i < x_j : g'(x) > 0 \Leftrightarrow \frac{f_{\phi_l}(x_i)}{f_\phi(x_i)} < \frac{f_{\phi_l}(x_j)}{f_\phi(x_j)}. \quad (4.11)$$

Without loss of generality it is assumed that the x_i are arranged in ascending order, i.e.,

$$\forall (i, j) : i < j \Rightarrow x_i \leq x_j, . \quad (4.12)$$

Initialization: For $n = 2$, $\mathbf{S}(\mathbf{n})$ can be written and reformulated

$$\mathbf{S}(2) \tag{4.13a}$$

$$\Leftrightarrow \frac{x_1 f_{\phi_l}(x_1) + x_2 f_{\phi_l}(x_2)}{f_{\phi_l}(x_1) + f_{\phi_l}(x_2)} > \frac{x_1 f_{\phi}(x_1) + x_2 f_{\phi}(x_2)}{f_{\phi}(x_1) + f_{\phi}(x_2)} \tag{4.13b}$$

$$\Leftrightarrow \frac{x_1(f_{\phi_l}(x_1) + f_{\phi_l}(x_2)) + (x_2 - x_1)f_{\phi_l}(x_2)}{f_{\phi_l}(x_1) + f_{\phi_l}(x_2)} > \tag{4.13c}$$

$$\frac{x_1(f_{\phi}(x_1) + f_{\phi}(x_2)) + (x_2 - x_1)f_{\phi}(x_2)}{f_{\phi}(x_1) + f_{\phi}(x_2)}$$

$$\Leftrightarrow x_1 + \frac{(x_2 - x_1)f_{\phi_l}(x_2)}{f_{\phi_l}(x_1) + f_{\phi_l}(x_2)} > x_1 + \frac{(x_2 - x_1)f_{\phi}(x_2)}{f_{\phi}(x_1) + f_{\phi}(x_2)} \tag{4.13d}$$

$$\Leftrightarrow \frac{f_{\phi_l}(x_2)}{f_{\phi_l}(x_1) + f_{\phi_l}(x_2)} > \frac{f_{\phi}(x_2)}{f_{\phi}(x_1) + f_{\phi}(x_2)} \tag{4.13e}$$

$$\Leftrightarrow \frac{f_{\phi_l}(x_1)}{f_{\phi_l}(x_2)} + 1 < \frac{f_{\phi}(x_1)}{f_{\phi}(x_2)} + 1 \tag{4.13f}$$

$$\Leftrightarrow \frac{f_{\phi_l}(x_1)}{f_{\phi}(x_1)} < \frac{f_{\phi_l}(x_2)}{f_{\phi}(x_2)} \tag{4.13g}$$

$$\Leftrightarrow g(x_1) < g(x_2), \tag{4.13h}$$

which is true according to Equation 4.11.

Inductive step: Assuming $\mathbf{S}(\mathbf{n})$ is true, it is shown that $\mathbf{S}(\mathbf{n} + 1)$ is true:

$$\mathbf{S}(\mathbf{n} + 1) \tag{4.14a}$$

$$\Leftrightarrow \frac{\sum_{i=1}^{n+1} x_i f_{\phi_l}(x_i)}{\sum_{i=1}^{n+1} f_{\phi_l}(x_i)} - \frac{\sum_{i=1}^{n+1} x_i f_{\phi}(x_i)}{\sum_{i=1}^{n+1} f_{\phi}(x_i)} > 0 \tag{4.14b}$$

$$\Leftrightarrow \left(\sum_{i=1}^{n+1} x_i f_{\phi_l}(x_i) \right) \left(\sum_{i=1}^{n+1} f_{\phi}(x_i) \right) > \left(\sum_{i=1}^{n+1} x_i f_{\phi}(x_i) \right) \left(\sum_{i=1}^{n+1} f_{\phi_l}(x_i) \right) \tag{4.14c}$$

$$\Leftrightarrow L_1 + L_2 + L_3 + L_4 > R_1 + R_2 + R_3 + R_4 \tag{4.14d}$$

where

$$\begin{aligned} L_1 &= \sum_{i=1}^n x_i f_{\phi_l}(x_i) \sum_{i=1}^n f_{\phi}(x_i), & R_1 &= \sum_{i=1}^n x_i f_{\phi}(x_i) \sum_{i=1}^n f_{\phi_l}(x_i), \\ L_2 &= f_{\phi}(x_{n+1}) \sum_{i=1}^n f_{\phi_l}(x_i) x_i, & R_2 &= f_{\phi_l}(x_{n+1}) \sum_{i=1}^n f_{\phi}(x_i) x_i, \\ L_3 &= x_{n+1} f_{\phi_l}(x_{n+1}) \sum_{i=1}^n f_{\phi}(x_i), & R_3 &= x_{n+1} f_{\phi}(x_{n+1}) \sum_{i=1}^n f_{\phi_l}(x_i), \\ L_4 &= x_{n+1} f_{\phi_l}(x_{n+1}) f_{\phi}(x_{n+1}), & R_4 &= x_{n+1} f_{\phi_l}(x_{n+1}) f_{\phi}(x_{n+1}). \end{aligned}$$

With $L_1 > R_1$ (according to inductive assumption $\mathbf{S}(\mathbf{n})$) and $L_4 = R_4$, we obtain

$$\mathbf{S}(\mathbf{n}) \wedge (L_2 + L_3 \geq R_2 + R_3) \Rightarrow \mathbf{S}(\mathbf{n} + 1). \tag{4.15}$$

Thus, it is sufficient to show:

$$L_2 + L_3 \geq R_2 + R_3 \quad (4.16a)$$

$$\Leftrightarrow f_\phi(x_{n+1}) \sum_{i=1}^n f_{\phi_l}(x_i) x_i + x_{n+1} f_{\phi_l}(x_{n+1}) \sum_{i=1}^n f_\phi(x_i) \quad (4.16b)$$

$$\geq f_{\phi_l}(x_{n+1}) \sum_{i=1}^n f_\phi(x_i) x_i + x_{n+1} f_\phi(x_{n+1}) \sum_{i=1}^n f_{\phi_l}(x_i) \quad (4.16c)$$

$$\Leftrightarrow f_{\phi_l}(x_{n+1}) \left(\sum_{i=1}^n x_{n+1} f_\phi(x_i) - \sum_{i=1}^n x_i f_\phi(x_i) \right) \quad (4.16c)$$

$$\geq f_\phi(x_{n+1}) \left(\sum_{i=1}^n x_{n+1} f_{\phi_l}(x_i) - \sum_{i=1}^n x_i f_{\phi_l}(x_i) \right) \quad (4.16d)$$

$$\Leftrightarrow f_{\phi_l}(x_{n+1}) \sum_{i=1}^n (x_{n+1} - x_i) f_\phi(x_i) - f_\phi(x_{n+1}) \sum_{i=1}^n (x_{n+1} - x_i) f_{\phi_l}(x_i) \geq 0 \quad (4.16d)$$

$$\Leftrightarrow \sum_{i=1}^n (x_{n+1} - x_i) \frac{f_\phi(x_i)}{f_\phi(x_{n+1})} - \sum_{i=1}^n (x_{n+1} - x_i) \frac{f_{\phi_l}(x_i)}{f_{\phi_l}(x_{n+1})} \geq 0 \quad (4.16e)$$

$$\Leftrightarrow \sum_{i=1}^n (x_{n+1} - x_i) \left(\frac{f_\phi(x_i)}{f_\phi(x_{n+1})} - \frac{f_{\phi_l}(x_i)}{f_{\phi_l}(x_{n+1})} \right) \geq 0 \quad (4.16f)$$

$$\Leftrightarrow \sum_{i=1}^n A_i B_i \geq 0, \quad (4.16g)$$

with

$$A_i = x_{n+1} - x_i, \quad B_i = \frac{f_\phi(x_i)}{f_\phi(x_{n+1})} - \frac{f_{\phi_l}(x_i)}{f_{\phi_l}(x_{n+1})}.$$

According to Equation 4.12,

$$\forall i, \quad A_i \geq 0. \quad (4.17)$$

Reformulating

$$B_i \geq 0 \Leftrightarrow \frac{f_\phi(x_i)}{f_\phi(x_{n+1})} \geq \frac{f_{\phi_l}(x_i)}{f_{\phi_l}(x_{n+1})} \quad (4.18a)$$

$$\Leftrightarrow \frac{f_{\phi_l}(x_{n+1})}{f_\phi(x_{n+1})} \geq \frac{f_{\phi_l}(x_i)}{f_\phi(x_i)} \quad (4.18b)$$

$$\Leftrightarrow g(x_{n+1}) \geq g(x_i), \quad (4.18c)$$

which is true for all i according to equations 4.11 and 4.12. Thus, with equations 4.17 and 4.18, Equation 4.16 is also true, which in turn proves the first case of Equation 4.9. \square

Remark

For sake of simplicity the above derivation assumed a monotonically increasing fitness landscape. Following an analogous approach it can be shown that equation 4.9 also holds for monotonically *decreasing* fitness landscapes. In that case, however, the selection differential is negative and $S_{x_l} - S_x < 0$ implies that learning accelerates evolution toward the higher fitness region. Thus, if $f'(z) < 0$ learning accelerates evolution if $g'(x) < 0$ and decelerates it if $g'(x) > 0$.

4.4 Extended Gain Function Framework

The gain function as formulated in Section 4.3.1 compares a learning versus a non-learning population and shows under what conditions the learning population evolves quicker (slower) toward a higher fitness region than the non-learning population.

In the following, the gain function framework is extended in order to predict how a change in a learning parameter impacts the influence of learning on evolution. In this section, the extended gain function is first formulated and then proved.

4.4.1 Formulation

The extended gain function framework assumes that there exists a learning parameter a that influences evolution. More generally, a can be interpreted as any kind of influence on the phenotype, such as an environmental influence during development, noise etc. In particular, it is assumed that phenotype z is determined by a and the genotype value x , i.e.,

$$z = \phi(x, a) , \quad (4.19)$$

and the corresponding fitness is,

$$f(z) = f(\phi(x, a)) . \quad (4.20)$$

For convenience $f(\phi(x, a))$ is denoted as $f_\phi(x, a)$. In the same fashion as in Equation 4.6, the expected change of the mean genotype in one generation is derived as

$$S_x = \frac{\sum_{i=1}^n x_i f_\phi(x_i, a)}{\sum_{i=1}^n f_\phi(x_i, a)} - \bar{x} , \quad (4.21)$$

where \bar{x} denotes the population mean genotype before selection. The influence of learning on evolutionary change in the population mean genotype can be predicted by analyzing its effect on S_x . If, e.g., an increase in a makes the selection differential larger (smaller), learning is predicted to accelerate (decelerate) evolution. Thus, an increase in learning parameter a , accelerates (decelerates) evolution if $\partial S_x / \partial a$ has the same sign as S_x . For example, if $S_x > 0$, learning accelerates evolution, if $\partial S_x / \partial a > 0$. Notice that $S_x > 0$ corresponds to the case of a fitness landscape that is increasing in positive x -direction, i.e., $\partial f(z) / \partial z > 0$.

In the next section, it is shown that the effect of learning on the rate of evolution is determined by $\partial^2 \log f_\phi(x, a) / \partial x \partial a$, in particular

$$\text{if, } \forall x \in]x_{\min}, x_{\max}[, \begin{cases} \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) > 0, \text{ then } \frac{\partial S_x}{\partial a} > 0 \\ \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) < 0, \text{ then } \frac{\partial S_x}{\partial a} < 0 \\ \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) = 0, \text{ then } \frac{\partial S_x}{\partial a} = 0 \end{cases}, \quad (4.22)$$

under the assumption that $f_{\phi}(x, a)$ is a monotonic function of x and differentiable within the population range $]x_{\min}, x_{\max}[$. The term

$$\boxed{\frac{\partial^2 \log f_{\phi}(x, a)}{\partial x \partial a}}$$

quantifies how a marginal gain in relative fitness from increased learning parameter a (the derivative of the logarithm) changes with the genotypic value and is therefore referred to as the *fitness gain derivative* (of the extended gain function framework). In the case that $\phi(x, a)$ is a stochastic function, $f_{\phi}(x, a)$ needs to be replaced by the expected fitness $\bar{f}_{\phi}(x, a)$.

For the sake of analysis (see next subsection) the finite population is approximated by a probability distribution of the population, $p(x)$. Thus, Equation 4.21 is reformulated as

$$S_x = \frac{\int_{x=x_{\min}}^{x=x_{\max}} xp(x)f_{\phi}(x, a)}{\int_{x=x_{\min}}^{x=x_{\max}} p(x)f_{\phi}(x, a)} - \bar{x}. \quad (4.23)$$

4.4.2 Proof

This section shows Equation 4.22 where S_x is defined as in Equation 4.23.

Outline

Equation 4.22 is split into three cases (A,B,C) which are treated separately in later steps of the proof.

$$\text{if, } \forall x \in]x_{\min}, x_{\max}[, \begin{cases} \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) > 0, \text{ then } \frac{\partial S_x}{\partial a} > 0 \text{ (case A)} \\ \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) < 0, \text{ then } \frac{\partial S_x}{\partial a} < 0 \text{ (case B)} \\ \frac{\partial^2}{\partial x \partial a} \log f_{\phi}(x, a) = 0, \text{ then } \frac{\partial S_x}{\partial a} = 0 \text{ (case C)} \end{cases}. \quad (4.24)$$

After defining

$$Q(x_0, x_1) = \left(\int_{x_0}^{x_1} p(x)f_{\phi}(x, a)dx \right) \left(\int_{x_0}^{x_1} xp(x)\frac{\partial f_{\phi}(x, a)}{\partial a}dx \right) - \left(\int_{x_0}^{x_1} xp(x)f_{\phi}(x, a)dx \right) \left(\int_{x_0}^{x_1} p(x)\frac{\partial f_{\phi}(x, a)}{\partial a}dx \right), \quad (4.25)$$

we obtain,

$$\frac{\partial S_x}{\partial a} = \frac{Q(x_{\min}, x_{\max})}{(\bar{f}_{\phi})^2}, \quad (4.26)$$

where \bar{f}_ϕ denotes the mean absolute fitness of the population. Thus, the sign of $\partial S_x/\partial a$ is determined by the sign of $Q(x_{\min}, x_{\max})$, and proving Equation 4.24 reduces to showing that $Q(x_{\min}, x_{\max})$ has the same sign as the fitness gain derivative, i.e.,

$$\text{sign}(Q(x_{\min}, x_{\max})) = \text{sign}\left(\frac{\partial^2 f_\phi(x, a)}{\partial x \partial a}\right). \quad (4.27)$$

In the following, it is first shown that the sign of $\partial^2 f_\phi(x, a)/\partial x \partial a$ determines the sign of the corresponding expression defined for a narrow interval within the distribution of x , $Q(x_0, x_0 + \delta)$, where $x_{\min} < x_0 < x_0 + \delta < x_{\max}$ and δ is small enough for the functions to be treated as linear. Then it is shown that, for any $x_{\min} < x_0, x_1 < x_{\max}$, widening the interval (i.e., increasing x_1 or decreasing x_0) does not change the sign of $Q(x_0, x_1)$, and so $Q(x_{\min}, x_{\max})$ has the same sign as $Q(x_0, x_0 + \delta)$.

Proof for a narrow x -interval

Within a narrow interval $]x_0, x_0 + \delta[$, the following linear approximations can be made:

$$f_\phi(x, a) = f_\phi(x_0, a) + (x - x_0) \frac{\partial f_\phi(x_0, a)}{\partial x}, \quad (4.28)$$

$$\frac{\partial f_\phi(x, a)}{\partial a} = \frac{\partial f_\phi(x_0, a)}{\partial a} + (x - x_0) \frac{\partial^2 f_\phi(x_0, a)}{\partial x \partial a}, \quad (4.29)$$

where $\partial f_\phi(x_0, a)/\partial x$ denotes $\partial f_\phi(x, a)/\partial x$ evaluated at $x = x_0$. The function $p(x)$ can be linearized in the same way, but it is advantageous here to express it as

$$p(x) = p(x_0) + \frac{x - x_0}{\delta} (p(x_0 + \delta) - p(x_0)). \quad (4.30)$$

With these substitutions, and after carrying out the integration and rearranging of terms, we obtain

$$Q(x_0, x_0 + \delta) = \frac{\delta^2}{72} (p^2(x_0) + 4p(x_0)p(x_0 + \delta) + p^2(x_0 + \delta)) \times \left[f_\phi(x_0, a) \frac{\partial^2 f_\phi(x_0, a)}{\partial x \partial a} - \frac{\partial f_\phi(x_0, a)}{\partial x} \frac{f_\phi(x_0, a)}{\partial a} \right]. \quad (4.31)$$

The term in the first set of brackets (upper row) is positive, so the sign of the expression on the right-hand side of Equation 4.31 depends on the term after the times sign (“ \times ”, bottom row). Notice, however that

$$\frac{\partial^2}{\partial x \partial a} \log(f_\phi(x, a)) = \frac{1}{(f_\phi(x, a))^2} \left[f_\phi(x, a) \frac{\partial^2 f_\phi(x, a)}{\partial x \partial a} - \frac{\partial f_\phi(x, a)}{\partial x} \frac{\partial f_\phi(x, a)}{\partial a} \right]. \quad (4.32)$$

Thus, the sign of $Q(x_0, x_0 + \delta)$ has the same sign as the fitness gain derivative, i.e.,

$$\text{sign}(Q(x_0, x_0 + \delta)) = \text{sign}\left(\frac{\partial^2 \log(f_\phi(x, a))}{\partial x \partial a}\right) \quad (4.33)$$

evaluated at x_0 .

Extending the x -interval

The proposition of Equation 4.24 requires that the same holds for $Q(x_{\min}, x_{\max})$, assuming that the sign of $\partial^2 \log(f_\phi(x, a)) / \partial x \partial a$ is constant throughout the interval. In other words, it needs to be shown that as the limits of the integrals in $Q(x_0, x_1)$ are extended from $]x_0, x_0 + \delta[$ to $]x_{\min}, x_{\max}[$, the sign of $Q(x_0, x_1)$ does not change.

Consider first **extending the upper limit** x_1 :

$$\begin{aligned}
 \frac{\partial Q(x_0, x_1)}{\partial x_1} &= x_1 p(x_1) \frac{\partial f_\phi(x_1, a)}{\partial a} \int_{x_0}^{x_1} p(x) f_\phi(x, a) dx + \\
 &\quad p(x_1) f_\phi(x_1, a) \int_{x_0}^{x_1} x p(x) \frac{\partial f_\phi(x, a)}{\partial a} dx - \\
 &\quad p(x_1) \frac{\partial f_\phi(x_1, a)}{\partial a} \int_{x_0}^{x_1} x p(x) f_\phi(x, a) dx - \\
 &\quad x_1 p(x_1) f_\phi(x_1, a) \int_{x_0}^{x_1} p(x) \frac{\partial f_\phi(x, a)}{\partial a} dx \quad .
 \end{aligned} \tag{4.34}$$

Notice that for the reformulation of the derivative the second fundamental theorem of calculus [85] is employed which is formulated as follows: If h is a function that is continuous on an open interval I and if b is any point in the interval I , then

$$\forall x \in I : \frac{\partial}{\partial x} \int_b^x h(y) dy = h(x) .$$

Equation 4.34 can be simplified by extracting $p(x_1)$, placing all other terms under a single integral and rearranging:

$$\begin{aligned}
 &\frac{\partial Q(x_0, x_1)}{\partial x_1} \\
 &= p(x_1) \left[\frac{\partial f_\phi(x_1, a)}{\partial a} \int_{x_0}^{x_1} (x_1 - x) p(x) f_\phi(x, a) dx - f_\phi(x_1, a) \int_{x_0}^{x_1} (x_1 - x) p(x) \frac{\partial f_\phi(x, a)}{\partial a} dx \right] \\
 &= p(x_1) \int_{x_0}^{x_1} (x_1 - x) p(x) \left[f_\phi(x, a) \frac{\partial f_\phi(x_1, a)}{\partial a} - f_\phi(x_1, a) \frac{\partial f_\phi(x, a)}{\partial a} \right] dx \\
 &= p(x_1) f_\phi(x_1, a) \int_{x_0}^{x_1} (x_1 - x) p(x) f_\phi(x, a) \left[\frac{1}{f_\phi(x_1, a)} \frac{\partial f_\phi(x_1, a)}{\partial a} - \frac{1}{f_\phi(x, a)} \frac{\partial f_\phi(x, a)}{\partial a} \right] dx \\
 &= p(x_1) f_\phi(x_1, a) \int_{x_0}^{x_1} (x_1 - x) p(x) f_\phi(x, a) \left[\frac{\partial \log f_\phi(x_1, a)}{\partial a} - \frac{\partial \log f_\phi(x, a)}{\partial a} \right] dx \quad .
 \end{aligned} \tag{4.35}$$

For $x = x_1$, the function under the last integral equals zero; for $x < x_1$, its sign is determined by the term in the last parentheses, which has the same sign as the fitness gain derivative, i.e.,

$$\text{sign} \left(\frac{\partial \log f_\phi(x_1, a)}{\partial a} - \frac{\partial \log f_\phi(x, a)}{\partial a} \right) = \text{sign} \left(\frac{\partial^2 \log f_\phi(x, a)}{\partial x \partial a} \right) . \tag{4.36}$$

Notice that

$$\forall x < x_1 : \frac{\partial^2 \log(f_\phi(x, a))}{\partial x \partial a} > 0 \Rightarrow \frac{\partial(f_\phi(x_1, a))}{\partial a} > \frac{\partial(f_\phi(x, a))}{\partial a} , \tag{4.37}$$

and vice versa. Hence, the sign of $\partial Q(x_0, x_1)/\partial x_1$ is the same sign as fitness gain derivative, i.e.,

$$\frac{\partial Q(x_0, x_1)}{\partial x_1} = \text{sign} \left(\frac{\partial^2 \log(f_\phi(x, a))}{\partial x \partial a} \right), \quad (4.38)$$

assuming that the sign of the latter is constant within interval (x_0, x_1) and that $p(x_1)f_\phi(x_1, a) > 0$.

Similarly, the effect of **extending the lower limit** x_0 is described by

$$\frac{\partial Q(x_0, x_1)}{\partial x_0} = p(x_0)f_\phi(x_0, a) \int_{x_0}^{x_1} (x_0 - x)p(x)f_\phi(x, a) \left(\frac{\partial \log f_\phi(x, a)}{\partial a} - \frac{\partial \log f_\phi(x_0, a)}{\partial a} \right) dx. \quad (4.39)$$

For $x > x_0$, the term in the last parentheses of Equation 4.39 has the same sign as the fitness gain derivative, i.e.,

$$\text{sign} \left(\frac{\partial \log f_\phi(x, a)}{\partial a} - \frac{\partial \log f_\phi(x_0, a)}{\partial a} \right) = \text{sign} \left(\frac{\partial^2 \log(f_\phi(x, a))}{\partial x \partial a} \right), \quad (4.40)$$

however, in Equation 4.39 the term $x_0 - x$ is negative, so the function under the integral has the opposite sign from $\partial^2 \log(f_\phi(x, a))/\partial x \partial a$ (the fitness gain derivative) for $x_0 < x < x_1$.

Conclusion

The above argument proves the proposition of Equation 4.24 as follows.

Case A: Consider first case A of Equation 4.24. The fitness gain derivative

$$\partial^2 \log(f_\phi(x, a))/\partial x \partial a > 0$$

for all $x \in]x_{\min}, x_{\min}[$, if $Q(x_0, x_1) > 0$ for any interval $]x_0, x_1[$ of width δ within $]x_{\min}, x_{\min}[$ (Equation 4.31). Furthermore, $\partial Q(x_0, x_1)/\partial x_0 \leq 0$ and $\partial Q(x_0, x_1)/\partial x_1 \geq 0$, so as the interval is extended in either direction (increasing x_1 toward x_{\max} or decreasing x_0 toward x_{\min}), $Q(x_0, x_1)$ remains positive (Equations 4.35 and 4.39). Hence, $Q(x_{\min}, x_{\max}) > 0$ and $\partial S_x/\partial a > 0$ (Equation 4.26), which proves case A of Equation 4.24.

Case B: Proof of case B of Equation 4.24 is analogous: The fitness gain derivative

$$\partial^2 \log(f_\phi(x, a))/\partial x \partial a < 0$$

for all $x \in]x_{\min}, x_{\max}[$, if $Q(x_0, x_1) < 0$ for a narrow interval of width δ ; furthermore, $\partial Q(x_0, x_1)/\partial x_0 \geq 0$ and $\partial Q(x_0, x_1)/\partial x_1 \leq 0$; hence, $Q(x_{\min}, x_{\max}) < 0$ and $\partial S_x/\partial a < 0$.

Case C: Finally, for case C of Equation 4.24: The fitness gain derivative

$$\partial^2 \log(f_\phi(x, a))/\partial x \partial a = 0$$

for all $x \in]x_{\min}, x_{\max}[$, if $Q(x_0, x_1) = 0$ for a narrow interval of width δ , and it remains zero as the interval is broadened because $\partial Q(x_0, x_1)/\partial x_0 = 0$ and $\partial Q(x_0, x_1)/\partial x_1 = 0$; hence $Q(x_{\min}, x_{\max}) = \partial S_x/\partial a = 0$.

□

4.5 Summary and Conclusion

In this chapter, a general framework which we call the *gain function* to predict the influence of learning on the rate of evolution has been presented. The gain function is formulated in terms of the effect of learning on the mapping from genotype to fitness. Figure 4.3 illustrates the analysis results.

In its first formulation in Section 4.3, the gain function can be used to predict the effect of *adding* individual learning to the evolutionary process. In its second formulation in Section 4.4, where it was formulated as the fitness gain derivative it can be used to predict the influence of changing a learning parameter on the rate of evolution. In the remainder of this thesis, the framework introduced in Section 4.3 is referred to as *basic gain function framework* and that of Section 4.4 as *extended gain function framework*.

The gain function analysis looks at the *effect* of learning and does not require to consider a particular learning scheme or algorithm. All that is needed is to know how learning influences fitness.

As mentioned earlier the mapping from genotype to fitness might be stochastic. In this case the gain function can be applied by calculating the expected fitness of a genotype. It should be noted here that there are various sources from which the stochasticity may originate. The stochasticity may originate from the mapping between genotype and innate phenotype (development), from the mapping between innate and learned phenotype (learning), or from the mapping between phenotype and fitness. The actual calculation of the corresponding expected fitness may be quite elaborate in many cases.

Furthermore, it has to be mentioned that the formulation of the genotype to fitness mapping in the absence of learning ($f_\phi(x)$) and in the presence of learning ($f_{\phi_l}(x)$) has been formulated based on a simplified model that does not consider how learning changes a phenotype *over time*. However, the gain function framework is not limited to such simplified models. As mentioned above, all that is needed is to know how learning influences fitness of the genotype.

In biological terms, the gain function only applies to *directional selection* (selection that moves the population toward higher fitness, as opposed to *disruptive* or *stabilizing* selection). In other words, the gain function considers how learning influences selection pressure. The gain function analysis is expectation-based and does not account for the variance of the population movement. Thus, the gain function does not allow to make predictions on the influence of learning on the time needed to cross a fitness valley toward a region with higher fitness. Such a prediction cannot be made based on analysis of the expected behavior, since fitness valley crossing requires an “unlikely” event. Thus, a stochastic analysis is required to predict the time to cross a fitness valley. It turned out that during the work on this PhD thesis, such an approach was suggested in the PhD thesis of Elhanan Borenstein [12] (also published in [13]). Using an abstract random-walk model (no population) Borenstein essentially shows that the time needed to cross a fitness valley is positively correlated with the depth of the fitness valley. Borenstein’s model does not allow predictions with regard to directional selection.

The gain function makes *exact short term predictions* of the mean genotype movement. The gain function framework does not allow exact predictions of the dynamics when a population that initially populates a fitness landscape region with a positive gain function derivative moves on to a region with a negative gain function derivative (the gain function may be no




Figure 4.3: Illustration of the main result of the gain function development. An increasing gain function (left panel) indicates that relative fitness differences between genetically weak and strong individuals are enlarged through learning. A decreasing gain function (right panel) indicates that relative fitness differences between genetically weak and strong individuals are reduced through learning.

longer monotonic within the range of the population). It does, however, allow *approximate long-term predictions* of the mean genotype movement, as the next chapter will show.

The gain functions framework allows to predict the results of several (but not all) models that have been reviewed in Section 4.1 which is demonstrated in Chapter 6.

Conditions for Learning-Induced Acceleration and Deceleration of Evolution

In Chapter 4, the gain function framework has been introduced as a general tool to predict whether for a given coupling of learning and evolution, learning is expected to accelerate or decelerate evolution. In this chapter, the *basic* gain function framework of Section 4.3 is applied in order to get a better understanding of the dynamics of coupled evolution and learning.

First, a general learning function is investigated (Section 5.1, based on [128] and [130]). Then, in Section 5.2, the special case where the fitness can be decomposed into an innate component and a learning component is analyzed with the gain function (based on [129]). Afterwards, on a more fine-grained level, it is shown that the shape of learning curves has an influence on the rate of evolution, as well (Section 5.3, again based on [129]). In Section 5.4, it demonstrated that also a non-monotonic gain function may be a good predictor for the population dynamics (based on [130]). Section 5.5 completes this chapter with summary and conclusion.

5.1 A General Learning Function

The gain function analysis concentrates on the *effect* of learning on phenotype and fitness, abstracting from the dynamics of learning and the underlying process, i.e., the learning technique or the learning algorithm. Therefore, in contrast to the existing literature, e.g., [64, 46, 30], no specific model of how learning occurs is introduced here. In principle, the result of any learning algorithm can be split into a "directional" and a "noise" part.

Firstly, learning in nature usually results in an improvement of function, behavior, skill, etc., and similarly in evolutionary computation, learning usually results in a higher solution

quality. In this thesis, this aspect of learning is called the directional part of learning, or *directional learning*.¹

Secondly, the results of learning may not always be the same, even if the learning procedure itself is identical. In evolutionary computation, probabilistic learning algorithms may produce noisy results. Also, in nature, not all learning efforts are immediately successful. Organisms might experience setbacks, are forgetful and new skills might interfere with previously learned ones. Furthermore, two individuals will have different experiences and often different degree of success even under identical learning schemes. In this thesis, this aspect of learning is called the noise part of learning, or *learning noise*. Another interpretation of the noise component is to treat it as developmental noise, i.e., the result of a development process from genotype x to phenotype z , $\phi(x) = z$, is usually noisy as well. Thus the results concerning the effect of noise also apply to developmental noise.

A general learning function which describes the effect of learning on the phenotype can be defined as

$$l(z) = z + \delta + \varepsilon \quad (5.1)$$

where z the one-dimensional real-valued phenotype value, δ is the directional component of learning (the average effect of learning on the phenotype) and ε is a random number sampled from a distribution with zero mean. Referring to the Nomenclature table of this thesis the learning parameter set of the general learning function could be defined as $a = (\delta, \sigma_\varepsilon)$, where σ_ε^2 is the variance of ε .

In the following, the two effects of learning are treated separately, first directional learning and then learning noise. For both types, the gain function framework is applied to study the effects of the respective learning type on the rate of evolutionary change. In particular, this analysis considers the shape of the fitness landscape. In the following the fitness landscape is referred to as *fitness function*.

This chapter concentrates on the influence of learning. Therefore, a simple development function ϕ is assumed that eases the analysis, in particular,

$$z = \phi(x) = x. \quad (5.2)$$

Thus, fitness in absence of learning is given by $f(x)$ and in the presence of learning by $f(l(x))$.

5.1.1 Directional Learning

Recalling Equation 5.1, a simple form of directional learning is defined as

$$l_\delta(z) = z + \delta, \quad (5.3)$$

where δ is a constant. It is assumed further that $\text{sign}(\delta) = \text{sign}(f'(z))$, i.e., that learning modifies behavior in the direction of higher fitness (otherwise learning would be maladaptive). This form of directional learning is illustrated in Figure 5.1.

¹ Notice that the term "directional learning" is also used in economic game theory [151]. There it describes the behavior of a decision maker that adjust the "direction" of his decision based on comparisons of his past decision and alternative decisions that he could have taken. However, the definition used in this chapter differs from the one used in economic science.




Figure 5.1: *Illustration of directional learning as defined in Equation 5.3. Directional learning (l_δ) involves a shift in the expected phenotype in the direction of higher fitness. With a non-linear fitness function, even the same directional effect of learning on the phenotype (δ) will result in different changes in fitness depending on the genotype value of the individual. With the concave fitness function shown in the figure a genetically weak individual (x_w) gains more from the same phenotype change due to learning than a genetically strong individual (x_s).*

First, the conditions for learning-induced acceleration and deceleration are derived with the simple gain function $g(x) = f_{\phi_l}(x)/f_\phi(x)$ as defined in Section 4.3. With Equation 5.2, the mapping from genotype x to fitness is $f(x)$ in absence of learning and $f(x + \delta)$ in presence of learning. Assuming a monotonical and continuously differentiable fitness function f , the sign of the derivative of the gain function is

$$\begin{aligned}
 \text{sign}(g'_\delta) &= \text{sign} \left[\frac{\partial}{\partial x} \left(\frac{f(x + \delta)}{f(x)} \right) \right] \\
 &= \text{sign} \left[\frac{\partial}{\partial x} \left(\frac{f'(x + \delta)f(x) - f(x + \delta)f'(x)}{f^2(x)} \right) \right] \\
 &= \text{sign} \left[\frac{\partial}{\partial x} \left(\frac{f'(x + \delta)}{f(x + \delta)} - \frac{f'(x)}{f(x)} \right) \right] \\
 &= \text{sign} \left[\frac{\partial}{\partial x} (\log(f(x + \delta)) - \log(f(x))) \right] \\
 &= \text{sign} \left[\delta \frac{\partial^2}{\partial x^2} \log(f(x)) \right] .
 \end{aligned} \tag{5.4}$$

The last equality follows from the relationship

$$\text{sign}(h'(x)) = \text{sign}((x_1 - x_2)(h(x_1) - h(x_2))) , \tag{5.5}$$

which holds for any monotonic function $h(x)$ and arbitrary x_1, x_2 with $x_1 \neq x_2$; in Equation 5.4, $h(x) = (\log(f(x)))'$.

Recall that learning accelerates evolution if $g'(x)$ has the same sign as $f'(x)$, and that above it was assumed $\text{sign}(\delta) = \text{sign}(f'(x))$. Therefore, directional learning as defined in Equation 5.3 is predicted to accelerate evolution if the logarithm of the fitness function is convex (positive second derivative). Conversely, if the logarithm of the fitness function is concave (negative second derivative) evolution slows down as a result of directional learning.

The same result can be obtained by applying the extended gain function framework as defined in Section 4.4. Since for all functions h ,

$$\frac{\partial}{\partial x}h(x+y) = \frac{\partial}{\partial y}h(x+y) ,$$

and

$$\frac{\partial^2}{\partial x \partial y}h(x+y) = \frac{\partial^2}{\partial x^2}h(x+y) ,$$

the fitness gain derivative (cf. Equation 4.22 and thereafter) can be rewritten as

$$\frac{\partial^2}{\partial x \partial \delta} \log f_\phi(x, \delta) = \frac{\partial^2}{\partial x \partial \delta} \log f(x + \delta) = \frac{\partial^2}{\partial x^2} \log f(x + \delta) = \frac{\partial^2}{\partial x^2} \log f(x) , \quad (5.6)$$

which confirms Equation 5.4 for positive δ . Notice that the case of negative δ cannot directly be treated with the extended gain function².

It is straightforward to calculate what a convex or concave logarithm of a function implies for the function itself.

$$\text{sign} \left[\frac{\partial^2}{\partial x^2} \log f(z) \right] = \text{sign} \left[\frac{f''(z)f(z) - (f'(z))^2}{f^2(z)} \right] = \text{sign} [f''(z)f(z) - (f'(z))^2] . \quad (5.7)$$

Assume f is monotonically increasing (the opposite case can be treated in analogous fashion). Thus, if f is concave ($f''(z) < 0$) the sign in Equation 5.7 is negative. However, if f is convex ($f''(z) > 0$) the sign in Equation 5.7 is not obvious.

In conclusion, directional learning accelerates evolution on all fitness functions with convex logarithm and decelerates evolution on all fitness functions with concave logarithm. This implies that evolution is decelerated on all concave fitness functions. It does, however, not imply that learning accelerates evolution on all convex fitness functions, but a *necessary* condition for accelerated evolution through directional learning is non-concavity of the fitness function.

² The case of negative δ cannot directly be treated with the extended gain function analysis because the latter approach considers a marginal change of the parameter δ in *positive* direction. Thus, taking the derivative of the monotonically increasing $f(x + \delta)$ with respect to δ represents a marginal decrease in the distance to the global optimum. In order to study the effect of a marginal increase in the distance to the global optimum through directional learning with the extended gain function one would have to reformulate the directional learning function such that $f(x - \delta)$. Taking the derivative with respect to δ on the reformulated function can be interpreted then as a marginal increase of the distance between x and the global optimum (again assuming f is monotonically increasing).

The same logic applies to a monotonically decreasing function. Given the above analysis, the calculations for the case that corresponds to a negative δ in Equation 5.4 are straightforward, and are therefore omitted here.




Figure 5.2: *Illustration of learning noise as defined in Equation 5.8. Learning noise (l_ε) adds to the phenotype of the individual a random number ε sampled from a distribution with a zero mean and a range $[\varepsilon_{min}, \varepsilon_{max}]$. For an individual with a genotype value x_i , a possible fitness loss is represented by the gray area above the curve, a possible fitness gain by the gray area below the curve. Thus, an individual has a positive expected overall fitness gain if the area above is larger than the corresponding area below the curve, i.e., the gain function $g(x) > 1$. If the distribution of ε is symmetric and the fitness function convex, $g(x) > 1$ for all x ; if the fitness function is concave (as in the figure), $g(x) < 1$, i.e., the noise on average leads to a fitness loss irrespective of x . However, the important issue with regard to the rate of evolution is, whether a genetically strong individual x_s on average gains more or loses less than the genetically weak individual x_w .*

5.1.2 Learning Noise

Next, the second component of the general learning function (cf. Equation 5.1), *learning noise*, is considered, which is defined as

$$l_\varepsilon(x) = x + \varepsilon, \quad (5.8)$$

where ε is a random number with zero mean and a symmetric probability distribution $p(\varepsilon)$ whose parameters are independent of x . This form of learning noise is illustrated in Figure 5.2.

In order to apply the gain function framework, the fitness corresponding to a given genotype value x has to be averaged over all phenotypes expressed by individuals with this genotype value. Using the Taylor series expansion, this expected fitness can be written as

$$\begin{aligned}\bar{f}(l_\varepsilon(x)) &= \int_{-\varepsilon_{\max}}^{+\varepsilon_{\max}} p(\varepsilon) f(x + \varepsilon) d\varepsilon \\ &= \sum_{i=0}^{\infty} \frac{f^{(i)}(x)}{i!} \int_{-\varepsilon_{\max}}^{+\varepsilon_{\max}} p(\varepsilon) \varepsilon^i d\varepsilon \\ &= \sum_{i=0}^{\infty} \frac{f^{(i)}(x)}{i!} \alpha_i.\end{aligned}\tag{5.9}$$

where α_i is the i 'th moment of $p(\varepsilon)$. Since $p(\varepsilon)$ is symmetric around 0, $\alpha_i = 0$ for all odd i , and the third order Taylor approximation gives

$$\bar{f}(l_\varepsilon(x)) \cong f(x) + \frac{\text{var}(\varepsilon)}{2} f''(x),\tag{5.10}$$

where $\text{var}(\varepsilon)$ is the variance of ε . Therefore, the simple gain function as defined in Section 4.3 can be approximated by

$$g(x) \cong 1 + \frac{\text{var}(\varepsilon)}{2} \frac{f''(x)}{f(x)},\tag{5.11}$$

and correspondingly

$$g'(x) \cong \frac{\text{var}(\varepsilon)}{2} \frac{f(x)f'''(x) - f'(x)f''(x)}{(f(x))^2},\tag{5.12}$$

which implies

$$\text{sign}(g'(x)) \cong \text{sign}(f(x)f'''(x) - f'(x)f''(x)).\tag{5.13}$$

Recall that learning is predicted to accelerate evolution if $g'(x)$ has the same sign as $f'(x)$, and decelerated if the signs are opposite. The above result holds for all symmetric noise distributions, as long as the fitness function can be sufficiently well approximated by the third-order Taylor series.

The same result can be obtained by applying the extended gain function framework as defined in Section 4.4. The fitness gain derivative (cf. Equation 4.22 and thereafter) can be rewritten as

$$\begin{aligned}\frac{\partial^2}{\partial x \partial \delta} \log f_\phi(x, \delta) &\cong f(x) + \frac{\text{var}(\varepsilon)}{2} f''(x) \\ &= 2 \frac{f(x)f'''(x) - f'(x)f''(x)}{(2f(x) + \text{var}(\varepsilon)f''(x))^2},\end{aligned}\tag{5.14}$$

thus,

$$\text{sign}\left(\frac{\partial^2}{\partial x \partial \delta} \log f_\phi(x, \delta)\right) \cong \text{sign}(f(x)f'''(x) - f'(x)f''(x)),\tag{5.15}$$

which confirms Equation 5.13.

Assuming $\forall x : f(x) > 0, f'(x) > 0$ part of Equation 5.13 can be proven without Taylor series approximation. In particular, it can be shown that

$$\forall x : f''(x) > 0 \wedge f'''(x) \leq 0 \Rightarrow g'(x) < 0 \quad , \quad (5.16a)$$

$$\forall x : f''(x) < 0 \wedge f'''(x) \geq 0 \Rightarrow g'(x) > 0 \quad , \quad (5.16b)$$

The proof for both cases of Equation 5.16 can be found in Appendix B.

5.1.3 Simulation

The above analysis delineates the conditions for learning to speed up or slow down evolution, but does not predict the magnitude of the effects. The latter issue is analyzed here with some simple computer simulations.

Simulation Set-Up

An asexual population of 10000 individuals is simulated, each characterized by a one-dimensional genotype value x , being a non-negative real number.

Recall (Equation 5.2 and thereafter) that in presence of learning fitness is given by $f(l(x))$ and in absence of learning by $f(x)$. Selection is simulated by Stochastic Universal Sampling [6], i.e., sampling (with replacement) of n offspring from n parents, where the probability of an individual being sampled is proportional to its fitness $f(l(x))$ respectively $f(x)$. Mutation is simulated by adding a random number from a normal distribution with parameters $\mu = 0$ and $\sigma = 10^{-3}$ to the genotype value x of each offspring (cut off at the genotype boundaries). Each simulation is initiated with all individuals having a genotype value equal to the lower boundary of the permitted range of x . The actual rate at which the mean genotype changes depends on the fitness function, the mutation strength, and the initial genotype distribution. The latter two parameters are not considered by the gain function analysis and are set to values that generate a visible population movement.

For the case of directional learning, the magnitude of learning is set to $\delta = 0.1$ (cf. Equation 5.3) and the following two fitness functions are studied

$$f_1(z) = e^{4z^2} \quad , \quad f_2(z) = z^{0.5} \quad .$$

According to Equation 5.4 directional learning is predicted to accelerate evolution on fitness function f_1 and to decelerate evolution on f_2 (notice that the logarithm of a function of the form $f(z) = z^a$ is concave for all $a > 0$).

For the case of learning noise, the learning function as defined in Equation (5.8) is implemented with ε sampled from a uniform distribution in the range $[-0.1, 0.1]$. To avoid negative phenotype values, x is constrained to $x \geq 0.1$. Two fitness functions are studied

$$f_3(z) = z^{0.4} \quad , \quad f_4(z) = z^{6.0} \quad .$$

According to Equation 5.13 learning noise should accelerate evolution on f_3 and decelerate evolution on f_4 . The functions f_1 to f_4 are chosen for the purpose of illustration and are not supposed to reflect a particular biological or computational scenario. For each setting 100 independent simulation runs are carried out, and the reported results are averaged over these simulation runs.

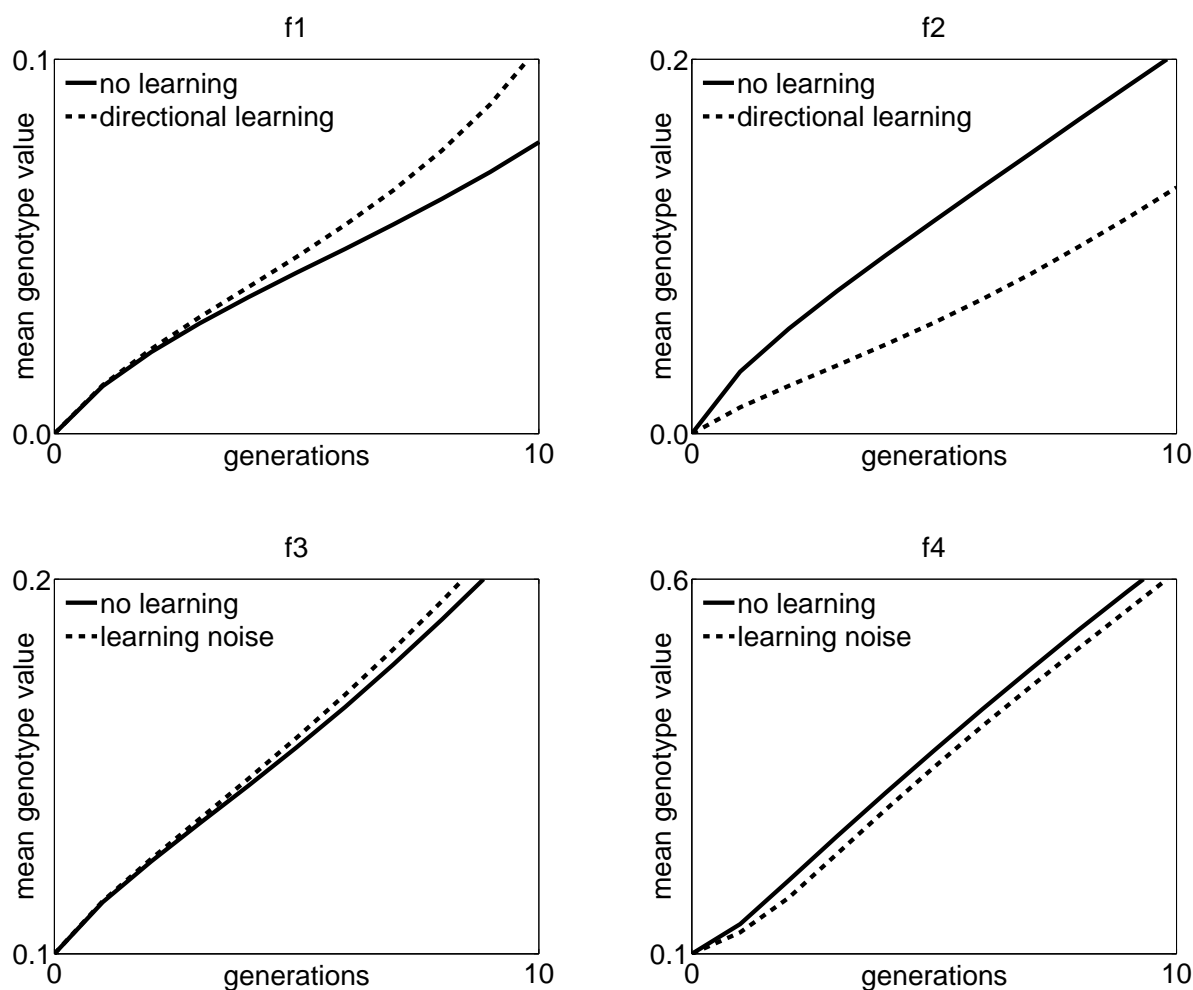


Figure 5.3: Simulation results of directional learning and learning noise. The evolutionary trajectories agree with the analytical predictions, i.e., directional learning accelerates evolution on f_1 and decelerates evolution on f_2 , learning noise accelerates evolution on f_3 , and decelerates evolution on f_4 . In these examples, the magnitude of the effect of learning noise is smaller than that of directional learning.

Simulation Results

Figure 5.3 shows the simulation results. The evolutionary trajectories agree with the analytical predictions. The simple form of directional learning accelerates evolution on function f_1 and decelerates evolution on function f_2 . Similarly, on function f_3 , learning noise accelerates, and on f_4 decelerates evolution. In these examples, the effect of learning noise is smaller than that of directional learning.

5.1.4 Conclusion

In this section, two components (a directional and a noise component) of learning have been identified that are common to most learning algorithms or procedures. The gain

function framework has been applied in order to derive general conditions under which these components accelerate or decelerate evolution. Based on simple assumptions on the effect of directional learning and learning noise the properties of the fitness function that determine whether evolution is accelerated have been identified. The results of the simulation study suggest that the effect of learning on the rate of evolutionary change is stronger in magnitude with directional than with learning noise. This supports the intuitive argument that directional learning is expected to change the mapping from genotype to fitness more drastically than symmetrically distributed noise.

5.2 Separable Fitness Components

In this section, a special case of coupled evolution and learning is analyzed, again with the gain function framework. As in the previous section, the most simple development function $z = \phi(x) = x$ is chosen here, in order to concentrate on the effect of learning. Thus fitness in case of learning is given by $f(l(x))$ and in absence of learning by $f(x)$ (the innate fitness).

In particular, it is assumed that the fitness of a learning individual is additively composed of an innate fitness component $f(x)$ and a learning component $f_L(x)$,

$$f(l(x)) = f(x) + f_L(x). \quad (5.17)$$

In the following, it is assumed that $f(x)$ is a positive and monotonically increasing function within the range of the population, i.e., $f(x) > 0, f'(x) > 0$. The gain function derivative can generally be calculated as

$$\begin{aligned} \frac{\partial g}{\partial x} &= \frac{\partial}{\partial x} \frac{f(x) + f_L(x)}{f(x)} \\ &= \frac{\partial}{\partial x} \left(1 + \frac{f_L(x)}{f(x)} \right) \\ &= (f(x))^{-2} \left(f(x) \frac{\partial f_L(x)}{\partial x} - f_L(x) \frac{\partial f(x)}{\partial x} \right) \\ &= \frac{f_L(x)}{f(x)} \left(\frac{\partial f_L(x)/\partial x}{f_L(x)} - \frac{\partial f(x)/\partial x}{f(x)} \right) \\ &= \frac{f_L(x)}{f(x)} \frac{\partial}{\partial x} (\log(f_L(x)) - \log(f(x))). \end{aligned} \quad (5.18)$$

If $f_L(x) > 0$ (and it is known that $f(x) > 0$), then

$$\text{sign} \left(\frac{\partial g(x)}{\partial x} \right) = \text{sign} \left(\frac{\partial(\log f_L(x))}{\partial x} - \frac{\partial(\log f(x))}{\partial x} \right). \quad (5.19)$$

Thus, if the first derivative of the logarithm of $f_L(x)$ is larger (smaller) than the first derivative of the logarithm of $f(x)$, then learning accelerates (decelerates) evolution. Notice that in the special case of $f_L(x) = f(x)$ learning has no influence on the rate of evolutionary change.

In the following, three categories of functions $f_L(x)$ are defined and further analyzed.

5.2.1 Positive, Decreasing $f_L(x)$

A positive, decreasing $f_L(x)$ implies that genetically weak individuals benefit more from learning than genetically strong ones. Intuitively, one would expect that learning decelerates evolution in this case. The following brief gain function analysis confirms this intuition. Since

$$f_L(x) > 0 \wedge \frac{\partial f_L(x)}{\partial x} < 0 \Rightarrow \frac{\partial(\log f_L(x))}{\partial x} < 0 \quad (5.20)$$

and

$$f(x) > 0 \wedge \frac{\partial f(x)}{\partial x} > 0 \Rightarrow \frac{\partial(\log f(x))}{\partial x} > 0 \quad (5.21)$$

and using Equation (5.19), one obtains $\frac{\partial g}{\partial x} < 0$. Therefore, for all scenarios with positive, decreasing function $f_L(x)$, learning decelerates evolution.

5.2.2 Constant $f_L(x)$

Next, the case when learning causes a constant fitness change, i.e., $f_L(x) = C$, is considered. Notice that this case is distinct from directional learning (Section 5.1.1) where a constant change takes place in genotype space rather than in fitness space.

With Equation (5.18), one obtains

$$\frac{\partial g}{\partial x} = -\frac{C}{f(x)} \frac{\partial(\log f(x))}{\partial x}, \quad (5.22)$$

and

$$\text{sign}\left(\frac{\partial g}{\partial x}\right) = \text{sign}(-C). \quad (5.23)$$

Therefore, in case of a constant fitness *increase* (positive C), evolution is *decelerated* through learning while for a constant fitness *decrease* (negative C) evolution is *accelerated* through learning. At first sight this may seem counter-intuitive. However, it is the *relative* fitness differences that determine the dynamics: A constant fitness increase implies a larger *relative* fitness gain for a weak individual (with small innate fitness) than for a strong individual (with large innate fitness). On the contrary, a constant fitness decrease (maladaptive learning) implies a larger *relative* fitness loss for a weak individual (with small innate fitness) than for a strong individual (with large innate fitness).

5.2.3 Positive, Increasing $f_L(x)$

Finally, the case of positive, increasing $f_L(x)$ is considered. For such functions, strong individuals always benefit more from learning than weak individuals (in terms of absolute fitness gain). Unfortunately, no simpler (and still general) formulation than Equation (5.19) can be derived for this case, without specifying either $f_L(x)$ or $f(x)$. Therefore two examples illustrate that functions of this category can either accelerate or decelerate evolution. If

$$f_L(x) = x^\alpha \quad (5.24)$$

and

$$f(x) = x^\beta, \quad (5.25)$$

then according to Equation 5.19,

$$\text{sign} \left(\frac{\partial(\log f_L(x))}{\partial x} - \frac{\partial(\log f(x))}{\partial x} \right) = \text{sign} \left(\frac{\alpha}{x} - \frac{\beta}{x} \right) = \text{sign}(\alpha - \beta) \quad (5.26)$$

determines whether evolution is accelerated ($\alpha > \beta$) or decelerated ($\alpha < \beta$).

5.2.4 Conclusion

The analysis of the three categories of learning-dependent fitness components, $f_L(x)$, has shown that (apart from maladaptive learning) evolution is only accelerated through learning in the case where the fitness of the learning-dependent component increases stronger in x than the fitness of the innate component.

5.3 Influence of Learning Curves on Evolution

So far, the mapping from genotype to fitness has been considered as a black box and it has not been discussed yet how lifetime fitness is actually attained. This was not necessary because for the gain function analysis, there is no need to know *how* the fitness was actually attained.

In artificial systems of coupled evolution and learning, the *result* of learning is often taken as the absolute fitness measure. One example for this is the evolution of artificial neural networks that are also individually trained. There, the neural network behavior after the training is the basis for the fitness assignment. Most of the papers cited in Yao's review of evolving artificial neural networks [188] follow this approach. In [101], the fitness assessment at the end of the individual's life has been named *posthumous fitness assessment*.

An alternative fitness assessment approach is to repeatedly evaluate an individual throughout its lifetime. In [101], this type of fitness assessment has been named *continual fitness assessment*. The latter type is biologically plausible and there are several artificial systems in which continual fitness assessment is applied. Several works in the field of evolutionary robotics [119] where a robot's (adaptive) control system is evolved and evaluated throughout the robot's lifetime, are examples of continual fitness assessment.

The gain function framework can handle both fitness assessment approaches as long as knowledge on the relative gain in fitness (the selection criterion) achieved through learning is available. This chapter focuses on the influence of learning curves on the rate of evolution. In particular, it will be analyzed, if and how the curvature of a learning curve influences evolution, also compared to the case when only the result of learning is the basis for selection (posthumous fitness assessment). For example, evolution of "early learners" may differ from the evolution of "late learners".

5.3.1 Extension of the Fitness Landscape Model

The traditional fitness landscape model that maps genotype to fitness or phenotype to fitness is not appropriate to visualize the influence of learning curves on absolute fitness. One way to

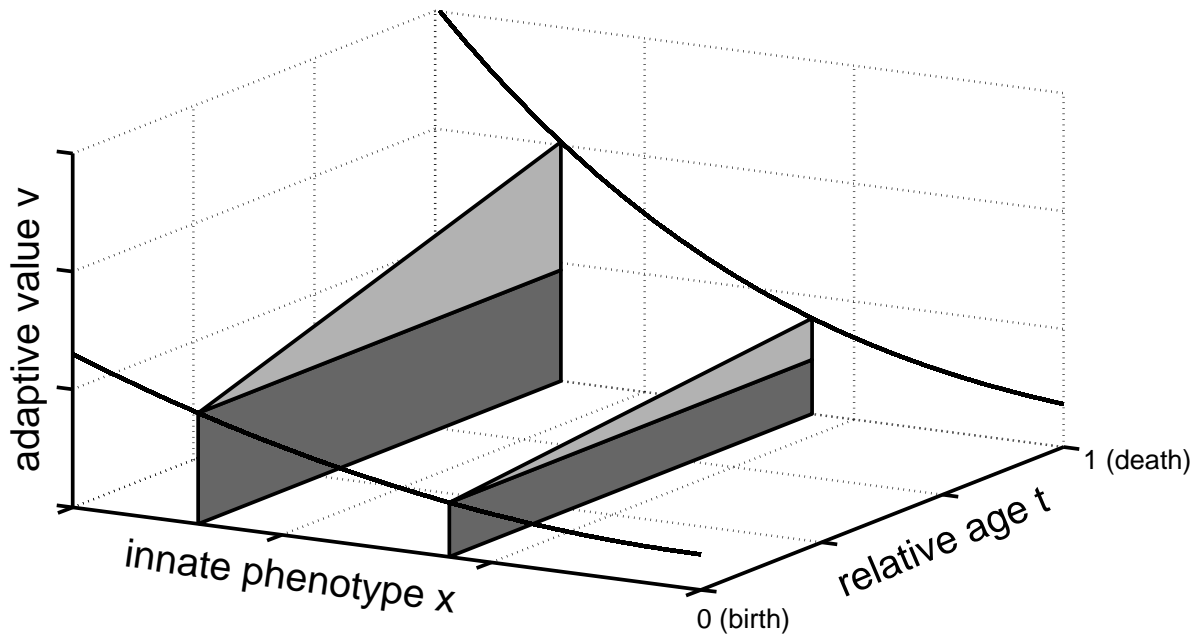


Figure 5.4: Extension of the fitness landscape model that accounts for learning curves.

define a learning curve of an individual with innate phenotype z_0 is the mapping from time (between birth and death) to the individual's adaptive value at a time. The average adaptive value achieved during lifetime can then be taken as the absolute fitness measure. Thus the adaptive value at a time t for a given innate phenotype can be visualized as in Figure 5.4

In order to concentrate on the effect of learning on evolution, again, a simple mapping from genotype $x \in \mathbb{R}$ to innate phenotype $z_0 \in \mathbb{R}$ is assumed

$$z_0 = \phi(x) = x . \quad (5.27)$$

Learning curves are defined w.r.t. the *relative* current age of an individual, i.e., between 0 (birth) and 1 (death). Accordingly, the (absolute) fitness f of an individual with genotype (and innate phenotype) x is given by

$$f(x) = \int_{t=0}^{t=1} v(x, t) dt . \quad (5.28)$$

In the absence of learning, the adaptive value is constant (in t -direction). In this case, the (absolute) fitness $f_\phi(x)$ is given by the size of the dark-gray area. In case of learning, (absolute) fitness $f_{\phi_l}(x)$ of an individual x is obtained by adding the size of the light-gray area (in Figure 5.4 a triangle) to the size of the dark-gray area. Posthumous assessment could also be visualized in the extended visualization of Figure 5.4. Since learning curves are not taken into account in this case, it is assumed that the maximum adaptive value is achieved immediately after birth. In the figure, the light-gray triangle would become a rectangle.

5.3.2 Modeling Learning Curves

In order to analyze the influence of the curvature of learning curves *ceteris paribus*, three functions $v_0(x)$, $v_1(x)$, and $h(t)$ need to be defined. $v_0(x)$ specifies the innate adaptive value

of an individual with genotype x , $v_1(x)$ specifies the adaptive value of x at the end of its life, and $h(t)$ specifies the curvature of the learning curve. $h(t)$ is limited to functions that are monotonic in t in the interval $t \in [0, 1]$.

Based on these definitions, the adaptive value function is defined as

$$v(x, t) = \frac{h(t) - h(0)}{h(1) - h(0)}(v_1(x) - v_0(x)) + v_0(x), \quad (5.29)$$

i.e., all individuals' learning curves have the same curvature, and $v(x, 0) = v_0(x)$ respectively $v(x, 1) = v_1(x)$.

5.3.3 Genotype-Independent Learning Curves

In the following, it is assumed that the genotype has no influence on the curvature of the learning curves, i.e., h does not depend on genotype x .

Posthumous versus Continual Fitness Assessment

If learning curves are not taken into account (posthumous fitness assessment) the gain function is given by

$$g(x) = \frac{v_1(x)}{v_0(x)}. \quad (5.30)$$

In the following, the gain function that accounts for the learning curves is denoted as $\tilde{g}(x)$. Notice that this is the simple gain function as introduced in Section 4.3.

$$\begin{aligned} \tilde{g}(x) &= \frac{\int_{t=0}^{t=1} v(x, t) dt}{v_0(x)} \\ &= \frac{\int_{t=0}^{t=1} \left(\frac{h(t) - h(0)}{h(1) - h(0)}(v_1(x) - v_0(x)) + v_0(x) \right) dt}{v_0(x)} \\ &= H \frac{v_1(x)}{v_0(x)} - H + 1 \\ &= Hg(x) - H + 1, \end{aligned} \quad (5.31)$$

where

$$H = \int_{t=0}^{t=1} \frac{h(t) - h(0)}{h(1) - h(0)} dt. \quad (5.32)$$

Straightforwardly,

$$\tilde{g}'(x) = Hg'(x), \quad (5.33)$$

and since $H > 0$,

$$\text{sign}(\tilde{g}'(x)) = \text{sign}(g'(x)). \quad (5.34)$$

In conclusion, there is no *qualitatively* different influence between posthumous and continual fitness assessment if the learning curves h do not depend on the genotype, i.e., if learning accelerates (decelerates) evolution with posthumous fitness assessment it also accelerates (decelerates) evolution with continual fitness assessment. However, the magnitude of acceleration or deceleration may differ for different curvatures.

Influence of a Curvature Change of the Learning Curves on Evolution

Although continual fitness assessment with genotype-independent learning curves does not qualitatively change the influence of learning on evolution, it may influence the magnitude of acceleration or deceleration that would be present with posthumous fitness assessment. In the following, the extended gain function framework of Section 4.4 is applied to study the influence of the curvature *change* of the learning curves on evolution.

First, the function that describes the curvature of the learning curve $h(t)$ is extended to $h(t, a)$, where a is a learning parameter that influences the curvature. Correspondingly, the fitness of an individual with genotype x and learning curve parameter a , is

$$\begin{aligned}
 f_\phi(x, a) &= \int_{t=0}^{t=1} v(x, a, t) dt \\
 &= \int_{t=0}^{t=1} \left(\frac{h(t, a) - h(0, a)}{h(1, a) - h(0, a)} (v_1(x) - v_0(x)) + v_0(x) \right) dt \\
 &= (v_1(x) - v_0(x))H(x, a) + v_0(x) \\
 &= v_1(x)H(a) + v_0(x)(1 - H(a)),
 \end{aligned} \tag{5.35}$$

where $H(a)$ is substituted,

$$H(a) = \int_{t=0}^{t=1} \frac{h(t, a) - h(0, a)}{h(1, a) - h(0, a)} dt. \tag{5.36}$$

The gain function derivative of the extended gain function framework can be reformulated

$$\begin{aligned}
 \frac{\partial^2}{\partial x \partial a} \log f_\phi(x, a) &= \frac{\partial^2}{\partial x \partial a} \log (v_1(x)H(a) + v_0(x)(1 - H(a))) \\
 &= \frac{\partial}{\partial x} \left(\frac{v_1(x)H'(a) - v_0(x)H'(a)}{v_1(x)H(a) + v_0(x)(1 - H(a))} \right) \\
 &= (v_1'(x)v_0(x) - v_1(x)v_0'(x))H'(a) \\
 &= (v_0(x))^2 \left(\frac{v_1(x)}{v_0(x)} \right)' H'(a) \\
 &= (v_0(x))^2 g'(x)H'(a),
 \end{aligned} \tag{5.37}$$

where $v_0'(x) = \partial v_0 / \partial x$, $v_1'(x) = \partial v_1 / \partial x$ and $H'(a) = \partial H / \partial a$. Thus,

$$\text{sign} \left(\frac{\partial^2}{\partial x \partial a} \log f_\phi(x, a) \right) = \text{sign}(g'(x)H'(a)), \tag{5.38}$$

Recall that the sign of $g'(x)$ indicates acceleration and deceleration of evolution in case of posthumous fitness assessment. If $H'(a) > 0$ and in case of $g'(x) > 0$, increasing learning parameter a accelerates evolution. If $H'(a) > 0$ and in case of $g'(x) < 0$, increasing learning parameter a decelerates evolution. So the influence of a learning curve parameter a on evolution is determined by the derivative of the integral of the normalized learning curve, $\int_{t=0}^{t=1} \frac{h(t, a) - h(0, a)}{h(1, a) - h(0, a)} dt$, w.r.t. a .

In the following example, learning parameter a determines the degree of convexity (concavity) of the learning curve, in particular,

$$h(t, a) = t^a . \quad (5.39)$$

With small a , h is increasing strongly for small t values which can be interpreted as “early learning”. In contrast, with large a , h is increasing strongly for large t values which can be interpreted as “late learning”. Notice that in the curvature of Equation 5.39, the limit of $a = 0$ corresponds to posthumous fitness assessment (everything is learned immediately), and the limit $a = \infty$ corresponds to the complete absence of learning.

Since $h(0, a) = 0$ and $h(1, a) = 1$ for all $a > 0$,

$$H'(a) = \frac{\partial}{\partial a} \int_0^1 t^a dt = \frac{\partial}{\partial a} (a + 1)^{-1} = -(a + 1)^{-2} . \quad (5.40)$$

Recalling Equation 5.38, one obtains

$$\text{sign} \left(\frac{\partial^2}{\partial x \partial a} \log f_\phi(x, a) \right) = \text{sign}(g'(x) - (a + 1)^{-2}) = \text{sign}(-g'(x)) . \quad (5.41)$$

Consider first the case of $g'(x) > 0$, i.e., where evolution is accelerated with posthumous fitness assessment. Increasing a (later learning) works to reduce the rate of evolution. In an analogous manner, consider $g'(x) < 0$, i.e., with posthumous fitness assessment, evolution is decelerated. Increasing a (later learning) works to increase the rate of evolution.

5.3.4 Genotype-Dependent Learning Curves

Now the case where the learning curves depend on the genotype value is considered. In this case, the curvature may not only influence the magnitude of acceleration and deceleration but even reverse the sign of the influence. An example of a learning curve that depends on the genotype is given as follows

$$h(t, a, x) = t^{a^{2x-1}} , \quad a > 0 , \quad x \in [0, 1] . \quad (5.42)$$

This learning curve is combined with

$$v_0(x) = x , \quad v_1(x) = 3x , \quad x \in [0, 1] , \quad (5.43)$$

which in case of posthumous fitness assessment produces a constant gain function of $g(x) = v_1(x)/v_0(x) = 3$ and thus learning with posthumous fitness assessment has no influence on evolution.

Figure 5.5 illustrates the corresponding (extended) adaptive value landscape for $a = 0.25$ and $a = 4$. For $a = 0.25$ the learning curves are convex for small x and concave for large x . Genetically weak individuals learn late while genetically strong individuals are early learners. The opposite is true for $a = 4$. Genetically weak individuals are early learners while genetically strong individuals learn late.

Figure 5.6 shows the corresponding gain functions which have been derived numerically. With $a = 0.25$ the genotype-dependent curvature of the learning curves causes a monotonically




Figure 5.5: Scenario with varying curvature of the learning curves. With learning parameter $a = 0.25$ (left panel) genetically weak individuals have a concave learning curve (“late learners”) and genetically strong ones have a convex learning curve (“early learners”), and vice versa with learning parameter $a = 4$ (right panel). Evolution is accelerated through learning with $a = 0.25$ and decelerated with $a = 4$.

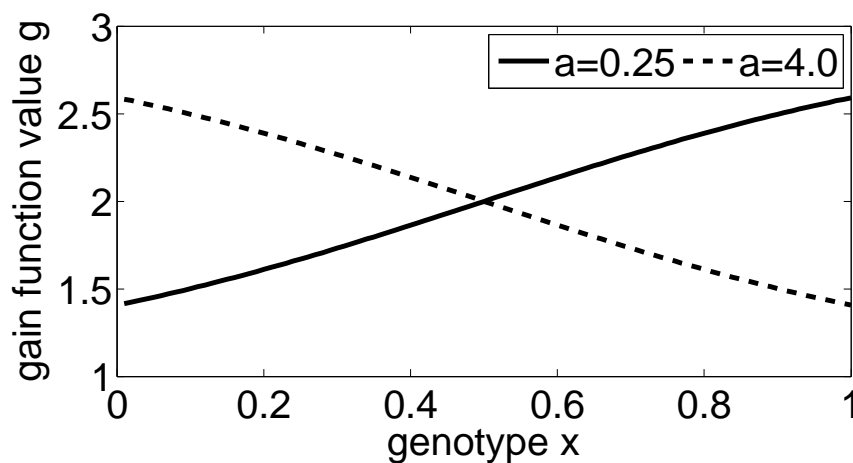


Figure 5.6: Gain functions corresponding to Figure 5.5. With $a = 0.25$ (left panel of Figure 5.5) the curvature of the learning curves causes a monotonically increasing gain function and therefore acceleration, while with $a = 4$ (right panel of Figure 5.5) the curvature causes a monotonically decreasing gain function and therefore deceleration.




Figure 5.7: *Simulation results corresponding to Figure 5.5. The simulations confirm the predictions of the gain function analysis (cf. Figure 5.6). Compared to posthumous fitness assessment, evolution is accelerated with continual fitness assessment and $a = 0.25$ and decelerated with continual fitness assessment and $a = 4$.*

increasing gain function, in contrast the gain function that corresponds to $a = 0.25$ causes monotonically decreasing gain function. Thus, with $a = 0.25$ the curvature of the learning curves accelerates evolution while in case of $a = 4$ the curvature of the learning curves decelerates evolution although with posthumous fitness assessment learning has no influence on evolution.

Additionally a simulation study with a population of 100 individuals and a Gaussian mutation with $\sigma = 10^{-4}$ is carried out. The remaining parameters of the experiment are set as in the simulation study of Section 5.1.3. The results as shown in Figure 5.7 confirm the predictions of the gain function analysis (cf. Figure 5.6): Compared to posthumous fitness assessment, evolution is accelerated with continual fitness assessment and $a = 0.25$ and decelerated with continual fitness assessment and $a = 4$.

5.3.5 Conclusion

Posthumous fitness assessment refers to the case when only the result of a learning process is taken as the basis for selection whereas continual fitness assessment refers to the scenario when learning curves are also taken into account. The gain function analysis has shown that genotype-independent learning curves only have an influence on the magnitude of learning-induced acceleration and deceleration, but not on the sign of the influence (compared to posthumous fitness assessment). However, genotype-dependent learning curves may also influence the sign of the influence, i.e., even if learning has no influence on evolution with posthumous fitness assessment, continual fitness assessment may cause acceleration or deceleration.

5.4 A Non-Monotonic Gain Function

An exact prediction of the population dynamics with the gain function analysis requires that the gain function is monotonic within the range of the population. In the following, it will be demonstrated that the gain function can also be used as an approximate predictor of the population dynamics even if the gain function is not monotonic.

5.4.1 Fitness, Learning and Gain Functions

Fitness Function

As a mapping from phenotype to fitness the sigmoid function

$$f(z) = \frac{1}{1 + e^{-z}} \quad (5.44)$$

is employed in the section (visualized in Figure 5.8(a)). This chapter concentrates on the influence of learning. Therefore development ϕ is defined as the identity function, i.e.,

$$z = \phi(x) = x . \quad (5.45)$$

The sigmoid function is convex for negative z values and concave for positive ones, and is monotonically increasing towards the asymptote $f = 1$.

Two types of directional learning (cf. Section 5.1) are applied to this function, namely, *constant directional learning* and *progressive directional learning*, producing different population dynamics.

Constant Directional Learning

Constant directional learning is the same type of learning as defined in Section 5.1.1, i.e.,

$$l_1(z) = z + \delta , \quad (5.46)$$

where in this section $\delta = 0.25$. Since we are not interested in the influence of the δ -value on evolutionary dynamics any other setting would be appropriate as well.

For this type of learning it has already been shown that for positive δ the gain function derivative is

$$\text{sign}(g'(x)) = \text{sign}((\log(f(x)))'') . \quad (5.47)$$

Applying Equation 5.47 to the sigmoid fitness one obtains

$$(\log(f(x)))'' = -e^{-x}(1 + e^{-x})^{-2} < 0 , \forall z . \quad (5.48)$$

Thus, it is expected that constant directional learning decelerates evolution on the sigmoid fitness function. The corresponding gain function is shown as solid line in Figure 5.8(b).

Progressive Directional Learning

Progressive directional learning is defined as

$$l_2(x) = x + e^x . \quad (5.49)$$

With this type of learning, individuals with larger x (genetically stronger individuals) learn more than genetically weak ones. In combination with the sigmoid fitness function, directional progressive learning produces a gain function as shown as dashed line in Figure 5.8(b). The gain function is non-monotonic with a maximum at genotype value $x = 0.14$. Thus, if a population is entirely located left of $x = 0.14$ progressive directional learning is predicted to accelerate evolution and in the case that it is entirely located right of $x = 0.14$ progressive directional learning is predicted to decelerate evolution. If, however, the gain function is non-monotonic *within* the range of the population it does no longer allow a precise prediction of the dynamics. However, a simulation study will demonstrate that the gain function is still useful to approximately describe the population dynamics. Thus, in addition to the gain function analysis, the sigmoid landscape (coupled with the two types of directional learning), will be studied empirically based on repeated computer simulations. All simulation experiments have been set up as described in the following.

5.4.2 Simulation

A population of 100 individuals, each characterized by a one-dimensional (real-valued) genotypic value x , evolves asexually, i.e., evolution is modeled as a cycle of mutation and selection. Linear fitness proportional selection (with respect to the sigmoid function value of the learned phenotype z) is simulated with the Stochastic Universal Sampling algorithm [6]³. To simulate mutation, a random number $X_{\phi(\mu,\sigma)}$ drawn from a normal distribution with parameters $\mu = 0$ and $\sigma = 10^{-3}$ was added to the genotypic value x of each offspring. In all simulations, the population's genotypes are initialized uniformly in the vicinity of -3 (in the interval $[-3.1, -2.9]$). For each setting 1000 independent simulation runs have been carried out. The presented results are averaged data over these runs.

5.4.3 Results

Constant Directional Learning

The simulation results for the case of constant directional learning are shown in Figures 5.8(c-d). Figure 5.8(c) shows the average trajectory of mean genotype evolution (\bar{x} in absence, \bar{x}_l in presence of learning) with *constant directional learning* and *no learning*. Figure 5.8(d) shows the average trajectory of mean genotype in case of learning, normalized by the mean genotype value in absence of learning, $\bar{x}_l - \bar{x}$. This normalized value is denoted *learning lead*. As predicted by the negative gain function derivative, constant directional learning decelerates evolution.

³ This algorithm implements sampling (with replacement) of n offspring from n parents, where the probability of an individual being sampled is proportional to its fitness $f(z)$, i.e., $f(x)$ without learning, $f(l(x))$ with learning).

Progressive Directional Learning

The simulation results for the case of progressive directional learning are shown in Figures 5.8(e-f). Figure 5.8(e) shows the average trajectory of mean genotype evolution and Figure 5.8(f) the normalized mean genotype value (cf. Figure 5.8(d)). Recall (Figure 5.8(b)) that first learning-induced acceleration and then deceleration is expected. These predicted dynamics are qualitatively confirmed by the simulation results.

The mean genotype of the learning population reaches the genotype that corresponds to the gain function maximum ($x = 0.14$) in generation 184. The maximum difference between learning and non-learning population has been reached already 25 generations earlier. However, during these 25 generations, the learning population has largely maintained its distance to the non-learning population.

5.4.4 Conclusion

The gain function analysis only allows an approximate prediction of the population dynamics over time. An exact prediction based on the gain function assumes that both learning and non-learning population have the same distribution in genotype space. In the example of progressive directional learning, the learning population moves quicker toward higher genotype values than the non-learning one, during the early phase of evolution. Thus, the learning individuals populate a different region in genotype space than the non-learning ones. Despite a positive gain function derivative the selection pressure might be stronger in the region of the non-learning population than in the region of the learning population.

Nevertheless, the evolutionary dynamics are quite well described by the gain function, as the example has demonstrated. In conclusion, the gain function approach can approximately predict the evolutionary dynamics even in the case where acceleration is followed by deceleration.

5.5 Summary and Conclusion

This chapter has demonstrated the generality of the gain function framework. It has been used here to deepen the understanding of the dynamics when evolution is coupled with learning.

In Section 5.1, two basic components of learning, namely a directional and a noise component, have been identified that are common to most learning algorithms or procedures, and general conditions have been derived under which these components accelerate respectively decelerate evolution. Directional learning accelerates evolution if the logarithm of the function that maps phenotype to fitness is convex and decelerates it if the logarithm is concave. It turned out that noise in the genotype-phenotype-mapping can actually accelerate the evolutionary process which is a somewhat non-intuitive result.

Then, in Section 5.2, the special case in which the fitness of a learning individual is additively composed of an innate component and a learning component has been analyzed. If learning produces a positive fitness gain (is not maladaptive) evolution is only accelerated through learning in the case where the fitness of the learning-dependent component increases stronger in x than the fitness of the innate component.

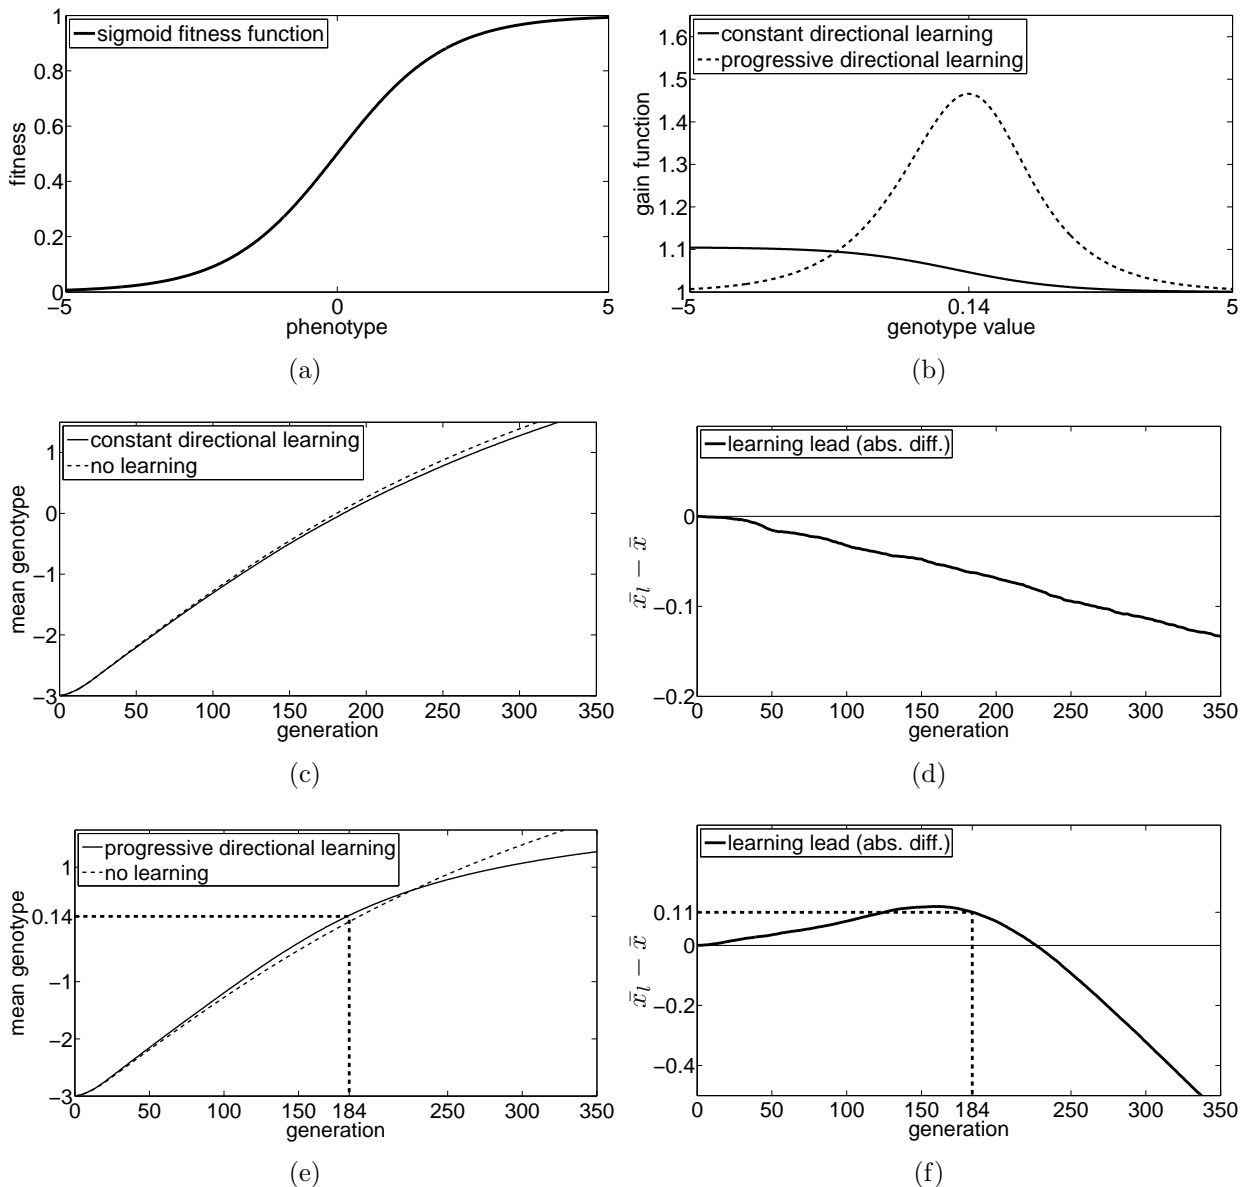


Figure 5.8: Evolution and learning on the sigmoid fitness function. (a): The sigmoid fitness function, (b): gain functions for constant directional learning and progressive directional learning, (c-f): averaged results of 1000 independent simulation runs with the sigmoid fitness function, in particular (c): mean genotype evolution with constant directional learning and no learning, (d): absolute difference of the curves in (c), i.e., mean genotype in case of learning and in the absence of learning, $\bar{x}_l - \bar{x}$, which is named “learning lead”, (e): same as (c) but with progressive directional learning, (f): same as (d) but with progressive directional learning.

Next, in Section 5.3, it has been investigated how fitness assessment of a learning individual influences evolution. In particular, learning curves have been modeled that describe the progress of a learning individual over its lifetime and it was compared how a continual fitness assessment during lifetime differs in its influence on evolution from a fitness assessment where only the result of learning is taken into account (posthumous fitness assessment). It was found with a gain function analysis that in the case that learning curves are genotype-independent, the curvature of these curves only has an influence on the magnitude of learning-induced acceleration and deceleration, but not on the sign of the influence (in relation to posthumous fitness assessment). However, genotype-dependent learning curves may also influence the sign of the influence, i.e., even if learning has no influence on evolution with posthumous fitness assessment, continual fitness assessment may cause acceleration or deceleration.

Finally, in Section 5.4, it has been demonstrated that the gain function can well describe the population dynamics even if it is not monotonic (which is one of the assumptions of the mathematical basis of the gain function).

The results of this section may not only be helpful for the design of optimization algorithms that couple evolution and learning. Furthermore they may shed some light on the results obtained by simulation studies in the field of artificial life and computational biology, or even real biological experiments, and provide a theoretical underpinning of some of the derived conclusions.

In the following chapter examples are presented that demonstrate how such a theoretical underpinning can be derived.

Gain Function Analysis of Other Models of Evolution and Learning

In this chapter, several models from the literature that investigate the influence of learning on evolution are revisited and analyzed with the gain function framework in order to derive a theoretical underpinning of the respective conclusions. Large parts of this chapter are based on [130].

6.1 Hinton and Nowlan’s In Silico Experiment

In 1987, the seminal paper of Hinton and Nowlan [64] presented the first computational model demonstrating the Baldwin effect (it will be referred to as H&N model). In this model, Hinton and Nowlan show “how learning can guide evolution” towards a global optimum, thereby giving an example of learning-induced acceleration of evolution.

6.1.1 Original Model Formulation

In the H&N model, an individual’s genotype x is represented by 20 gene loci (elements) with alleles ‘0’, ‘1’ and ‘?’, i.e.,

$$x \in \{0, 1, ?\}^{20}. \quad (6.1)$$

that is mapped to the phenotype z , which is represented by a bit-string of length 20, i.e.,

$$z \in \{0, 1\}^{20}. \quad (6.2)$$

Hinton and Nowlan suggest to interpret the phenotype as the synapse weight specification of a neural network. In the phenotype space of size $2^{20} = 1048576$, there exists exactly one phenotype with the correct specification. In the words of Hinton and Nowlan this phenotype (the “good net”) is like a “needle in a haystack”. Hinton and Nowlan do not specify the optimal phenotype but without loss of generality it will be assumed in this section that it is the “all ones” phenotype, i.e., $z^* = 11111111111111111111$.

In the mapping from genotype x to innate phenotype z , each phenotype element z_i corresponds to a locus of the gene x_i , and is defined as

$$\forall i : z_i(x_i) = \begin{cases} x_i & \text{if } x_i \in \{0, 1\} \\ X_{\{0,1\}} & \text{otherwise,} \end{cases} \quad (6.3)$$

where $X_{\{0,1\}}$ is a random number sampled from a Bernoulli distribution with $p = 0.5$, i.e., 0 and 1 are equally likely to be drawn. If the generated phenotype matches the “all ones” phenotype the individual is assigned an absolute fitness of $f = 20$. If it does not match the optimal phenotype the individual starts guessing for it (formally, Equation 6.3 is repeatedly executed). The individual stops guessing after 1000 trials or if it has found the optimal phenotype. The individual “guessing” process is interpreted as individual learning.

The fitness of an individual with genotype x is given by

$$f(n_x) = 1 + \frac{19n_x}{1000}, \quad (6.4)$$

where n_x is the number of remaining trials after the optimum has been found. Thus, the earlier the good net is found, the higher the fitness.

Hinton and Nowlan do not completely specify the evolutionary algorithm used in their simulation but mention that they employ a version of the genetic algorithm proposed by Holland [66] with a population size of 1000, a fitness proportional selection scheme and the 1-point crossover [66]. In the original paper, mutation is not mentioned and thus it must be assumed that no mutation has been simulated. The same interpretation can be found in the secondary literature, as for instance in [138].

Hinton and Nowlan present a figure (Figure 2 in their article) which shows the trajectories of the simulated evolution (it is unclear if the figure shows averaged data or the data of a single run). A reimplementation of the H&N model produced similar evolutionary trajectories that are presented in Figure 6.1. The *number of incorrect alleles* is the average number of '0's, the *number of correct alleles* is the average number of '1's, and the *number of undecided alleles* is the average number of '?'s in the genotype. It can be seen that the number of incorrect alleles is decreasing quickly and at the same time the number of correct alleles is increasing. Then, the number of undecided alleles decreases further, i.e., undecided alleles are replaced by correct ones. However, a certain fraction of undecided alleles is not replaced even if the evolution is run for several thousand generations (not shown).

According to [138] and [59] one reason for the “persistent question-marks” [59] is that in the original H&N simulation no mutation is used: Once the variation in one locus of the gene is lost through genetic drift there is no way to change this locus.

However, even if mutation would be used there is a reason why on average some '?' loci would persist. With mutation the population reaches a stable state (equilibrium) after a finite number of generations. Mutation keeps introducing non-optimal genotypes to the population and selection works to remove them. Thus, on average there will be a certain fraction of non-optimal genotypes with '?' or '0' loci. (This phenomenon - the formation of a population at so-called mutation-selection balance - is well known as *quasi-species* in biology [37, 38]. The quasi-species is not the central issue of this chapter and it is referred to Chapter 7 where this concept is discussed in more detail.)




Figure 6.1: Simulation with a reimplementation of the H&N model. Evolutionary trajectories are averaged over 100 independent simulation runs and are very similar to the results presented in Figure 2 of the original article [64].

Hinton and Nowlan state that “the same problem was never solved by an evolutionary search without learning.” It is not explicitly stated how the simulation was done without learning but it is reasonable to assume that simply the ‘?’ has been removed from the set of alleles and that the genotype-phenotype mapping is

$$\forall i : z_i(x_i) = x_i . \quad (6.5)$$

Although not explicitly stated in [64] it is assumed here that the optimal genotype in absence of learning is assigned a fitness of 20 and all non-optimal genotypes are assigned a fitness of 1.

With this setting, and in the absence of mutation it is indeed very unlikely that the population “accidentally” discovers the optimal phenotype (the needle in the haystack) before genetic drift has removed the genetic variation. However, even if some mutation is added the time until the needle is found is extremely long, and even if the optimum was found it is likely to be lost due to the disruptive cross-over effect.

Hinton and Nowlan presented the first computational example of the Baldwin effect. Since it takes an extremely long time until the optimum has been populated in absence of learning and only few generations in presence of learning it can be argued that learning accelerates evolution here.

In the following, the gain function framework is applied to the H&N model in order to produce an analytical argument for the observed learning-induced acceleration of evolution. Before the gain function is applied, however, a reformulation of the original model is required.

6.1.2 Model Reformulation

In the original model, genotype x is defined as $x \in \{0, 1, ?\}^{20}$ in the case of learning and $x \in \{0, 1\}^{20}$ in the absence of learning. However, the gain function analysis requires that genotype and phenotype have the same representation and that learning can be “added”. To achieve this the H&N model is reformulated.

In the reformulated model, a genotype is now defined as

$$x \in \{0, 1, ?_0, ?_1\}^{20} . \quad (6.6)$$

In brief, alleles '0' and '1' encode the phenotype directly, whereas alleles '?₀' and '?₁' map either to '0' or '1' after a learning period, but learning starts at 0 in case of '?₀' and at 1 in case of '?₁'. Formally, the mapping from genotype to innate phenotype z_0 (development) is defined as

$$\forall i : z_{0i}(x_i) = \begin{cases} 0 & \text{if } x_i \in \{0, ?_0\} \\ 1 & \text{otherwise} , \end{cases} \quad (6.7)$$

and the phenotype changes according to

$$\forall i : z_i(x_i) = \begin{cases} x_i & \text{if } x_i \in \{0, 1\} \\ X_{\{0,1\}} & \text{otherwise} . \end{cases} \quad (6.8)$$

The difference between learning and non-learning individuals in the reformulation of the model is that learning individuals are allowed to perform 1000 random guesses, whereas for non-learning individuals the genotype translates directly to the phenotype and no further improvement is possible. This modification does not substantially change the H&N model and produces the same evolutionary dynamics as the original formulation, cf. Figure 6.2. The gain function framework can now be applied to the reformulated model.

6.1.3 Gain Function Analysis

To apply the gain function (of the basic framework of Section 4.3) three classes of genotypes are distinguished.

First, if there exists one or more '0' alleles in the genotype, the optimal phenotype will not be found in either case, with or without learning. This means the gain function is a constant equal to one.

Second, if the genotype is composed of alleles '1' and '?₁', the optimal phenotype will be generated in both cases with or without learning, which also implies a constant gain function of one.

Thus, in both cases, the gain function is a constant and learning has no influence on evolution.

In the third case, the genotype is composed of at least one locus with '1' or '?₁' allele, at least one locus with '?₀' allele and no locus with '0' allele. In the following, it will be shown that the gain function is increasing toward the optimum in this case:

If q denotes the number of question marks of an individual's genotypes (sum of '?₀' and '?₁' loci) and all other loci are '1', the expected absolute fitness in case of learning can be

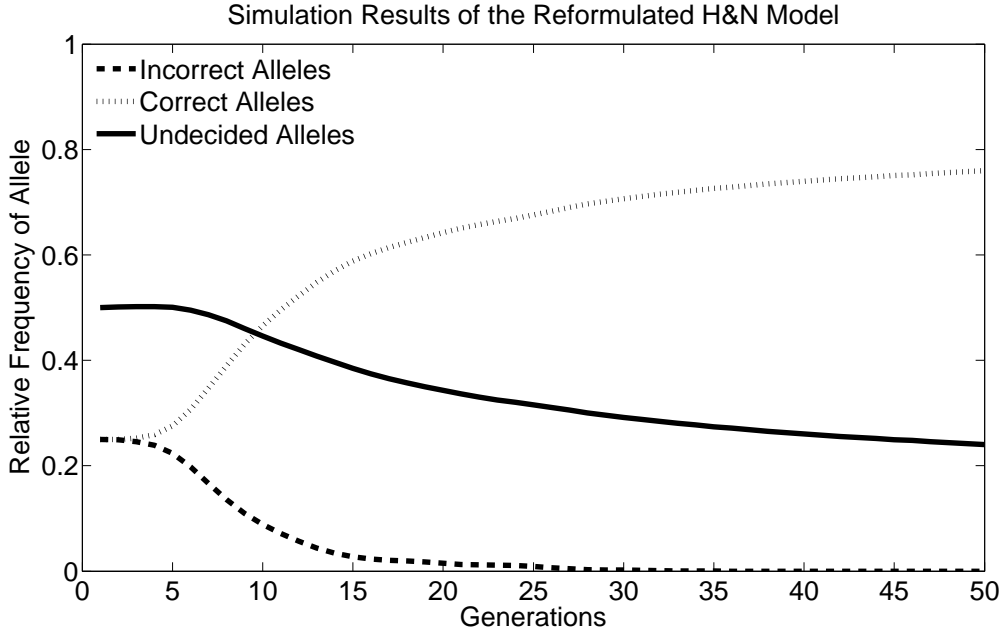


Figure 6.2: Simulation with the reformulation of the H&N model. The comparison to the evolutionary dynamics in Figure 6.1 (where the simulation parameters are identical) shows that the original model and the reformulation are equivalent. The number of undecided alleles is the sum of η_0 and η_1 alleles.

calculated as follows: The probability of guessing the all-ones phenotype in one trial is 2^{-q} and the probability to guess it exactly at the k 'th guess is

$$p(k, q) = (1 - 2^{-q})^{(k-1)} \cdot 2^{-q} . \quad (6.9)$$

Thus, the expected fitness \bar{f}_l of a learning individual with q question marks and at most 1000 learning trials is

$$\bar{f}_l(q) = \sum_{k=1}^{1000} p(k, q) f(1000 - k) + (1 - 2^{-q})^{1000} f(0) . \quad (6.10)$$

f is defined as in Equation 6.4.

In the absence of learning, it is impossible for an individual of the third category to find the optimum (since there is at least one η_0 allele) and according to Equation 6.4,

$$\bar{f}(q) = 1 . \quad (6.11)$$

Based on this the gain function of Section 4.3 can be formulated as

$$g(q) = \frac{\bar{f}_l(q)}{\bar{f}(q)} = \bar{f}_l(q) . \quad (6.12)$$

Figure 6.3 shows this gain function $g(q)$ plotted against reversely ordered q and the corresponding differential $g(q-1) - g(q)$. The gain function is increasing (its differential is positive) towards the fitness optimal genotype with $q = 0$, thus the gain function analysis confirms the simulation results of Hinton and Nowlan [64].




Figure 6.3: Gain function and gain function differential in the H&N model. The increasing gain function (positive differential) towards the fitness optimal genotype (no '?' alleles) indicates learning-induced acceleration of evolution in the model of Hinton and Nowlan [64].

6.1.4 Discussion

In the literature, several papers have commented on Hinton and Nowlan's results, however, a selection pressure argument is sufficient to explain Hinton and Nowlan's result. In the H&N model, individuals have a genetic predisposition toward the optimal phenotype. In the absence of learning, these differences between genetic predispositions are invisible for selection. Learning amplifies or actually unveils these differences. A learning induced amplification of genetic predispositions is exactly the conclusion that follows from a positive gain function derivative.

More generally, it is expected that in extreme fitness landscapes with large plateaus learning accelerates evolution.

6.2 Papaj's In Silico Experiment of Insect Learning

In biology, computer simulations of evolution are often used as a research tool to support evolutionary theory. An example of this is Papaj's simulation of evolution and learning in insects which he presents in the first part of [133]. Based on an earlier work [75] Papaj describes a scenario in which the environment of a population of insects suddenly changes such that only one host species is available. An insect behavior is only to a certain extent genetically specified and partly plastic. Hence, to what extent an insect is able to exploit this host species, depends on both its genetic configuration and the ability to learn.

The result of Papaj's simulations are that learning inhibits the evolution of genetically (innately) strong individuals. In the following, a gain function analysis of Papaj's simulation model is done in order to get a better theoretical understanding of his results.

Papaj points out that the arguments derived from this model should apply more generally and equally well to other kinds of behaviors. Indeed, as will be shown in the following, the model is formulated quite generally.

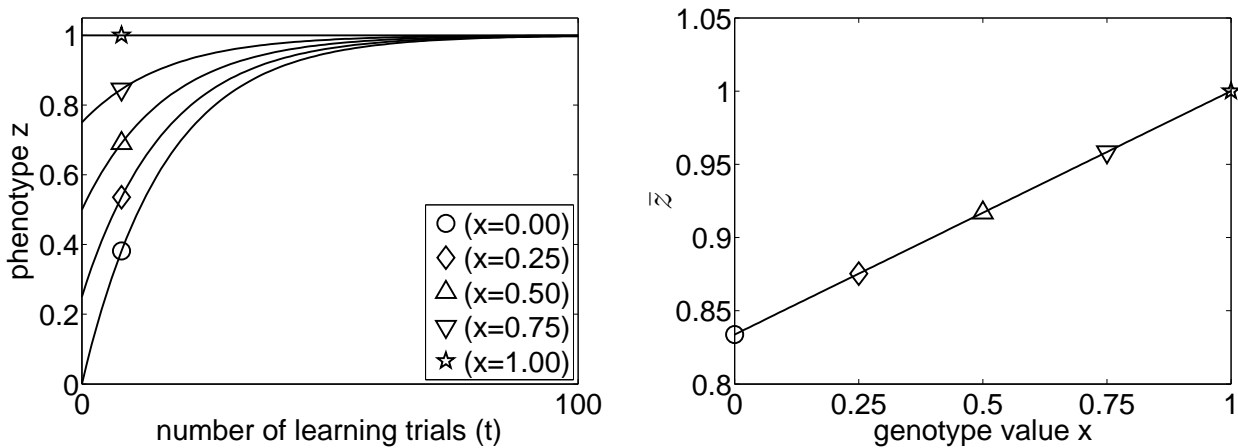


Figure 6.4: Phenotype change over lifetime in Papaj's model of evolution and learning in insects [133]. The left panel shows learning curves for a learning parameter $L = 0.06$ and different genotypic values (equals innate phenotype) $x \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$, cf. Equation 6.13. All individuals have a strong progress in learning. Those with higher genotypic values have a better starting position to reach the learning target, but the “genetically weak” ones “catch up” during learning. In the right panel, the average phenotype over $T = 100$ learning trials with learning parameter $L = 0.06$ is shown, as calculated using Equation 6.15.

6.2.1 Model Formulation

An insect's behavior (the phenotype) is represented by a real-valued response number $z \in [0, 1]$. The innate behavior is directly encoded as genotypic value $x \in [0, 1]$. Learning depends on two parameters $a = (L, T)$. T is the duration of learning (lifetime of an insect). The behavioral change over time t ($t = 0..T$) is influenced by learning parameter $L \in \mathbb{R}_0^+$ (in [133], $L \in [0, 0.1]$). Thus, the phenotype at a time depends on t , x and L , and is specified as

$$z(x, L, t) = x + (1 - x)(1 - e^{-Lt}) = 1 + (x - 1)e^{-Lt}, \quad (6.13)$$

which can be interpreted as a learning curve. Notice, however that this is not the same type of learning curves as in Section 5.3. The learning curves of Section 5.3 were defined as the mapping from time to adaptive value and not (as in this section) as the mapping from time to phenotype.

Equation 6.13 is visualized in the left panel of Figure 6.4 for $a = 0.06$, for five different genotypic values x . Presumably, Papaj chose this type of learning curve because it guarantees that insect behavior at birth is solely specified by the genotype, i.e., $z(x, L, 0) = x$, and because in the T consecutive learning trials z converges asymptotically toward the optimal phenotype $z = 1$, which is a typical animal learning curve according to [133]. All individuals have a strong progress in learning, those with higher genotypic values have a better starting position to reach the learning target quicker, but the genetically weak ones “catch up” during learning. Fitness in Papaj's experiment is determined by a function f that is applied to the average phenotype of an individual's lifetime, in particular

$$f(\bar{z}) = 1 - (1 - \bar{z})^2 \quad (6.14)$$




Figure 6.5: Definition of fitness in Papaj's model [133], a concave function defined on the mean individual phenotype, cf. Equation 6.14.

which is an inverted parabola with maximum at $\bar{z} = 1$, i.e., a concave function on $\bar{z} \in [0, 1]$, cf. Figure 6.5 where one half of the parabola is shown. Thus, the optimal behavior is achieved with $z = 1$ and the optimal fitness with $\bar{z} = 1$ respectively. Notice that an alternative (perhaps more intuitive) definition would have been to define an adaptive value function $v(z(x, L, T))$ and to measure fitness as the integral over this function in the limits of its lifetime. This approach is equivalent to the one taken in Section 5.3 of this thesis. Nevertheless, since in Papaj's formulation individual lifetime changes are taken into account, his approach to determine fitness can be considered as a *type* of continual fitness assessment (cf. Section 5.3).

6.2.2 Gain Function Analysis

In order to calculate the expected fitness in presence and in absence of learning, the average phenotype value \bar{z} needs to be calculated. Since Papaj uses a discrete time model an exact calculation would involve taking the sum over the different phenotype values of an individual's lifetime. For the sake of simple analysis this sum is approximated with the corresponding integral, i.e.,

$$\bar{z}(x, L, T) = \begin{cases} x & , \text{if } T = 0 \\ \frac{1}{T} \int_{t=0}^T z(x, L, t) dt = 1 + \frac{1-x}{LT} (e^{-LT} - 1) & , \text{if } T > 0 \end{cases} \quad (6.15)$$

The resulting average phenotype (for $T = 100$ and $L = 0.06$) is shown in the right panel of Figure 6.4.

With equations 6.14 and 6.15 the expected fitness of an individual with genotype x , learning parameter L and lifetime $T > 0$ is in the presence of learning given by

$$f_{\phi_l}(x, L, T) = f(\bar{z}(x, L, T)) = 1 - \left((x - 1) \frac{e^{-LT} - 1}{LT} \right)^2, \quad (6.16)$$

and in the absence of learning simply

$$f_{\phi}(x) = f(x). \quad (6.17)$$

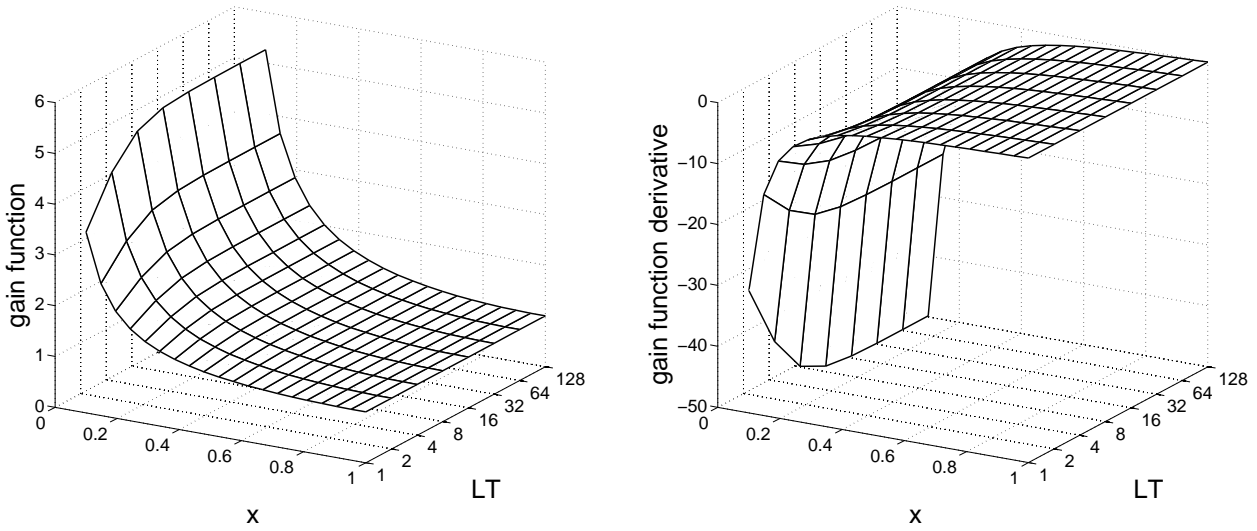


Figure 6.6: Basic gain function of Papaj's experiment [133]. The left panel shows the gain function $g(x)$ plotted against genotypic value x for different values of the product of lifetime and learning parameter LT (logarithmic scale), cf. Equation 6.18. The right panel shows its derivative with respect to x , cf. Equation 6.19. For all possible parameter combinations LT , the gain function is negatively sloped toward the optimum at $x = 1$, which corresponds to a negative gain function derivative.

Thus, the basic gain function is derived as

$$g(x) = \frac{f_{\phi_l}(x, L, T)}{f_{\phi}(x)} = \frac{1 - \left((x-1) \frac{e^{-LT} - 1}{LT} \right)^2}{1 - (1-x)^2}. \quad (6.18)$$

After differentiation with respect to x and some straight-forward calculations the gain function derivative can be formulated as

$$g'(x) = \frac{2(1-C)}{(x^2 - 2x)^2} (x-1) \quad (6.19)$$

with $C = \left(\frac{e^{-LT} - 1}{LT} \right)^2$.

Since $L > 0$ and $T \geq 0$, the product $LT \geq 0$ can be interpreted as one variable. Since $C \in]0, 1[$ for $LT > 0$, one can see that $g'(x) < 0$ for all $x \in]0, 1[$. The gain function (Equation 6.18) and its derivative (Equation 6.19) are visualized in Figure 6.6. For all parameter combinations LT , the gain function is negatively sloped toward the optimum at $x = 1$, which corresponds to a negative gain function derivative. Parameter combinations for small values of LT and x are omitted to avoid numerical difficulties since the gain function is not defined for $x = 0$ and $LT = 0$.

The negative gain function derivative supports and explains Papaj's simulation results. Learning indeed suppresses the evolution of genetic predisposition toward high fitness.

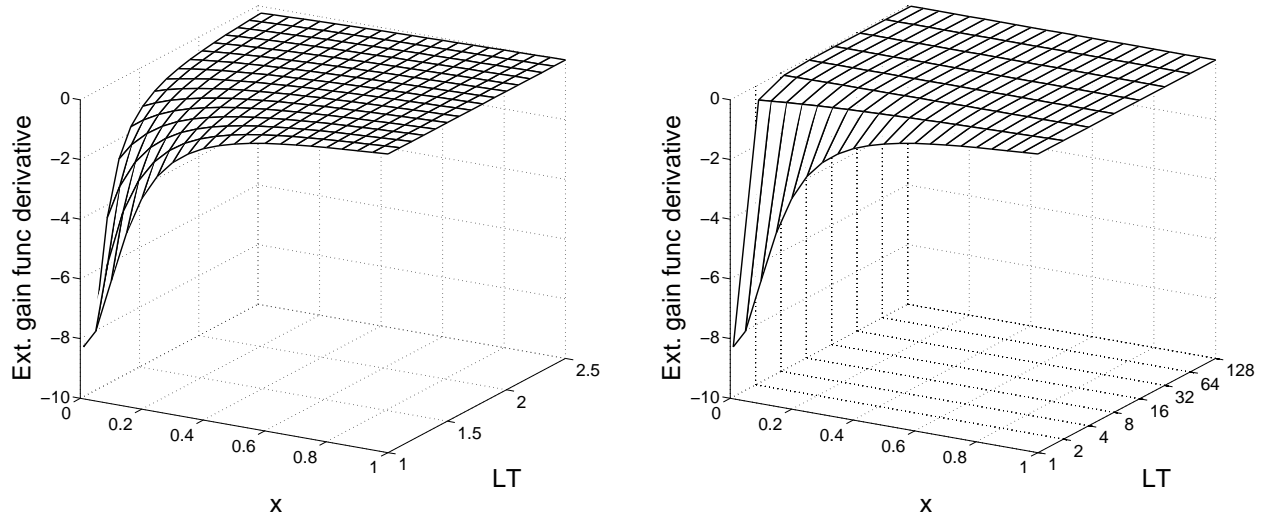


Figure 6.7: Extended gain function derivative of Papaj’s experiment [133] plotted against genotypic value x and different values of the product of lifetime and learning parameter LT (logarithmic scale) as specified in Equation 6.21. The left panel zooms into the range of low values of LT , the right panel shows a larger LT range. For all combinations of LT and x the extended gain function derivative is negative, but almost zero for larger LT values. This implies that an increase in LT slows down evolution but no substantial further deceleration can be expected above a certain threshold of LT .

6.2.3 Extended Gain Function Analysis

In Section 6.2.2, an analysis based on the basic gain function framework was used to support the simulation results of Papaj that the addition of learning slows down the evolution of genetically strong individuals. Now the extended gain function framework of Section 4.4 is used to gain further insights into the effect of learning on evolution in Papaj’s model.

Recalling Equation 6.16

$$f_{\phi_l}(x, L, T) = 1 - \left((x - 1) \frac{e^{-LT} - 1}{LT} \right)^2, \quad (6.20)$$

and that the product LT can be interpreted as one variable, the gain function derivative of the extended framework is calculated as

$$\frac{\partial^2}{\partial x \partial (LT)} \log f_{\phi_l}(x, L, T) = \frac{4LT e^{2LT} (-1 + e^{LT}) (-LT + e^{LT} - 1) (x - 1)}{(2e^{LT} (x - 1)^2 - (x - 1)^2 + e^{2LT} ((LT)^2 - (x - 1)^2))^2}. \quad (6.21)$$

A step-by-step derivation of this equation is presented in Appendix C. This gain function derivative is shown in Figure 6.7. Clearly, for all combinations of LT and x the right-hand side of Equation 6.21 is negative. However, the extended gain function derivative additionally reveals that for larger values of LT ($LT > 2.5$) the derivative is almost zero. If, as in the previous section, $T = 100$ is assumed, we learn from the extended analysis that increasing the learning parameter beyond $L = 0.025$ does not substantially accelerate evolution further.

6.2.4 Continual versus Posthumous Fitness Assessment

In Section 5.3, the concepts of continual and posthumous fitness assessment have been introduced. As mentioned above the fitness assessment model of Papaj's formulation can be considered as a type of continual fitness assessment even though Papaj does not introduce an "adaptive value" function.

If in Papaj's model only the result of learning is taken into account (posthumous fitness assessment), the fitness in case of learning is given by

$$f(z(x, L, T)) = 1 + (x - 1)e^{-LT} , \quad (6.22)$$

(cf. equations 6.13 and 6.14). Now, assume that posthumous fitness is the reference case and we want to investigate how accounting for learning curves influences the rate of evolution, compared to the case of posthumous fitness assessment. This can again be done using the basic gain function framework. In particular, the fitness in case of posthumous fitness assessment becomes the denominator of the gain function and the fitness in case of continual fitness assessment (Equation 6.16) becomes the numerator of this special gain function which is denoted as $g^*(x)$,

$$g^*(x) = \frac{1 - \left((x - 1) \frac{e^{-LT} - 1}{LT} \right)^2}{1 - ((1 - (1 + (x - 1)e^{-LT}))^2)} = \frac{1 - \left((x - 1) \frac{e^{-LT} - 1}{LT} \right)^2}{1 - (x - 1)^2 e^{-2LT}} \quad (6.23)$$

With $C_1 = (e^{-LT} - 1)^2 / (LT)^2$ and $C_2 = e^{-2LT}$ the corresponding gain function derivative can be written as

$$\frac{\partial g^*}{\partial x} = \frac{2(x - 1)^2 (C_1 - C_2)}{1 - C_2(x - 1)^2} . \quad (6.24)$$

The gain function of Equation 6.23 and its derivative in Equation 6.24 are shown in Figure 6.8 for various combinations of LT . The gain function is increasing (has a positive derivative) in x direction. This means, compared to posthumous fitness assessment as the reference case continual fitness assessment (as Papaj has done) accelerates evolution. Recall that accounting for learning with continual fitness assessment decelerates evolution compared to the complete absence of learning. Thus, deceleration caused by learning with continual fitness assessment is weaker than deceleration caused by posthumous fitness assessment. Here, "weaker deceleration" is equivalent to "acceleration".

6.2.5 Discussion

Similar to the analysis of the Hinton and Nowlan model in the previous section, the gain function analysis of Papaj's experiment [133] allows to derive a clear analytical argument for the observed simulated evolutionary dynamics. In contrast to Hinton and Nowlan's model, the gain function derivative is negative and evolution is decelerated through individual learning in Papaj's experiment. The gain function analysis does not only confirm the simulation results but using the extended framework it is also possible to identify the "interesting" regions of the model parameter space in which learning has a substantial influence on evolution. Furthermore and beyond Papaj's results the influence of learning curves under continual




Figure 6.8: Basic gain function that compares continual fitness assessment with posthumous fitness assessment (as the reference case in the denominator of the basic gain function) in Papaj’s experiment. The gain function is increasing (left panel), i.e., has a positive derivative (right panel) in x direction. Thus, compared to the case of posthumous fitness assessment accounting for learning curves (as Papaj has done) accelerates evolution.

fitness assessment compared to the case when learning curves are not taken into account under posthumous fitness assessment is determined. It is found that learning curves accelerate evolution in Papaj’s experiment.

6.3 Mathematical Models with Developmental Noise

Most mappings from genotype to phenotype have a random component. This holds for virtually all species in nature, but also for many artificial systems. This random component is often called developmental noise. The influence of developmental noise on evolution has been studied in a few papers. In this section, these models are revisited and analyzed with the gain function framework of Section 5.1.2.

6.3.1 Existing Models

At least three papers [18, 5, 3] look at the influence of developmental noise on the rate of evolution. All three papers assume a Gaussian fitness landscape of the form

$$f(x) = ce^{-s(x-x_{\text{opt}})^2} \quad (6.25)$$

(cf. Figure 6.9) and conclude that developmental noise slows down genetic evolution. In all cases normally-distributed developmental noise is assumed, hence Equation 5.13 (or alternatively Equation 5.15) which requires symmetric noise can be applied to determine the sign of the gain function derivative.

$$f(x)f'''(x) - f'(x)f''(x) = (x - x_{\text{opt}})8s^2c^2e^{-2s(x-x_{\text{opt}})^2} \quad (6.26)$$




Figure 6.9: Gaussian fitness function as used in [18, 5, 3], see Equation 6.25 with parameters $c = 1$, $s = 1$, $x_{opt} = 2$.

For $x < x_{opt}$, f is increasing, and $f(x)f'''(x) - f'(x)f''(x) < 0$ which implies that $g'(x) < 0$. The same argument applies to $x > x_{opt}$ where $g'(x) > 0$, hence $\text{sign}(g(x)) = \text{sign}(-f'(x))$ which implies learning-induced deceleration.

Thus, the gain function analysis confirms the results of [18, 5, 3] who took a different analytical approach to derive the same conclusion.

6.3.2 Discussion

It should be noted that the conclusion of [18, 5, 3] that developmental noise slows down evolution resulted from their choice of a Gaussian fitness function. In this thesis, it has been shown in Section 5.1.2 that (symmetric) noise can also accelerate evolution. The only requirement is that $(f(x)f'''(x) - f'(x)f''(x))$ is positive.

6.4 Biological Data - An Inverse Gain Function Application

In the models that have been investigated so far, knowledge about the fitness landscape and the learning algorithm was given and this knowledge was used in the gain function framework to predict the evolutionary dynamics. However, the logical equivalence in Equation 4.9 tells that an “inverse” approach is also possible:

Given some evolutionary data (in absence and presence of learning), one can derive the sign of the gain function. In other words, we learn something about the effect of learning on fitness and learn something about the learning mechanism. In the following, this is done in a rather qualitative way with data from the first biological (*in vitro*) experiment that demonstrated the Baldwin effect [107] in the evolution of resource preference in fruit flies.

6.4.1 In Vitro Evolution of Resource Preference

In this experiment, Mery and Kawecki studied the effect of learning on resource preference in fruit flies (*Drosophila melanogaster*). For details of the experiment, it is referred to [107]. In

the following only a brief qualitative description is provided: The flies had the choice between two substrates (pineapple and orange) to lay their eggs on, but the experimenters took only the eggs laid on pineapple to breed the next generation of flies which are (after grown up) given the same choice for their eggs. Measuring the proportion of eggs laid on pineapple, one could see that a stronger preference for pineapple evolved, from 42 percent in the first generation to 48 percent in generation 23. To test the Baldwin effect another experiment was done, where also eggs laid on pineapple were selected to breed the next generation, but flies could previously learn that pineapple is the “good” substrate. To allow for learning, several hours before the experimenter took away the eggs for breeding, the dis-favored orange was supplemented with a bitter-tasting chemical for some time (and replaced with a “fresh” orange after that). If flies learned to avoid orange, they would lay fewer eggs on it later, i.e., show a stronger preference for pineapple. After 23 generations of learning, the innate preference (measured in absence of the bitter chemical) evolved to 55 percent, significantly more than the 48 percent that evolved in the absence of learning. Thus, in this experiment learning accelerated evolution. According to Equation 4.9 the gain function has a positive derivative.

Mery and Kawecki did the same experiment with orange as the favored substrate, i.e., eggs for breeding were taken from orange, and pineapple was supplemented with the bitter-tasting chemical in case of learning. In 23 generations the innate preference for orange evolved from initially 58 percent to 66 percent in presence of learning but to even more, 72 percent, in absence of learning. Thus, in this setting, learning decelerated evolution. According to Equation 4.9 the gain function has a negative derivative. The first row of Table 6.1 summarizes the experimental results. As in [107] the cases when pineapple was the favored resource is referred to as *Learning Pineapple* in case of learning and *Innate Pineapple* in absence of learning, and correspondingly *Learning Orange* and *Innate Orange* when orange was the favored resource.

6.4.2 A Qualitative Gain Function Analysis

The following analysis aims to shed some light on these - seemingly contradictory - results. If the relationship between innate resource preference and success of the resource preference learning is independent of what the high-quality resource currently is, the experimental results can be interpreted as follows: When evolution starts from a relatively weak innate preference for the favored fruit (42 percent as in the first experiment with pineapple as the high-quality resource), this leads to learning induced acceleration. However, if evolution starts from a relatively strong innate preference for the favored fruit (58 percent as in the second experiment with orange as the high-quality resource) this leads to learning induced deceleration of evolution. Therefore, if evolution started further away from the evolutionary goal, then learning accelerated evolution, implying an increasing gain function, and if it started closer to the evolutionary goal, learning decelerated evolution, implying a decreasing gain function. Thus, in principle one can expect a gain function that is increasing for a weak innate preference for the target fruit and decreasing for a strong innate preference for the target fruit. This implies a maximum gain function value at an intermediate innate preference for the target fruit and lower gain function values for weak and strong innate preferences.

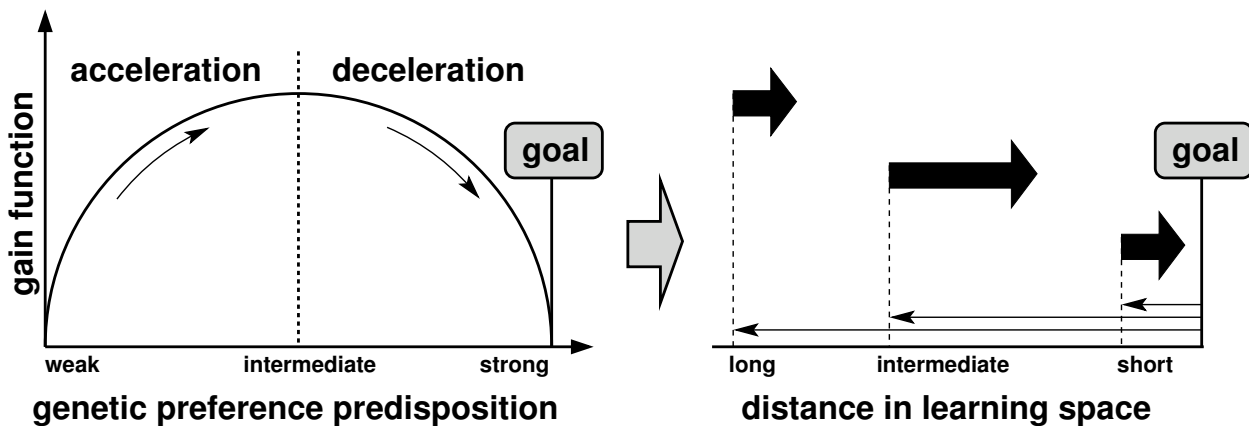


Figure 6.10: Illustration of the qualitative gain function analysis of the fruit fly experiment. The biological data in [107] suggest that learning is most successful for an intermediate distance between individual genetic predisposition and the target predisposition. This leads to a gain function as shown in the left side, increasing first and then decreasing. This gain function implies a learning pattern as shown in the right side. The length of the thin arrows indicates the initial distance to the learning target and the length of thick arrows indicate the corresponding success of learning. Evidence for such a learning pattern is supported by findings in [137] and [122].

Recalling that the gain function $g(x) = f_l(x)/f(x)$ reflects the relative fitness gain due to learning, it seems that learning is not very effective when the starting point of learning is far away from or very close to the learning goal (low gain function values), and is probably most effective for a starting point with an intermediate distance to the learning goal.

Besides these conclusions from the experimental results, there are other arguments for such a relationship:

For an individual that already shows strong innate preference for a high-quality resource, its learning success might be low because perfection is usually difficult (and requires large resources), or simply because the preference cannot be increased beyond 100 percent.

In contrast, there is scope for a large effect of learning in individuals that show a weak preference for the high-quality resource, i.e., strong preference for the low-quality resource. However, there are two reasons why such individuals with strong innate preference for the low-quality resource might be slow in changing their preference toward the high-quality resource. Firstly, because of their strong initial preference for the one resource, individuals will only rarely sample the other one, and thus rarely have a chance to find that the other resource is in fact better. Secondly, even if they occasionally sample the other resource, their strong innate preference for the first one may be difficult to overwrite. This argument is supported by experiments with *phytophagous* insects (organisms that feed on plants), e.g., [137] and also with humans [122].

Figure 6.10 illustrates the conclusion of the qualitative gain function analysis.

6.4.3 In Silico Evolution of Resource Preference

To test these conclusions, the biological experiment is studied *in silico*, i.e., simulated using an artificial evolutionary system of resource preference. In the simulation model, the innate preference for orange is genetically encoded as $x \in [0, 1]$ and represents the probability to choose orange in a Bernoulli trial. If the individual fails to choose the high-quality resource, it does not produce offspring. However, if the high-quality resource is chosen, the "digital fly" receives a fitness score of 1, which results in a high probability to produce offspring for the next generation (assuming a linear-proportional selection scheme).

Thus, if pineapple is the high-quality resource, the expected fitness in absence of learning f^P is given by $f^P(x) = 1 - x$ (*innate pineapple*). Since learning is on average beneficial, the fitness in presence of learning $f_l^P(x)$ must be larger, i.e., $f_l^P(x) \geq f^P(x)$ (*learning pineapple*). Correspondingly, if orange is the high-quality resource, we obtain $f^O(x) = x$ (*innate orange*), and $f_l^O(x)$ (*learning orange*), where $f_l^O(x) \geq f^O(x)$.

In the model, populations are initialized with $x \in [0.55, 0.61]$, and with an average orange preference of $\bar{x} = 0.58$. This is the same mean preference as observed in the initial generation of the biological experiment [107]. For the simulation, a population size of 150 is chosen, which is similar to the biological experiment. Mutation is simulated by adding a random number from a normal distribution with mean 0 and standard deviation $5 \cdot 10^{-5}$, i.e., a small effect of mutation on resource preference is assumed.

A gain function that is increasing for weak, maximal for intermediate, and decreasing for strong innate preference for the high-quality resource is given by a linear transformation of the Gaussian function $\phi(x, \sigma)$:

$$g(x, \alpha, \sigma) = a_1(\alpha, \sigma) + a_2(\alpha, \sigma) \phi(x, \sigma), \quad (6.27)$$

where $a_1(\alpha, \sigma) = 1 - \frac{\alpha \phi(0, \sigma)}{\phi(0.5, \sigma) - \phi(0, \sigma)}$ and $a_2(\alpha, \sigma) = \frac{\alpha}{\phi(0.5, \sigma)}$, such that g is 1 at the genotype boundaries and maximal in the center of the genotype space ($x = 0.5$). Parameter a reflects the maximum relative fitness gain (at $x = 0.5$) that can be achieved through learning. In the biological experiments of Mery and Kawecki [107], the fitness gain due to learning was assessed by comparing the innate preference and the preference after learning (given by the proportion of eggs on the fruit substrate) at generation 23. Depending on *if* and *what* the ancestor populations have learned, and what the target resource in the assay was, the fitness gain varied widely in the biological experiment. Among the different settings the maximum fitness gain due to learning was an increase from 45 to 57 percent of eggs laid on the high-quality resource, i.e., a fitness gain of $(57 - 45)/45 = 0.27$. For the gain function of the simulation, Equation 6.27, a similar value $\alpha = 0.25$ was chosen. The only remaining parameter σ was tuned to get a maximally steep gain function in the preference region where evolution starts (satisfying that $f_l(x)$ is still monotonic) resulting in $\sigma = 0.075$.

Having defined the gain function $g(x)$ and the expected fitness in absence of learning $f(x)$ ($f^P(x) = 1 - x$ in case of pineapple selection) the fitness in case of learning $f_l(x) = g(x)f(x)$ (cf. Equation 4.9) can be derived. Figure 6.11 shows how learning influences the fly's probability to choose orange and the resulting gain function. Based on these properties of the evolutionary system a simulation study can be done.

Figure 6.12 shows the simulated evolution of the mean innate preference for orange. The innate preference for orange evolves faster in the absence of learning (*Innate Orange*) than




Figure 6.11: Simulation model for the evolution of resource preference of fruit flies. The left panels shows the influence of learning on the fly's probability to choose orange for different values of the innate preference for orange x (the probability to choose pineapple is $1 - p_{orange}$) in the experiment with simulated evolution. The right panel shows the gain function, which is identical for learning orange and learning pineapple. The horizontal axis shows the genetic predisposition of the target fruit.




Figure 6.12: Simulation results of the evolution of resource preference of fruit flies. The figure shows the evolution of mean innate preference for orange (averaged over all individuals and 50 independent evolutionary runs, with \pm one standard error). Notice that the preference for pineapple is one minus the preference for orange. If orange is the high quality resource, learning decelerates evolution, however, if pineapple is the high quality resource, learning accelerates evolution. As in the biological experiment, a set of control runs have been carried out in which the high-quality food changes every generation between orange and pineapple.

in the presence of learning (*Learning Orange*). However, the innate preference for pineapple evolves faster in case of learning (*Learning Pineapple*) than without learning (*Innate Pineapple*). The short error bars (of the length of two standard-errors) indicate the statistical significance of the difference in evolved preferences. This qualitatively confirms the results of the biological experiment of [107]. In Table 6.1, the experimental results of the artificial evolution are directly compared to the results of the biological evolution. The numbers in brackets are normalized with respect to the initial preference. First of all, it can be seen that the effects of acceleration and deceleration are *qualitatively* identical. In both cases, with and without learning, and for both, orange and pineapple selection, evolution proceeds quicker in the natural evolution experiment. However, with regard to the normalized values, the relative difference between evolution with and without learning is very similar in the natural and artificial evolution.

6.4.4 Discussion

The aim of this experiment was not to quantitatively replicate the results of the biological experiment. Too many assumption need to be made in order to simulate evolution of natural fruit-flies realistically. For example, as a gain function simply a Gaussian function with a maximum at $x = 0.5$ was chosen. The biological data suggested that the maximum of the gain function lies between 0.42 and 0.58. No attempt has been made here to tune the simulation model, but simply the middle, 0.5, was chosen. If evolution starts at $x = 0.42$ (selection for pineapple), this means that the genotype interval in which evolution is accelerated is rather small. Certainly a larger optimal x -value allows to produce stronger learning-induced acceleration. Furthermore the biological gain function may not be symmetric. Thus acceleration (selection for pineapple) may have a different magnitude than deceleration (selection for orange). Direct knowledge about the mutation strength and the mutation symmetry in the biological experiment is not available, but the same strength of symmetric mutation over the entire genotype space was assumed in the artificial evolution. This may not correspond to reality either. For example in the absence of learning in the biological experiment, selection for orange produced a shift from 0.58 preference to 0.72 while selection for pineapple produced a shift from 0.42 to only 0.48 (in 23 generations).

Despite this, the gain function argument may not be the only explanation. Mery and Kawecki [107] discuss several other reasons in detail.

This shows that the gain function approach can be applied “inversely” in order to get a better understanding of the effects of learning on fitness. Of particular interest might be the insect learning pattern of the type illustrated in Figure 6.1, which might also apply to many artificial learning system.

6.5 Mathematical Models on the Fitness-Valley-Crossing Ability

Fitness landscapes are often characterized by a number of local fitness optima which are connected with evolutionary pathways that require to pass a local fitness minimum which is

Table 6.1: Experimental results for the *in vitro* evolution [107] and the *in silico* evolution of fruit flies. For both cases the average innate preference for orange after 23 generations is shown.

Selection for Orange				
orange preference	initial	evolved	w/o learning	with learning
<i>in vitro</i> evolution	0.58 (100%)		0.72 (124%)	> 0.66 (114%)
<i>in silico</i> evolution	0.58 (100%)		0.61 (105%)	> 0.59 (102%)

Selection for Pineapple				
pineapple preference	initial	evolved	w/o learning	with learning
<i>in vitro</i> evolution	0.42 (100%)		0.48 (114%)	< 0.55 (130%)
<i>in silico</i> evolution	0.42 (100%)		0.46 (109%)	< 0.48 (114%)

commonly named *fitness valley*. A population's evolutionary movement through this local minimum toward another optimum is called *fitness-valley-crossing*.

One way to analyze the population dynamics of fitness-valley-crossing in the case of coupled evolution and learning is an experimental study. This has been done in [109, 172] for a particular set of fitness landscapes. In both papers, it has been found that for the observed fitness landscapes and learning schemes, learning smooths the fitness landscape and thereby improves the fitness-valley-crossing ability. Experimental studies often face the difficulty to derive general conclusions since only a particular set of scenarios can be studied. For a more general understanding a theoretical analysis or at least a theoretical underpinning of experimental results is sometimes more desirable.

The gain function framework is an example of a theoretical analysis, however, as mentioned in Section 4.5, the gain function framework applies to *Directional Selection*. Directional selection is a selection type that moves the population toward higher fitness. The mathematical proofs of the two gain function frameworks in sections 4.3 and 4.4 are based on expected values. Thus, the gain function predicts the *expected* dynamics. Fitness-valley crossing, however, essentially requires an *unlikely* event. Hence an expectation-based approach can not (directly) be used to analyze the fitness valley crossing process. A stochastic analysis approach seems to be more appropriate for the prediction of a fitness valley crossing event. However, a stochastic analysis of an evolutionary population movement (with or without individual learning) has great (partly unresolved) challenges which are described in the following.

Evolution as it is commonly modeled, is a discrete-time stochastic process which has the *Markov property*, i.e., the system state in the next time step solely depends on the current state [50]. A Markov-Chain analysis therefore seems to be an appropriate tool for such a stochastic analysis. A state of the Markov chain for evolutionary dynamics needs to completely specify the current population. Also, the transition probabilities between all states of the models need to be derived. Two main difficulties arise from such an analysis

1. Large state space,
2. Deriving the transition probabilities,

which are briefly discussed in the following. A more thorough review can be found in [146].

6.5.1 Problem of Large State Spaces in Markov-Chain Analyses

An exact Markov Chain analysis of an evolutionary process with a population of n individuals and genotype space X that allows to express $|X|$ different genotypes requires to model a state space S with size

$$|S| = \binom{n + |X| - 1}{n - 1} \quad (6.28)$$

states (cf. [161]) and $|S|^2$ transition probabilities. This becomes impracticable even for relatively small n and $|X|$.

One way to circumvent the high computational cost is an infinite population model as, e.g., proposed by Vose [176]. There, the population is represented by a vector p with $\sum_{i=1}^{|X|} p_i = 1$, that describes the fraction p_i of individuals with genotype i and evolution can be modeled as a dynamical system. Like most other infinite set models an infinite population model has some drawbacks. For example, if one wants to investigate how much time a population needs to cross a fitness valley, an infinitely large population will depending on its initial state but independent of the characteristics of the fitness landscapes very quickly have a non-zero fraction that has crossed the fitness valley. The difficulty lies in the definition of thresholds that determine the “successful” fitness valley crossing.

Another solution to reduce the computational cost and still assume a finite population has been suggested in [161] and [111]. They suggest an *approximate* Markov Chain analysis which basically attempts to “lump” similar states together in order to reduce the large state space and the corresponding transition matrix. The methods in [161] and [111], however, still require to calculate the exact transition matrix in the beginning and are hence still impracticable for reasonably large n and $|X|$. An improved approach has been suggested in [112]. They use an idea from [176] to create the transition matrix of the “lumped states” directly. While making some scenarios computationally tractable, “lumping” naturally reduces the accuracy of the prediction of the evolutionary dynamics. It is unclear how severe this accuracy reduction is for the study of multi-modal fitness landscapes. The improved lumping method has only been applied to uni-modal functions yet where it (already) produces noticeable deviations from the original Markov-model.

6.5.2 Difficulty of Deriving the Transition Probabilities in Markov-Chain Analyses

Despite its computational cost, deriving the transition probabilities between all states is mathematically not straightforward. Yet, this has only been done for simple evolutionary models, predominantly for simple genetic algorithms [117, 34, 48, 170, 63, 161, 176, 112, 111] and models with population size 1 (cf. [146]). This difficulty may explain why no exact Markov chain analysis of a population-based model that couples evolution and individual learning has been published yet. Adding individual learning may also increase the state space because a genotype-phenotype distinction is required then.

A simple way to model learning in evolving populations would be to calculate for each genotype or innate phenotype in advance the expected fitness after learning and use the resulting “after learning fitness landscape” for a Markov-chain analysis.




Figure 6.13: *Illustration of Borenstein's Drawdown (Figure reproduced from [12] with permission from Elhanan Borenstein). Qualitatively, the Drawdown is defined as the largest fitness descent that the random walk has to pass on the fitness landscape in order to reach the global fitness maximum.*

6.5.3 Borenstein's Approach

Recently Borenstein et al. [13, 12] have proposed a stochastic model of evolution and learning that circumvents the above described difficulties. However their model has some other drawbacks. In their model, the genotype space is discrete and one-dimensional. A population is represented by one value of the genotype space (whether this value represents the average genotype value of the population or something similar is not specified) and the population's mutation-selection movement is modeled by a one-dimensional non-symmetric random walk [163]. Learning is modeled as fitness landscape change. More precisely, all forms of learning investigated in [13, 12] lead to a smoothing of the fitness landscape. The simplified evolutionary model allows a general conclusion that is valid for arbitrary fitness landscapes:

Borenstein et al. generally derive the expected time E_0^N the non-symmetric random walk needs to move from an initial (low-fitness) point at the left border of the one-dimensional genotype space ($x = 0$) to the global fitness maximum at the right border of the genotype space ($x = N$). Using this formulation it is shown that E_0^N is asymptotically bounded by $0.5 N^2(1 + R)$ where R is the so-called *Drawdown* [13, 12] which is also referred to as *negative extent* in [121]. The Drawdown is (qualitatively) defined as the largest fitness descent that the random walk has to pass on the fitness landscape in order to reach the global fitness maximum. See Figure 6.13 for a qualitative illustration of the Drawdown.

In an unpublished Master thesis [52] which revisits Borenstein et al.'s model, examples are presented in which the Drawdown is not a good indicator for the estimation of E_0^N . Thus, while the Drawdown theory allows general conclusions its simplified model of evolution and learning may oversee some properties of evolutionary dynamics. Furthermore, the Drawdown is unable to predict the evolutionary dynamics on a monotonic fitness landscape (where the Drawdown is zero).

Since the gain function applies to directional selection but not to fitness valleys, and the Drawdown vice versa, the two approaches in a way complement each other.

6.5.4 The Role of the Gain Function

As mentioned earlier the gain function framework applies to *Directional Selection* and can not predict the time needed to cross a fitness valley. However, the gain function may still provide some relevant information.

If the gain function is decreasing toward a local fitness maximum, learning reduces the selection pressure toward this local optimum. Hence, a population movement away from the optimum (possible towards the global optimum) becomes more likely.

6.5.5 Discussion

This section has reviewed the existing approaches to analyze evolutionary dynamics and shown that none of these allows to generally investigate the influence of learning on the fitness-valley-crossing ability of a population in a practical way. The suggested models are either not computationally or mathematically tractable, or they are not simple enough to derive general conclusion with respect to fitness landscape properties. These models only allow conclusions with respect to a particular fitness landscape rather than certain landscape properties. One exception is the model of Borenstein et al. [13, 12]. However, the generality of the derived conclusion comes at the expense of accuracy in the analysis. The gain function may provide some hints to predict the evolutionary dynamics near fitness valleys but does not allow a quantitative analysis.

At the current state, an experimental study may therefore be a reasonable alternative to investigate multi-modal fitness landscapes.

6.6 Summary and Conclusion

This chapter started with an analysis of Hinton and Nowlan's simulation [64] from the gain function point of view. Their results have been explained with different arguments in the literature, however, it has been shown that a selection pressure argument is sufficient to explain the main result. More generally, it is expected that also in other extreme fitness landscapes with large plateaus learning accelerates evolution.

Similar to the analysis of the Hinton and Nowlan model, the gain function analysis of Papaj's experiment [133] allowed to derive a clear analytical argument for the observed simulated evolutionary dynamics. However, in contrast to Hinton and Nowlan's model the gain function derivative is negative and evolution is decelerated through individual learning in Papaj's experiment. The gain function analysis did not only confirm the simulation results of Papaj but using the extended framework it was also possible to identify the "interesting" regions of the model parameter space in which learning has a substantial influence on evolution.

Next, the gain function analysis was used to confirm the results of some analytical approaches [18, 5, 3] to predict the influence of developmental noise on evolution. However, the conclusion of [18, 5, 3] that developmental noise decelerates evolution is not generally true but results

from their choice of a Gaussian fitness function. Fitness landscapes in which developmental noise accelerates evolution have been discussed in Section 5.1 of this thesis.

Then, the gain function was used to shed some light on evolutionary data of a biological *in vitro* experiment [107]. There the gain function framework was applied inversely to learn something about the learning mechanism in fruit flies that learn a resource preference.

Finally, some known approaches that may be useful to investigate the influence of learning on the fitness-valley-crossing ability and one model that actually provides a theoretical analysis on the influence of learning in this scenario have been reviewed. At the current state, all methods have some drawbacks and an experimental study may be a reasonable alternative to investigate the influence of learning on the fitness-valley-crossing ability.

In this chapter, several models from different fields have been reassessed with the gain function approach. The results of these analyses do not only confirm the universality of the gain function approach but also provide clear explanations for the evolutionary dynamics that have been observed in the models thereby adding to a more general understanding of the dynamics of coupled evolution and learning. Also the limits of the gain function approach have been discussed and an alternative use of simulation studies in such cases has been suggested.

Balancing Evolution and Learning

So far, in this thesis, the dynamics of evolution and learning have been investigated from the perspective of how learning influences selection pressure. Selection pressure is an important determinant of evolutionary genotypic change. This chapter studies the issue of balancing the two levels of adaptation. The assumption for this study is that an increase in the intensity of learning is associated with a decrease in the generation turnover. The generation turnover indicates the reproduction frequency. From a computational point of view, this trade-off arises from the allocation of computational resources, such as CPU time. Since both, evolution and learning demand computational resources, a conflict arises: When resources are limited, a high degree of learning implies a low degree of reproduction, and vice versa. In addition to the influence of selection pressure, the generation turnover is an important factor that influences genotypic change. Other things equal, a reduction of the generation turnover implies a slow down of genotypic change.

Considering an evolutionary process in presence of a resource conflict and complete absence of learning, one could ask, whether the addition of learning at the expense of a decreased generation turnover brings about an advantage with regard to the overall adaptation behavior. If this is the case, another question is, up to which point trading evolution for learning is beneficial.

In order to observe an overall advantage of coupled evolution and learning over pure evolution, the adaptational benefits arising from the inclusion of learning must over-compensate its costs. Examples for both a positive and negative cost/benefit balance have been studied with the gain function framework in preceding chapters. There, the decreased generation turnover has not been considered. Recalling these results, one might have doubts whether the inclusion of learning in presence of the generation turnover reduction can yield an overall adaptational benefit.

However, it is worthwhile taking a closer look at the conclusions that follow from a gain function analysis. A negative gain function derivative implies a slow-down of genotypic change. In the stationary environments studied so far, the slowed change has been interpreted as an adaptational disadvantage. However, this interpretation does not necessarily apply to

all environments. In changing environments, a slower rate of genotypic change does not always cause an adaptational disadvantage. The rationale behind this hypothesis is that an adaptational disadvantage in the short run may produce an advantage in the long run, and vice versa. Furthermore, aspects such as the maintenance of diversity play a different role in changing environments as opposed to stationary conditions.

In the following, the evolution/learning trade-off is studied with a simulation model. In this model, the cost of learning are modeled as the associated consumption of computational resources. The model employs two very similar trial-and-error adaptation mechanism that only differ from each other in that one is applied to the population (evolution) and the other to the individual (learning). The central parameter of the proposed simulation model allows to adjust the ratio of computational resources allocated by evolutionary adaptation steps and by individual learning. In addition to the computational plausibility of learning cost, the simulation model reflects several properties of evolution and learning from a biological point of view (cf. Section 7.1). To the best knowledge of the author of this thesis, yet no model has been published that incorporates the *trade-off* between evolutionary adaptation and learning.

The remainder of this chapter is organized as follows. Section 7.1 provides a background on computational and biological evolution/learning trade-offs. Section 7.2 discusses related work. In Section 7.3, the simulation model is introduced. Section 7.4 investigates the effects of varying the balance between evolution and learning on the population dynamics. In Section 7.5, conditions are investigated under which a balance with intermediate ratios of evolution and learning produces an adaptational advantage over the extreme strategies “pure evolution” and “pure learning”. Section 7.6 completes the chapter with summary and conclusions. Parts of this chapter are based on [126] and [127].

7.1 Computational and Biological Evolution/Learning Trade-Offs

As mentioned above, an evolution/learning trade-off is apparent in evolutionary computation models that are coupled with a local learning procedure: Computational resources can only be assigned once, either to the level of evolutionary adaptation or to the level of individual learning. Thus, with limited computational resources, an increase in the degree of learning, e.g., the number of learning iterations, implies a reduction of the computational resources that can be allocated for evolution. See Figure 7.1 for an illustration of this trade-off. With 0 learning steps, i.e., with only 1 evaluation applied to an individual, 100 percent of the resources are consumed by evolution. By an increase of learning steps from 0 to 1, i.e., with 2 instead of 1 calls to the evaluation function, the number of genetic mutations per time (mutation frequency) is reduced by 50 percent. If the number of learning steps increases further, the mutation frequency is reduced further. The mutation frequency indicates the intensity of evolutionary adaptation. Recall that in Section 2.2.3 it has been argued that a similar trade-off also exists in biology in the case that the population size is approximately constant.

In evolutionary computation, the actually realized point on the trade-off curve (Figure 7.1) is usually predefined. However, in nature the average individual lifetime is the emergent




Figure 7.1: A trade-off between population-based evolutionary adaptation and individual learning that can be found in evolutionary computation when the overall computational budget is constant.

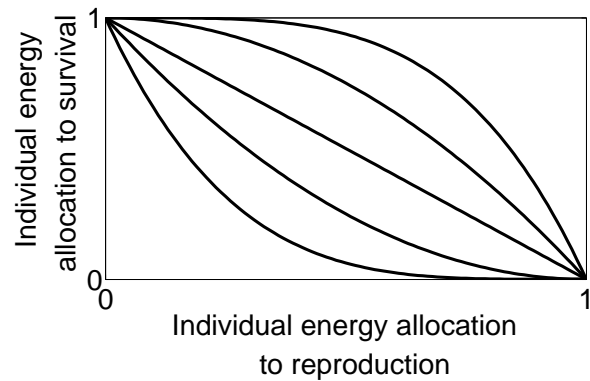


Figure 7.2: A biological evolution/learning trade-off: A qualitative sketch of possible trade-offs between reproduction and survival as can be found similarly in many species in nature is shown.

result of an evolutionary self-organization process. One possibly important mechanism arises from individual energy allocation (cf. Section 2.2.3). With a finite amount of individual energy, a trade-off between reproduction and survival efforts which both demand energy is apparent, see, e.g., [96, 156, 165, 116]. An increase in reproduction effort is associated with less lifetime (learning) and vice versa. In other words, longer lifetime – or more learning – implies less evolutionary adaptation. Figure 7.2 visualizes this kind of evolution/learning trade-off. Notice that the figure is a qualitative sketch of the a trade-off and does not refer to measured data.

In relation to the trade-offs depicted in Figures 7.1 and 7.2, it is worth noting that the distinction between “individual adaptation strategies” and “population adaptation strategies” has been used to categorize the diversity of natural adaptation mechanisms [4]. In Figures 7.1 and 7.2, population adaptation and individual adaptation would correspond to a point on the trade-off curve near the left and right border, respectively. It is evident that in some species – like bacteria – evolutionary adaptation is the dominant adaptation mechanisms, while other species – like mammals – rely much more on individual learning. So, it seems that in nature there is also an evolution/learning trade-off between species.

7.2 Related Work

Several computational models of evolution attempt to model the cost of individual learning. In some of these models, the cost of learning are taking into account as a fitness reduction term. This approach is referred to as *explicit learning cost*. A prominent example is Hinton and Nowlan’s [64] simulation study which has been analyzed using the gain function framework of Section 6.1. In their model, an increase in the length of the learning period causes a fitness decrease as well, cf. Equation 6.4. This has been taken into account in the gain function formulation in Equation 6.12. The gain function framework can be applied if the cost of learning can be modeled as a reduction of individual absolute fitness. Other examples for

explicit modeling of learning cost are Mayley's [102, 103] simulation studies with Kauffman's NK fitness landscapes [79] which show that learning may work to hide genetic differences between individuals and thereby decelerate evolution. Each performed learning step reduces the fitness by a fixed predefined amount.

It is questionable whether an explicitly specified parameter realistically reflects the cost of learning and the resulting fitness reduction. However, it should be noted that in most of the studies employing explicit cost, computational or biological plausibility is not the primary goal. Instead, these studies often intend to induce some selection pressure toward a reduction of the degree of individual learning.

To some extent, a few memetic algorithms (MAs) studies (cf. Section 2.4.3) contribute to the understanding of the optimal balance between evolution and learning. In MAs, the distribution of computational effort between evolution and local search has been identified as a critical design issue of optimization algorithms [57, 72, 58, 86]. The extent to which these studies contribute to this chapter is limited by a common property of MAs: The main motivation for the application of MAs is that they allow to employ distinct search techniques for coarse-grained evolutionary search and fine-grained individual search. Therefore, it is difficult to discriminate to what extent the performance improvement arises from the specialization of the techniques and to what extent from employing two instead of one search level. Investigating the latter aspect is the focus of this chapter. Therefore, findings in most MAs cannot be used to answer the question of this chapter which considers the coupling of two very similar adaptation mechanisms. Furthermore, MAs are commonly applied to stationary optimization problems under a Lamarckian inheritance scheme. In this chapter, however, changing environmental conditions and a biological (non-Lamarckian) inheritance mechanism are investigated.

Besides memetic optimization algorithm studies, there are some examples in literature on the evolution of artificial neural networks that implicitly contribute to the study of balancing evolution and learning. Ku et al. [90] investigate the influence of learning on evolution for the optimization of recurrent neural networks [99]. In their study they combine a cellular genetic algorithm [184] with different local hill climbing methods for the optimization of the synapse weights of the recurrent neural network. In one section of the paper a runtime analysis that includes all cost (for evolution and learning) is presented that compares coupled evolution and learning with evolution in the absence of learning. Accounting for the computational cost shows that learning decelerates evolution. The same authors confirm these results in a similar study [91].

Only few papers study the interplay between evolution and learning in dynamic environments. Among these are the ones discussed in Section 3.1 that compare Lamarckian and biological inheritance. However, these papers do not compare the presence of learning under biological inheritance with the absence of learning. Among other papers that consider dynamic environments only few account for learning. A notable exception is Dopazo et al. [31] who extend the model of Hinton and Nowlan [64]. In their model, the optimal phenotype switches between two states. As in Hinton and Nowlan's model, each additional learning step reduces the fitness by a pre-specified amount. Dopazo et al. observe that evolution works to reduce plasticity up to a point that still allows to learn either of the two optimal phenotypes. Suzuki and Arita [171] come to a similar conclusion. They study the evolution of agents that play the iterated prisoners dilemma in which individual strategies have varying degrees of

plasticity. The authors claim that the interaction of agents produces a dynamic environment. Evolution produces strategies with an intermediate degree of plasticity. It seems that this result can be explained in a similar fashion as in Dopazo et al. Plasticity incurs cost since the individual has to undergo a learning period. Yet, under dynamic conditions, plasticity has the benefit of providing flexibility. Thus, evolution works to optimize the cost-benefit balance.

Although Anderson [5] does not consider the cost of learning in his quantitative genetics study on evolution and learning, his results are to be mentioned. In a simulation study, he observes the evolution of a high degree of learning along with the emergence of a high degree of genotypic diversity. In Section 7.4, the influence of learning on population diversity will also be studied. This issue has recently been studied in (at least) two other papers. In the work by Curran and O’Riordan [27], cultural learning is added to a genetically evolving population in simple game playing scenarios. It is found that this type of learning increases diversity. Although the Hiding effect is not mentioned by the authors of [27], it is an intuitive explanation for the results. Recall that the Hiding effect can be observed when learning reduces individual fitness differences. This corresponds to a negative gain function derivative. Through cultural or any other form of learning, individuals with genetic differences can achieve a similar fitness which results in a higher diversity in the population than in the absence of learning. In another paper by the same authors [26], it is shown that the addition of cultural learning leads to a better adaptation behavior under dynamic environments. An explanation for this may also be the increased level of diversity induced by the Hiding effect. Notice however that a higher degree of diversity is not always favorable. Whether diversity leads to an adaptational advantage can be assessed only from the perspective of *what level* of diversity is most appropriate for a given dynamic environment with regard to the exploration *and* exploitation capability of the population, as will be seen later. This issue is not discussed in [26] and [27].

7.3 Simulation Model

A population is composed of n individuals, each represented by a 3-tuple

$$I = (x, z, t_b) , \quad (7.1)$$

where x is a one-dimensional real-valued *genotype*, z is a (variable) one-dimensional real-valued *phenotype*, and t_b is the time when the individual is born. How well an individual’s phenotype z is adapted to the environment at time t is specified by its *adaptive value*. The latter is given by a time-dependent environment function $v(z, t)$. Selection is based on the adaptive values of the individuals and takes place every time step.¹ The default population size is $n = 1000$.

Figure 7.3 shows the loop of the evolutionary and the learning aspects of the adaptation process, representing one time step in the discrete time model. Adaptation steps 1 to 3 represent evolutionary adaptation steps (right-hand side of the figure), whereas adaptation steps 4 to 6 represent steps of individual learning (left-hand side of the figure). The genotype-phenotype-

¹ Notice that in evolutionary optimization, v is usually called *fitness*. The usage of *fitness* as reference to the “temporary basis of selection” would be inconsistent with the biological lifetime concept of fitness. This chapter applies the terminology introduced in Chapter 2.

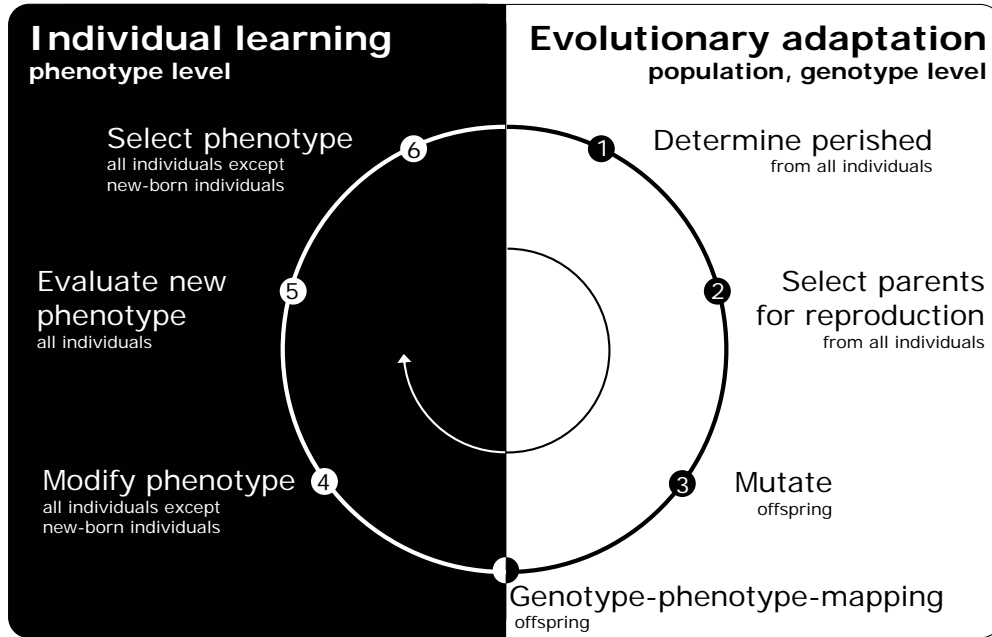


Figure 7.3: Adaptation loop that takes place every time unit in the analysis model of coupled evolution and learning.

mapping represents the transition from individual learning to evolutionary adaptation. In the following, the adaptation steps are described in detail.

7.3.1 Evolutionary Adaptation

Determine perished (Step 1): Individuals have a limited lifetime which is specified by the model parameter L . Based on an individual's time of birth t_b , lifetime L , and the current time t , the number of individuals that reach the end of their life is determined (condition $t - t_b = L$). Let m be the number of perished individuals in one evolutionary cycle. Notice that L can be interpreted as a learning parameter.²

Select parents for reproduction (Step 2): Based on the individual's current adaptive value $v(z, t)$, m individuals are selected to produce offspring to replace the m perished ones. The selection probability is linearly proportional to the current adaptive value. This selection scheme is implemented by the *stochastic universal sampling* algorithm proposed in [6]. The union of the original population with the m perished individuals removed and the m new offspring forms the new population. Therefore the population size is constant over time.

Mutate (Step 3): The genotype of the new offspring is mutated by adding a random number sampled from a normal distribution $X_{\phi(0, \sigma_G)}$ with mean $\mu = 0$ and standard deviation $\sigma = \sigma_G$ to the original (parent's) genotype x . Formally,

$$x' = x + X_{\phi(0, \sigma_G)} \quad , \quad (7.2)$$

where x' is the genotype of the offspring.

² In the nomenclature on page IX, learning parameters are denoted a , thus L is one dimension of a .

7.3.2 Genotype-Phenotype-Mapping

This step is only applied to offspring (new-born individuals). The innate phenotype is specified by the genotype value (identity function to their genotype).

Using the formal definition of Equation 7.1, the whole process of (asexual) reproduction of an offspring I' from a parent I at time t is given by

$$(I, t) = ((x, z, t_b), t) \mapsto (x', x', t) = I' \quad , \quad (7.3)$$

where x' is calculated according to Equation 7.2. The offspring's lifetime L is equal to its parent's lifetime L and the time of birth is set to the current time.

7.3.3 Individual Learning

All individuals (except the new-born offspring³) try to increase the adaptive value using a simple learning algorithm. This algorithm employs a rudimentary lifetime memory that stores the best phenotype found so far and the corresponding adaptive value that has been achieved at the time when the phenotype value was explored. The learning procedure comprises the following three steps.

Modify phenotype (Step 4): Individuals explore their environment by modifying their current phenotype. This is modeled by adding a random number sampled from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = \sigma_P$. Formally,

$$z_t = z_{t-1} + X_{\phi(0, \sigma_P)} \quad , \quad (7.4)$$

where z_t denotes the phenotype in time t . σ_P can be seen as a learning parameter (σ_P is one dimension of a , cf. nomenclature on page IX).

Evaluate new phenotype (Step 5): The new phenotype (for new-born individuals the innate phenotype) of each individual is evaluated in its current environment. An environment is specified by a time-dependent adaptive value function, i.e., the adaptive value becomes $v(z_t, t)$. In this chapter, four different environments are studied, as illustrated in Figures 7.6, 7.7, 7.11, and 7.12. In addition to the mapping from phenotype z to the adaptive value v , the time dimension (t) is included for visualization of the environmental dynamics. This adaptive value becomes the selection criterion for this time step regardless of whether it is kept or not in the next step.

Select phenotype (Step 6): The adaptive value of the new phenotype of each individual is compared to its adaptive value at time $t - 1$. If the new adaptive value is lower than the memorized one, the memorized phenotype remains the current phenotype.

7.3.4 Remarks

The implementation of this kind of learning can be seen as a (1+1)-EA [150] and is very similar to the EA employed on the population level. Notice that in both levels of adaptation the same adaptation step-size is chosen, i.e., $\sigma_G = \sigma_P$ which is intended to avoid a methodological

³ This exception allows a precise distribution of adaptive value function calls between evolution and learning; during the first lifetime unit, only the innate phenotype is evaluated and no modification of it takes place.

bias towards a particular adaptation mechanism, in particular $\sigma_G = \sigma_P = 0.01$ is set in all simulation studies of this chapter.

The model allows to distribute the adaptation effort between evolution and learning by setting the lifetime parameter L . With population size n , the average generation turnover, i.e., the average number of individuals that perish and are replaced in one time unit, is n/L . For example, with $n = 1000$ individuals, $L = 20$ causes an average generation turnover of 50 individuals per time unit, i.e., 5 percent of all adaptation steps are genetic mutations and 95 percent are individual learning steps. Setting $L = 1$ would cause an average generation turnover of 1000 individuals per time unit, i.e., 100 percent of all adaptation steps are evolutionary steps. The case of $L = 1$ represents the special case of *pure* evolution, the case of $L = 10$ represents a case of *coupled* evolution and learning. Hence, L can be interpreted as a learning parameter.

7.4 Influence of Lifetime on Population Dynamics

The balance between *exploration* and *exploitation* is the key issue for adaptive systems [67, 35]. In this section, it is investigated how a change in the lifetime parameter L that adjusts the distribution of evolution and learning influences population dynamics with a focus on the exploration/exploitation balance. There are two obvious effects of increased lifetime/learning.

First, as mentioned earlier, *an increase in lifetime L reduces the average generation turnover*. Recall that the average generation turnover of a population of n individuals is $\frac{n}{L}$.

Secondly, *an increase in lifetime/learning L smooths the effective fitness landscape*. Learning influences the average adaptive value over an individual's lifetime. The mapping from genotype to its average adaptive value is called *effective fitness* here.

In principle, an increase of learning intensity can either increase or decrease selection pressure as we have seen in Chapter 4. However, there are only rare cases where the selection pressure is actually increased. The model of Hinton and Nowlan [64] is such an example. Borenstein et al. refer to such scenarios as *extreme fitness landscapes* [13]. In all other cases, an increase in learning intensity decreases selection pressure, and an increase in L smooths the effective fitness landscape.

It is well known that the exploration and exploitation ability is strongly influenced by the degree of *diversity* which is an important concept of this chapter. A second important concept of this chapter is the *quasi-species*. In the following, both concepts are briefly introduced.

Diversity

In [98] diversity is described as “the variety and abundance of organisms at a given place and time”. However, it is less clear how to measure it [22, 98]. Numerous diversity measurements have been suggested in the biology literature, but it is shown in [145] that no single measurement is universally superior. In artificial life and evolutionary computation, diversity has been studied, too [114, 100]. Every diversity measurement reflects at least one of the two aspects, namely, *richness* and *evenness*. Richness defines how many different elements exist in an element set, e.g., how many species live in the ecosystem or how many different genotypes or phenotypes are present in a population. Evenness refers to the distribution of

a given set of elements with respect to certain element properties. A uniform distribution produces the maximum evenness. Some diversity measurements focus on either of the two properties, e.g., the Hurlberts Probability of Interspecific Encounter [71] quantifies evenness. Most measurements aggregate both aspects, such as the Shannon entropy [154, 136] and the Simpson index [157]. In this chapter, the intuitive and widely used Simpson index is adopted for measuring diversity, as recommended in [145]. The Simpson index reflects the probability for two randomly sampled individuals to fall into the same partition class of a partitioned continuous space and is defined as

$$H = 1 - \sum_{i=1}^m \left(\frac{n_i}{n}\right)^2, \quad (7.5)$$

where n is the population size, and n_i the number of individuals in partition class i (out of m partition classes). H increases with both, evenness and richness which is a desired property. For all studies, Shannon-Evenness, Hurlberts Probability of Interspecific Encounter, the standard richness measure, and the average Euclidean distance have also been used as diversity measures with which, without exception, qualitatively equal results have been obtained.

The Quasi-Species

Here, the concept of the quasi-species is introduced briefly. Refer to [37, 38, 15] for a detailed description of this concept.

If the environment is static for a sufficiently long time, the population moves to a local or global optimum where the selection pressure is usually decreasing and the loss of diversity is slowed down. At the same time, mutation keeps introducing new (non-optimal) genotypes to the population. At some point in time the two forces selection and mutation that reduce and increase diversity simultaneously are equally strong and cancel each other out. This well-known phenomenon has been termed *mutation-selection balance* [54, 25]. Under mutation-selection balance the population forms a “cloud” of genotypes around the so-called *wild-type*. This genotype cloud at mutation-selection balance is called *quasi-species* [37, 38].

The following illustration is inspired by the description of the quasi-species in [15]. Assume that there are only two genotypes, an optimal and a non-optimal one where x denotes the fraction of optimal genotypes in the population, implying that $(1 - x)$ is the fraction of non-optimal genotypes. Further, h denotes the fitness of the optimal genotype in multiples of the non-optimal genotype’s fitness, i.e., h indicates the degree of selection pressure. The variable p denotes the probability that a non-optimal genotype mutates to the optimal one which is equal to the probability that an optimal genotype mutates to the non-optimal one. If reproduction is linear-proportional to fitness, evolutionary dynamics are described by a difference equation that describes the increase of the fraction of optimal genotypes, x_t , over time,

$$x_{t+1} = (1 - p) \frac{hx_t}{hx_t + (1 - x_t)} + p \frac{1 - x_t}{hx_t + (1 - x_t)} = \frac{(1 - p)hx_t + p(1 - x_t)}{hx_t + (1 - x_t)}. \quad (7.6)$$




Figure 7.4: The fraction of optimal genotypes in a simple evolution model over time (Equation 7.7). Mutation p and selection pressure h determine the mutation-selection-balance (quasi-species formation). Generation turnover λ has no influence in the long run.

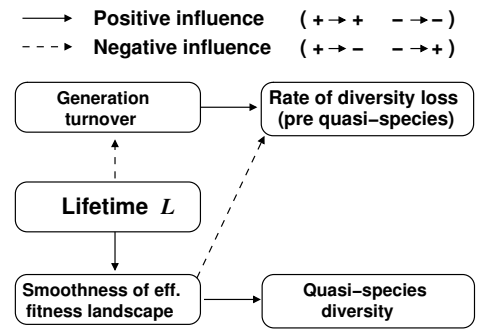


Figure 7.5: An increase in lifetime/learning causes a decrease in the rate of diversity loss and an increase in quasi-species diversity.

In the absence of mutation ($p = 0$) and for $h > 1$, x increases asymptotically toward 1. However, with mutation ($p > 0$) genotypes mutate *toward* and *away* from the optimum.

As an extension of the common quasi-species model the generation turnover is included in the difference equation. This is an extension of the common quasi-species model. With $\lambda = 1/L$ denoting the relative generation turnover (percentage of individuals that are replaced), one obtains

$$x_{t+1} = (1 - \lambda)x_t + \lambda \frac{(1 - p)hx_t + p(1 - x_t)}{hx_t + (1 - x_t)}. \quad (7.7)$$

Figure 7.4 shows how, according to Equation 7.7, the fraction of optimal genotypes evolves over time for different parameters. A lower generation turnover λ leads to a slowed convergence, i.e., to a slower loss of diversity and later formation of the quasi-species, but has no influence on the mutation-selection-balance. Hence, the composition of the quasi-species is solely determined by mutation probability p and selection pressure h .

In this simple model, diversity can be measured in the sense of evenness as $1 - |2x - 1|$, i.e., maximal diversity is given by $x = 0.5$. In Figure 7.4, we see that a small h (a smooth fitness landscape) leads to a higher quasi-species diversity. In conclusion, decreasing selection pressure leads to an increased quasi-species diversity regardless of the generation turnover. However, the generation turnover influences the rate of diversity loss before quasi-species formation. See Figure 7.5 for an illustration of this conclusion. Notice that this model does not account for finite population effects.

Both concepts, diversity and quasi-species, play an important role in the analysis of the following simulation studies.

7.4.1 Influence of Learning on Diversity (Environment 1)

The influence of lifetime/learning on diversity is investigated with a simulation study in Environment 1 (Figure 7.6) which is defined by the adaptive value function

$$v_1(z, t) = e^{-z^2}. \quad (7.8)$$




Figure 7.6: *Environment 1: A uni-modal, stationary Gaussian function.*




Figure 7.7: *Environment 2: A composition of two Gaussian functions, where the optimum moves from 0 to 1 at time 10000.*

Function v_1 is a Gaussian function centered at $z = 0$. Environment 1 is stationary, i.e., the mapping from z to f is independent of t . In the following simulations the population is initially distributed uniformly on $[-2, 2]$. Figure 7.8 shows the population dynamics of typical evolutionary runs. Each thick black dot represents the genotype of one individual at a time, each thin gray dot represents a phenotype. Notice that for this visualization the original population size of 1000 has been reduced to 100.

With a lifetime of $L = 1$ (pure evolution), the population quickly converges to a stable quasi-species state. With a lifetime of $L = 20$, the quasi-species formation takes about 5 time units. In case of $L = 20$ (coupled evolution/learning), this takes significantly longer and the quasi-species is less stable. After 500 time units, the diversity seems to be slightly higher with coupled evolution/learning than with pure evolution. From these observations the following hypotheses are derived:

1. Higher lifetime slows the speed of genotypic diversity loss.
2. Higher lifetime increases quasi-species diversity.

A second simulation study confirms these hypotheses. Figure 7.9 shows how diversity ⁴ averaged over 500 independent evolutionary runs evolves over time. The thin black line shows the average genotype (equals phenotype) diversity in case of pure evolution ($L = 1$). The case of coupled evolution/learning ($L = 20$) is denoted with a thick black line showing the average genotype and a thick gray line showing the average phenotype diversity, respectively. The trajectory resulting from an additional experiment is shown as dashed line. In this additional experiment, all individuals have a lifetime of $L = 20$ but learning is disabled, thereby avoiding the smoothing of the effective fitness landscape (Hiding effect). Hence, an individual's phenotype value is equal to its genotype throughout its lifetime. This additional experiment allows to separate the influence of reduced generation turnover and fitness landscape smoothing.

⁴ Simpson diversity, cf. Equation 7.5. Notice the space is discretized into partition classes $]-\infty, -3],]-3, -2.75],]-2.75, 2.5], \dots,]2.75, 3],]3, +\infty[$.




Figure 7.8: Evolutionary dynamics of a typical evolutionary run in Environment 1 in case of pure evolution ($L = 1$, top panel) and coupled evolution/learning ($L = 20$, bottom panel). Each thick black dot represents the genotype of one individual (out of a population of 100 individuals) at a time, each thin gray dot represents a phenotype.




Figure 7.9: Evolutionary dynamics in Environment 1. Left panel: Comparing the average diversity evolution in Environment 1 in case of pure evolution ($L = 1$) and coupled evolution/learning ($L = 20$). Coupling evolution and learning causes a slower genotypic diversity than pure evolution. Coupled evolution/learning also results in a higher genotypic quasi-species diversity and a lower phenotypic quasi-species diversity than pure evolution. Right panel: Mean distance to the optimum. After formation of the quasi-species, the population with coupled evolution/learning has on average a smaller phenotype distance to the optimum but a larger genotypic distance.

From Figure 7.9, it can be seen that with evolution/learning ($L = 20$), the rate of genotypic diversity loss is indeed lower than with pure evolution (compare the slopes of the thin and the thick black lines). The extent to which this is caused by the reduced generation turnover is represented by the difference between the thin black line and the dashed line. The extent to which the increased rate of genotypic diversity loss is caused by the smoothing of the effective fitness landscape, is represented by the difference between the slopes of the dashed line and the thick black line.

Meanwhile, it can be seen that a higher lifetime ($L = 20$) leads to a more diverse quasi-species. The average time of the formation of a quasi-species – all curves remain more or less constant – is approximately 15 in case of $L = 1$ and 300 in case of $L = 20$. The explanation for the higher quasi-species diversity is that L causes a smoothing of the effective fitness landscape that shifts the mutation-selection balance.

Although the phenotype is strongly dependent on the genotype, phenotypic diversity is lower after quasi-species formation. An explanation for this finding is that genetically different individuals may adapt to a similar phenotype during lifetime which directly reduces phenotypic quasi-species diversity. The latter argument is further supported by additional simulation results presented in the right panel of Figure 7.9. There the mean genotype and phenotypic distance to the optimum is shown (averaged over 1000 independent simulation runs). The population with evolution/learning has on average a smaller phenotypic distance to the optimum despite a larger genotypic distance.

In agreement with these findings, the experimental results of Curran and O’Riordan [27] show that the coupling of evolution with cultural learning produces a higher *genotypic* diversity than pure evolutionary adaptation. However, in disagreement with the findings of this section, Curran and O’Riordan find that the inclusion of cultural learning also leads to a higher *phenotypic* diversity. One reason for the disagreement might be that in Curran’s and O’Riordan’s model genotype and phenotype are represented in different domains and the authors employed different diversity measurements for these domains which prohibits a direct comparison.

In summary, an increase in the degree of learning, a) slows down the loss of genotypic diversity, and b) causes a higher genotypic quasi-species diversity despite a lower phenotypic quasi-species diversity. With regard to exploration, a high diversity is desired. However, with regard to exploitation a high adaptation velocity (loss of diversity) is desired. The following section shows how exploration and exploitation are affected by an increase in learning intensity.

7.4.2 Influence of Learning on Exploration and Exploitation (Environment 2)

Figure 7.7 shows Environment 2 which is defined by the time-dependent adaptive value function v_2 ,

$$v_2(z, t) = h e^{-\left(\frac{z - z_{\text{opt}}(t)}{\sigma_{\text{opt}}}\right)^2} + e^{-\left(\frac{z - (1 - z_{\text{opt}}(t))}{0.25}\right)^2}, \text{ with } h > 1, \quad (7.9)$$

$$\text{and } z_{\text{opt}}(t) = \begin{cases} 0, & \text{if } t < 10000 \\ 1, & \text{otherwise} \end{cases},$$

where h is a height factor that determines the difference of relative adaptive value between local and global optima. For instance, $h = 2$ means the global optimum is twice as high as the local optimum. This environment is designed in such a way that the basins of attraction of the two optima have an equal size between the optima, i.e., in the interval $[0, 1]$. This is realized by adjusting σ_{opt} with respect to h . The respective σ_{opt} can be derived numerically. A detailed description of this is presented in Appendix D. In this environment, the adaptive value function changes only once in $t = 10000$. Then, the global optimum z_{opt} changes from 0 to 1 where it remains for the rest of the simulation time. Notice that z_{opt} is an environment parameter.⁵ The population is expected to form a quasi-species around the optimum 0 well before $t = 10000$. The evolutionary dynamics immediately after the change at $t = 10000$ provides insights into how the balance between evolution and learning affects exploration and exploitation in this model.

The population dynamics in Environment 2 are investigated with the following experiment. For a range of constant lifetime settings, evolution is run for 1000 times and in each evolutionary run two performance indicators, namely *discovery time* and *transition time* are measured.

Definition 7.1 (Discovery time). *The time that the population needs to reach the interval $[0.5, 1.5]$ with at least one individual after the environmental change, i.e., the time needed to discover the neighborhood of the global optimum.*

The discovery time can be seen as an indicator for the exploration ability.

Definition 7.2 (Transition time). *The time that the population needs to populate the neighborhood of the global optimum (interval $[0.5, 1.5]$) after the discovery with at least 50 percent of the population.*

The transition time can be seen as an indicator for the exploitation ability.

Figure 7.10 shows the two properties for the tested range of lifetime settings. The discovery time is first decreasing with an increasing lifetime. This is due to an increase in genotypic quasi-species diversity (cf. Section 7.4.1). With increasing diversity it is more likely to discover a neighboring optimum. When the lifetime increases further, the discovery time starts to increase at some point. This phenomenon can be explained as follows: Despite a further increase in genotypic quasi-species diversity, the generation turnover decreases with increasing lifetime, thereby reducing the number of “trials” to find the new optimum. The latter effect seems to be stronger than the former for large lifetimes and vice versa. The discovery time is an indicator for exploration.

The transition time increases monotonically with the lifetime. This is due to the decreasing generation turnover, i.e., the less individuals are replaced the longer it takes to populate the new optimum. The transition time is a measure for exploitation.

If the environment changes repeatedly, the interplay between discovery (exploration) and transition (exploitation) determines the overall adaptation success of the population. The following section investigates this aspect in detail.

⁵ With regard to the nomenclature on page IX, z_{opt} is one dimension of the Environment parameter vector e .




Figure 7.10: *Evolutionary dynamics in Environment 2. The discovery time is an indicator for exploration, where transition time indicates exploitation ability. The discovery time, and the transition time, averaged over 1000 evolutionary runs, suggest that there exists a non-trivial optimal lifetime with regard to the exploration/exploitation balance.*

7.5 Existence of an Optimal Evolution/Learning Balance

This section presents the simulation results of Environments 3 and 4. It is shown that for Environment 3, the optimal adaptation behavior is achieved when no individual learning is included. An increasing degree of learning decreases the degree of evolution and deteriorates the overall adaptation capability of the population. In contrast, it is shown that for Environment 4 an increasing degree of learning at the expense of evolutionary adaptation brings about an adaptational advantage. However, with too much learning, this advantage vanishes. Hence, the optimal balance is given for intermediate degrees of evolution and learning.

7.5.1 Optimality of Pure Evolution (Environment 3)

Figure 7.11 shows Environment 3, that is defined by the time-dependent adaptive value function v_3 ,

$$v_3(z, t) = e^{-(z - z_{\text{opt}}(t))^2} \quad \text{with } z_{\text{opt}}(t) = 0.2 \lfloor t/T \rfloor . \quad (7.10)$$

The uni-modal function that maps phenotype to adaptive value moves gradually in positive z direction where T (the length of the change interval) determines the velocity of this movement. Notice that T is an environment parameter.⁶ The following experiment demonstrates that pure evolution is the best adaptation strategy in Environment 3. Pure evolution ($L = 1$) and coupled evolution/learning ($L=20$) are compared and experiments are done for three different settings of the change interval, $T \in \{1, 10, 100\}$, representing rapidly changing, moderately changing, and slowly changing environments, respectively. Figure 7.13 shows the population mean adaptive value, averaged over 100 independent simulation runs, for the first 400 time

⁶ With regard to the nomenclature on page IX, T is one dimension of the Environment parameter vector e .




Figure 7.11: *Environment 3: A uni-modal Gaussian function that moves gradually in positive z direction, where T (the length of the change interval) determines the velocity of this movement.*




Figure 7.12: *Environment 4: The mapping from phenotype to adaptive value at a time is identical to Environment 2, however, in Environment 4 the optimum changes periodically with an expected change interval of length T .*

units of evolution. In the rapidly changing environment ($L = 1$), the population mean adaptive value is going down to zero quickly in both settings, $L = 1$ and $L = 20$, although slower in case of $L = 1$. On the contrary, in the slowly changing environment ($T = 100$) a high mean adaptive value level of the population is maintained for both, $L = 1$ and $L = 20$. However, in the environment with an intermediate change velocity ($T = 10$) the population mean adaptive value is decreasing in case of coupled evolution/learning ($L = 20$) while it remains at a high level with pure evolution ($L = 1$).

This result is explained as follows: If the environment changes slowly ($T = 100$) both adaptation strategies allow to follow the monotonic movement of the optimum, although small differences in the rate of adaptation to the population with $L = 1$ produces a slightly better adaptive behavior. In the environment with an intermediate change velocity ($T = 10$) the population mean adaptive value is decreasing in case of coupled evolution/learning while it remains at a high level with pure evolution. This means that at some change velocity above $T = 10$, the coupled evolution/learning strategy fails, because the population can not follow the moving optimum. If the dynamics are monotonic as in this example, pure evolution is the best adaptation strategy. A higher degree of (lifetime-induced) diversity is not needed for adaptation, and is actually detrimental because of its negative effect on the exploitation of a new optimum. If the environment changes even quicker as in case of $T = 1$ (top panel in Figure 7.13) neither of the two adaptation strategies allows to follow the optimum, although with pure evolution, the optimum is lost later.




Figure 7.13: Evolution of the population mean adaptive value in Environment 3 for selected settings. If the environment is changing too quickly ($T = 1$), neither of the populations (with $L = 1$ and $L = 20$) can maintain a high mean adaptive value. However, for an intermediate change rate ($T = 10$), the population employing pure evolutionary adaptation ($L = 1$) has an advantage.

7.5.2 Optimality of an Intermediate Degree of Learning (Environm. 4)

Figure 7.12 shows Environment 4 which is defined by the time-dependent adaptive value function v_4 ,

$$v_4(z, t) = h e^{-\left(\frac{z - z_{\text{opt}}(t)}{\sigma_{\text{opt}}}\right)^2} + e^{-\left(\frac{z - (1 - z_{\text{opt}}(t))}{0.25}\right)^2}, \text{ with } h > 1, \quad (7.11)$$

$$\text{and } z_{\text{opt}}(t) = \begin{cases} 0, & \text{if } (z_{\text{opt}}(t-1) = 1 \wedge X_{\text{Uni}[0,1]} < \frac{1}{T}) \vee \\ & (z_{\text{opt}}(t-1) = 0 \wedge X_{\text{Uni}[0,1]} \geq \frac{1}{T}) \\ 1, & \text{otherwise,} \end{cases}$$

where $X_{\text{Uni}[0,1]}$ is a random number drawn from a uniform probability distribution on the interval $[0, 1]$. In Environment 4, the mapping from phenotype to adaptive value at a time is identical to Environment 2, however, in Environment 4 the optimum changes *periodically* with an expected *change interval* of length T . The actual time between changes is uniformly, stochastically distributed and can vary strongly.

The following experiment investigates the evolutionary dynamics in this environment for height factors $h = 2$ and $h = 5$ with a range of constant lifetimes, and for the environmental change intervals 20, 50, 100, 200. The genotype population is initially distributed uniformly on $[-0.5, 1.5]$. The overall adaptation quality is assessed by measuring the mean population fitness over time for 200 independent evolutionary runs. The results are shown in Figure 7.14.

These results show that the slower the environmental change, the higher is the mean adaptive value for the population. For height factor $h = 2$ (left panel), the optimal lifetime is approximately $L = 75$ for an expected change interval of $T = 200$, however, for change intervals lower than that ($T \in \{20, 50, 100\}$), the optimal lifetime is at the boundary of the tested range ($L = 1000$). There seems to be a threshold for the rate of environmental change below which an intermediate lifetime is optimal. For a height factor of 5 (right panel), this threshold lies between an expected change interval of 20 and 50. For a change interval of $T = 20$, a maximally high lifetime $L > 1000$ is optimal, for slower changing environment $L = 25$ (in case of $T = 50$) and $L = 30$ (in case of $T = 100$ and $T = 200$) is optimal. The existence of a threshold for the rate of environment change below which an intermediate lifetime is optimal has been confirmed in several other settings of h .

Figure 7.15 shows the population dynamics of typical runs in Environment 4 for the non-trivial optimal balance between evolution and learning. As an example, the case of height factor $h = 5$ and the change interval $T = 200$ is studied. This corresponds to the dotted line in the right panel of Figure 7.14. Figure 7.15 shows four different degrees of learning (lifetime L) for this setting, a low degree $L = 1$ which produces a rather low mean adaptive value, an intermediate degree $L = 30$ which produces approximately the maximum mean adaptive value, and high degrees $L = 200$ and $L = 1000$ which produce rather low mean adaptive values. The thick gray line shows the trajectory of the global optimum, the thick black dots show the genotype values, and the small gray dots show the phenotype values present in the population at a time.

With pure evolution ($L = 1$) the population quickly converges to the global optimum. The population maintains diversity with mutation-selection balance, however, this degree of diversity is not sufficient to discover another global optimum. This shows that the discovery




Figure 7.14: Mean adaptive value (over time, individuals and simulation runs) for different constant lifetimes in Environment 4 for change intervals $T \in \{20, 50, 100, 200\}$ and height factors 2 (left panel) and 5 (right panel), respectively. There exists an optimal lifetime that depends on the environmental dynamics and height differences between local and global optimum.




Figure 7.15: Typical evolutionary runs in Environment 4. The thick gray line shows the global optimum, the thick black dots show the genotype values, and the small gray dots show the phenotype values present in the population at a time. With $L = 1$ (pure evolution) the population only occasionally discovers a new global optimum. For long lifetimes $L = 200$ and $L = 1000$ the population is not flexible enough to move the majority of individuals to the current global optimum before the next environmental change occurs. Only in the intermediate case of $L = 30$, a good balance between exploration and exploitation is achieved and the population follows the environmental dynamics.

time is too long for the given environmental dynamics. In some evolutionary runs, a population transition occurred occasionally. Next, the cases $L = 200$ and $L = 1000$ are considered.

With a high degree of learning ($L = 200$), evolution has only a weak influence on the overall adaptation process. The genotypes (black dots) remain relatively wide-spread in genotype space and individuals are able to adapt to one of the two optima during the lifetime. Due to the high degree of diversity maintained throughout the simulation time, discovery time is very short. The transition time, however, is too long to move the majority of individuals to the current global optimum before the next environmental change occurs. With $L = 1000$ the slow transition time is even more evident: Because of the extremely low generation turnover, selection takes place rarely in this case, and evolution is virtually disabled.

However, in the intermediate case of $L = 30$, the population follows the environmental dynamics. Evolution and learning are well balanced. As a result, it is possible for the population to discover a new optimum after an environmental change and to exploit it in a relatively short period of time. This gives the population an adaptational advantage over the populations with a too low or too high degree of learning.

In a preliminary study, Environment 4 is defined with deterministic environmental changes. Although in principle this leads to the same conclusion, some interesting phenomena can be observed. Since these findings are not central for the understanding of this chapter, they are relegated to Appendix E. The example of Environment 4 has shown that there are dynamic environments in which adding individual learning to the population can result in better overall population adaptation. However, too much or too little learning results in a worse overall adaptation behavior.

7.6 Summary and Conclusion

A trade-off between individual learning and generation turnover is evident not only in evolutionary computation but also in nature. In the presence of this trade-off, the degree of learning influences the overall adaptation behavior, not only by means of change in selection pressure but also by a decreased generation turnover. Other things equal, a decrease in the generation turnover implies a slow down of genotypic change.

The issue of balancing evolution and learning toward an optimal overall adaptation behavior has been studied with a simulation model. Unlike many other models in which the cost of learning are explicitly assigned, the cost of learning are implicitly given by the associated consumption of computational resources. The model employs two very similar trial-and-error adaptation mechanisms that only differ from each other in that one is applied to the population (evolution) and the other to the individual (learning). The central parameter of the proposed simulation model - individual lifetime - allows to adjust the ratio of computational resources allocated by evolutionary adaptation steps and the ratio allocated by individual learning.

It turned out that an increase in individual lifetime,

- a) allows the population to maintain a higher degree of diversity,

and at the same time

- b) reduces the generation turnover.

The balance between a) and b) affects the exploration/exploitation behavior of the overall adaptation process. Hence, the adjustment of the evolution/learning balance indirectly influences the exploration/exploitation balance of the entire adaptation process.

Using simulations it has been shown that in an environment with monotonic dynamics (Environment 3), pure evolutionary adaptation is the best adaptation strategy. In such environments, diversity is not needed for adaptation and can actually be detrimental because of its negative effect on the exploitation of a new optimum.

A different result has been found in an environment in which the population has to cross fitness valleys repeatedly (Environment 4). There, exploration ability is of importance. It turned out that the learning-induced increase in diversity improves the exploration ability in the right way, such that a coupled evolution/learning strategy has an adaptational advantage over pure evolution. If, however, the degree of learning increases beyond a certain point, thereby increasing exploration ability and reducing exploitation ability, the adaptational advantage vanishes. Thus, an intermediate degree of learning which allows for both, exploration and exploitation of a new optimum, is the optimal adaptation strategy.

There is a good reason to believe that this finding is not limited to environments in which the global optimum switches between only two values: In the case where an intermediate lifetime is optimal, the transition from the old to the new optimum occurs mostly after quasi-species formation, i.e., at a time when the population has completely moved to one optimum and has “forgotten” the old one. Thus, even if future optima appear at different locations, the right balance exploration and exploitation is of great importance.

The control of the exploration/exploitation balance by means of adjusting the learning intensity (lifetime) has not been mentioned in the literature. Unlike the Baldwin effect, an improvement in overall adaptation behavior through individual learning, can only be observed in dynamic environments.

In this chapter, the balance between evolution and learning has been studied from a purely adaptational advantage point of view. It must be mentioned here that in nature other factors and constraints may also play a role.

In nature, the optimal balance can not be set externally, instead it is either constrained by natural laws, an emergent property of the evolution, or a mix of both. Similarly, in evolutionary computation the optimal balance between evolutionary adaptation and learning may not be known in advance and it is then desired that the right balance is found in a self-adaptive way. The following chapter investigates if and under what conditions a near-optimal overall adaptation behavior can emerge in a self-adaptation process.

Self-Adaptation of the Evolution/Learning Balance

The previous chapter has shown that there is a potential advantage of coupling evolution and learning in dynamic environments even in the case of a trade-off between the two means of adaptation. Accounting for the evolutionary dynamics in the long run, there is an optimal balance between evolution and learning. In nature, the optimal balance can not be set externally. Instead it is either constrained by natural laws, an emergent property of the evolution, or a mix of both. Similarly, in computational evolution the optimal balance between evolution and learning may not be known in advance. Therefore, it is desired that the optimal balance emerges from a self-adaptation process. In this chapter, it will be shown under what conditions self-adaptation of the evolution/learning balance can lead a near-optimal overall adaptation behavior. Large parts of this chapter are based on [127].

8.1 Related Work

As reviewed in Section 7.2, no model has yet been published that accounts for a trade-off between evolution and learning. Not surprisingly there is also no work that deals with the evolutionary self-adaptation of this trade-off. However, there are a few papers that are still to a certain degree related to this chapter.

Life history evolution [166, 143] is a branch of evolutionary biology that studies the evolution of the reproductive cycle of individuals including properties like time to maturity, time to first reproduction etc. Only recently, life history evolution has been studied for the first time in evolutionary computation.

In Bullinaria's [17] study on evolution of artificial neural networks for classification tasks, the age of maturity is an important property of lifetime history. Individuals are protected by their parents until they reach the age of maturity. Testing different ages of maturity, it turns out that lifetime learning is more effective, the higher the age of maturity. Despite some cost of late maturity for both parents and offspring relatively high ages of maturity associated

with a high degree of learning evolve. Noteworthy, the environment in Bullinaria's study is stationary.

8.2 Extension of the Analysis Model

The analysis model used here is an extended version of the one introduced in Section 7.3, where an individual was formally defined in Equations 7.1 to 7.4. Individual lifetime L is the central parameter that determines the balance of evolution and learning. In the extended model, L is now individually encoded in the genotype. Hence, the genotype x is no longer given by a simple scalar $x \in \mathbb{R}$. Instead, it is composed of one variable that encodes the innate phenotype z_0 and one variable that encodes the lifetime L , i.e.

$$x = (z_0, L) . \quad (8.1)$$

The mutational change from parent genotype x to offspring genotype x' ,

$$x = (z_0, L) \mapsto (z'_0, L') = x' \quad (8.2)$$

is defined by

$$z'_0 = z_0 + X_{\phi(0, \sigma_G)} , \quad (8.3)$$

where $X_{\phi(0, \sigma_G)}$ is a Gaussian random number (σ_G is also known as adaptation step size, cf. Equation 7.2), and

$$L' = \begin{cases} L + 1 , & \text{if } 0.00 \leq X_{\text{Uni}[0,1]} < 0.05 \\ L - 1 , & \text{if } 0.05 \leq X_{\text{Uni}[0,1]} < 0.10 \wedge L > 1 \\ L & , \text{ otherwise ,} \end{cases} \quad (8.4)$$

where $X_{\text{Uni}[0,1]}$ is a random number drawn from a uniform probability distribution on the interval $[0, 1]$. Equation 7.3 which describes the reproduction transition is still valid. Thus, the individual lifetime L can evolve, thereby enabling self-adaptation of the evolution/learning trade-off by means of mutation and selection.

8.3 An Initial Experiment of Lifetime Evolution

The extended model is applied to Environment 4 of Chapter 7 as defined with a change interval of $T = 200$. Adaptation step-sizes are again set to $\sigma_G = \sigma_P = 0.01$. Recall that the optimal balance between evolution and learning has been found at a lifetime of approximately $L = 30$.

According to the formal definition in Section 8.2 the lifetime is encoded in the genotype of each individual. In the initial population, the lifetimes are assigned randomly to the individuals with respect to a uniform probability distribution over $[1, 5]$. Thus, the population starts with a low expected lifetime. (Later, evolution that starts with a high expected lifetime is studied as well.) Lifetime mutation is realized according to Equation 8.4.

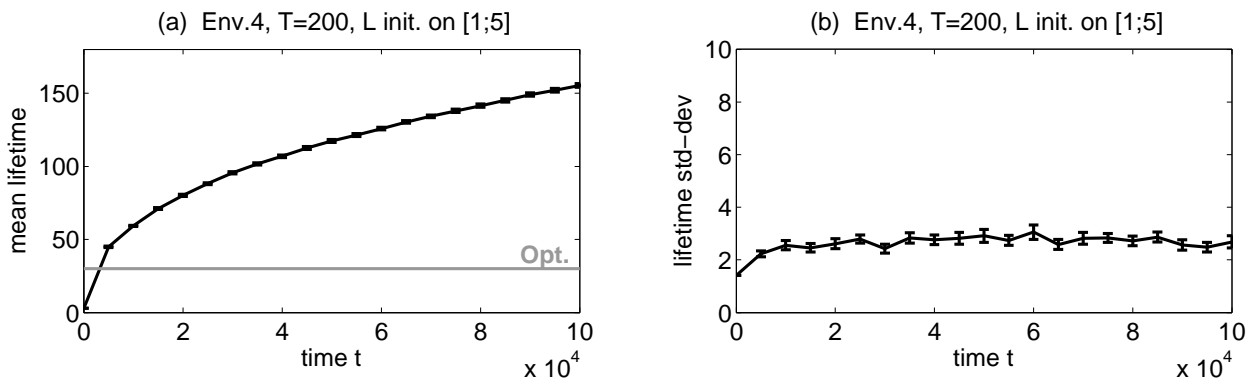


Figure 8.1: Evolution of lifetime in Environment 4 ($h = 5$) at a change interval of $T = 200$. (a) shows the mean lifetime in the population over time and (b) the standard deviation of the lifetime present in the population. Error bars indicate the standard-error over 30 independent simulation runs. According to Section 7.5.2 the optimal lifetime is 30, as marked as a gray line in (a). Simply encoding lifetime parameter L leads to an unbounded increase of the average lifetime.

Figure 8.1 shows the result of the first 100000 time steps of 30 independent simulation runs. Figure 8.1(a) shows the evolution of population mean lifetime, averaged over the 30 simulation runs with error bars.¹ Figure 8.1(b) shows the corresponding standard deviation of the lifetime within the population, again averaged over the 30 simulation runs with error bars. The mean lifetime increases to a value far beyond the optimal lifetime of 30 and seems to grow infinitely. The variation of lifetime within the population is relatively small, cf. Figure 8.1(b).

Apparently the optimal lifetime does not emerge from a self-adaptation process. How can the infinite growth of lifetime be explained?

The following theoretical considerations illuminate this issue. It is assumed in the following that a population of n individuals with genotypes $\{x_i\}_{i=1\dots n}$ is given. The corresponding phenotype changes throughout the lifetime and produces (for each individual) a vector of realized adaptive values. $\bar{v}(x_i)$ denotes the average adaptive value of individual x_i (over lifetime).

\bar{L} denotes the mean lifetime of all individuals in the population. The average generation turnover is n/\bar{L} . Hence the average expected number of offspring of individual with genotype x_i at a time, is calculated as

$$\frac{n}{\bar{L}} \frac{\bar{v}(x_i)}{\sum_{j=1}^n \bar{v}(x_j)} = \frac{n}{\bar{L}} \frac{\bar{v}(x_i)}{n \bar{v}} = \frac{\bar{v}(x_i)}{\bar{L} \bar{v}}, \quad (8.5)$$

where \bar{v} denotes the mean adaptive value of the population during the lifetime of individual x_i . The expected number of offspring $w(x_i)$ of individual x_i over its entire lifetime L_i is given by

$$w(x_i) = L_i \frac{\bar{v}(x_i)}{\bar{L} \bar{v}}. \quad (8.6)$$

This equation shows that the expected number of offspring increases with lifetime L_i . This means individuals with a longer lifetime have an implicit reproductive advantage. In short,

¹Error bars have the length of +/- one standard error.

long living individuals reproduce more because they have more opportunities to do so. Hence in the long run, individuals with extremely long lifetimes overwhelm. As shown earlier, extremely long lifetimes produce an adaptational disadvantage with respect to the overall population behavior. Moreover, long lifetimes are biologically infeasible.

The evolution of a very long lifetime can be attributed to the fact that there is no *individual* trade-off between average reproduction probability and the lifetime of individuals. In nature, such a trade-off is evident as reviewed in Section 7.1. In the following, it is shown how a trade-off between reproduction and lifetime can be implemented in the proposed model and how it influences the evolution of lifetime.

8.4 Lifetime Evolution with a Trade-Off between Reproduction and Lifetime

In the previous section, it has been shown that the reproductive advantage increases with lifetime in absence of a negative effect of lifetime on reproduction (Equation 8.6). In order to neutralize this undesired effect, a trade-off between average reproduction probability and lifetime is introduced. Lifetime L_i reduces the probability to reproduce as follows

$$V(x_i) = \frac{v(x_i)}{L_i} . \quad (8.7)$$

Function $V(x_i)$ denotes the new adaptive value of x_i that accounts for the trade-off. The lifetime mean of V of an individual x_i , as denoted $\bar{V}(x_i)$, is calculated as

$$\bar{V}(x_i) = \frac{\bar{v}(x_i)}{L_i} . \quad (8.8)$$

Recall that $\bar{v}(x_i)$ denotes lifetime mean of the original adaptive value, v , of an individual x_i . In analogy to Equation 8.6, and accounting for Equation 8.8, the expected number of offspring in presence of a trade-off, $W(x_i)$, is derived as

$$W(x_i) = L_i \frac{\bar{V}(x_i)}{\bar{L} \bar{V}} = \frac{\bar{v}(x_i)}{\bar{L} \bar{V}} . \quad (8.9)$$

The function $\bar{V}(x_i)$ denotes the population mean with respect to the new adaptive value V (over the entire lifetime of individual x_i). $W(x_i)$ denotes the expected number of offsprings of x_i in presence of the proposed trade-off.

Equation 8.9 shows that the expected number of offspring is independent of the individual lifetime L_i , if the trade-off between reproduction and individual adaptation as defined in Equation 8.7 is taken into account. A shorter lifetime increases the probability to reproduce at a time. On the other hand, individuals with high lifetime have more reproduction opportunities. However, the reproduction trade-off ensures that no overall advantage arises from a certain lifetime.

8.4.1 Evolution of the Optimal Lifetime in Environment 4

With this model modification and under otherwise identical conditions as in Section 8.3, simulated evolution is repeated for Environment 4 of Chapter 7. The results are presented in Figure 8.2. Recall the optimal lifetime of 30 as found in Section 7.5.2 where the lifetime was predefined and kept constant during evolution. First, as shown in Figure 8.2(a) evolution starts with a low mean lifetime, initialized randomly on $[1, 5]$ w.r.t. a uniform probability distribution, i.e., with an average lifetime of 3. We see that in the presence of the reproduction/lifetime trade-off, a near-optimal lifetime between 30 and 35 evolves. In a follow-up experiment, as shown in Figure 8.2(c), the population starts to evolve with high lifetimes, initialized on $[30, 70]$, i.e., with an average lifetime of 50. Again, the population evolves a near-optimal lifetime. The results show that independent of the starting conditions, self-adaptation toward a near-optimal evolution/learning balance works robustly.

For the simulations presented in Figures 8.2(a) and (c), the variation of the lifetime present in the population (measured as standard deviation) is also shown in Figures 8.2(b) and (d), respectively. The variation is low in both experiments, indicating that there is a stable population movement toward the optimal lifetime, and that the population mean lifetime does not “average out” the actual population dynamics.

In another experiment with the model that incorporates the trade-off, the population is initialized with lifetimes uniformly distributed on $[1, 5]$. However, now the environment changes on average every 20 time units. The result is presented in Figure 8.3. Now, the average population evolves a lifetime of 100 during the first 100000 time steps of evolution and even longer lifetimes in succeeding time steps (not shown). This corresponds to the findings of Section 7.5.2 where a very long lifetime (larger 1000) turned out to be optimal if the environment changes with an expected change interval of 20.

8.4.2 Evolution of the Optimal Lifetime in Environment 3

Simulated evolution is also repeated for Environment 3 of Chapter 7 using the model that accounts for the reproduction/lifetime trade-off. Here as well, simulation parameters are set as in Section 7.5.1.

Recall that in Environment 3, the adaptational challenge is to follow a quickly moving optimum. In Section 7.5.1, it is shown that pure population adaptation, i.e., $L = 1$ is the best adaptation strategy for this type of environmental dynamics (cf. Figure 7.13). Figure 8.4 shows the evolution of lifetime in Environment 3 with an environmental change interval of 10. In this example as well, a near-optimal degree of population adaptation (near $L = 1$) emerges from a self-adaptation process.

8.5 Summary and Conclusion

The preceding chapter has shown that depending on the environmental dynamics there exists a certain balance of evolution and learning that is optimal with respect to the mean population fitness measured over time. Following these conclusions this chapter has investigated if an optimal or near-optimal balance can emerge by means of self-adaptation.

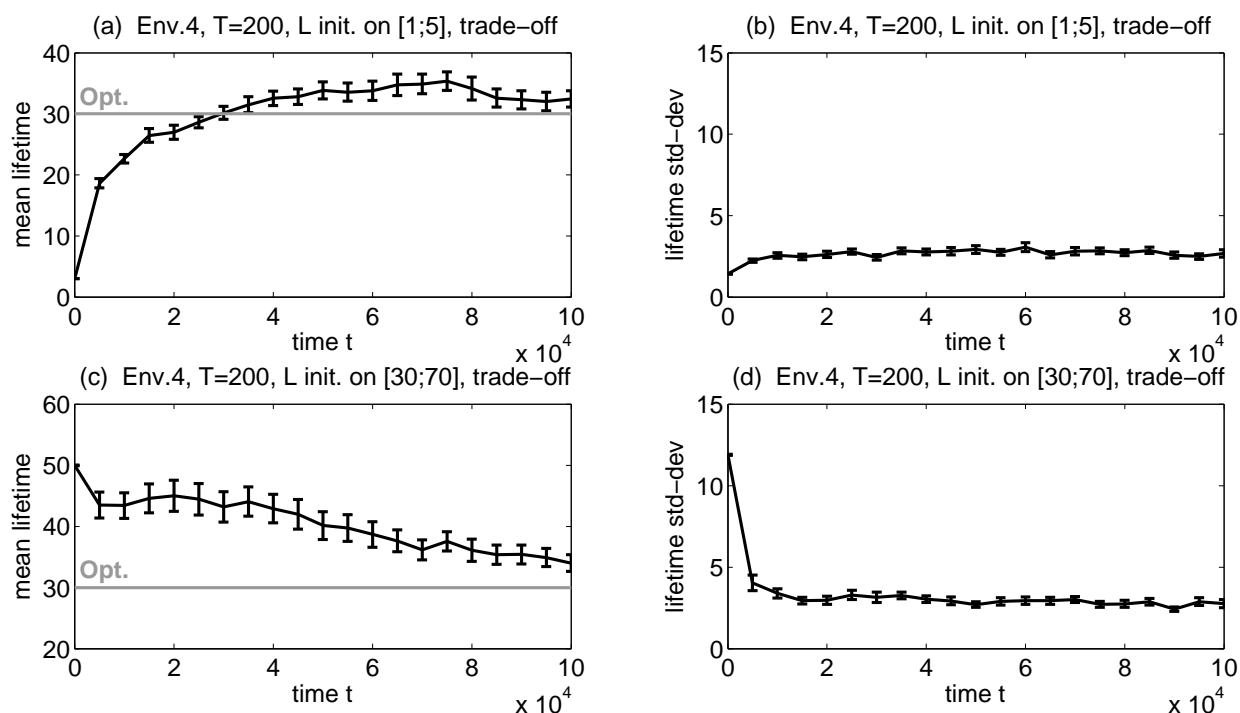


Figure 8.2: Evolution of lifetime in Environment 4 ($h = 5$) at a change interval of $T = 200$ in presence of a reproduction/lifetime trade-off. (a) shows the mean lifetime in the population over time where evolution starts at low mean lifetime and (b) the corresponding standard deviation of the lifetime present in the population. (c) shows the mean lifetime in the population where evolution starts at high mean lifetime and (d) the corresponding standard deviation. Error bars indicate the standard-error over 30 independent simulation runs. According to Section 7.5.2 the optimal lifetime is 30, as marked as a gray line in (a) and (c). Independent of the initialization the near-optimal lifetime evolves robustly.




Figure 8.3: Evolution of lifetime in Environment 4 ($h = 5$) at a change interval of $T = 20$ in presence of a reproduction/lifetime trade-off. (a) shows the mean lifetime in the population over time and (b) the corresponding standard deviation of the lifetime present in the population. Error bars indicate the standard-error over 30 independent simulation runs. According to Section 7.5.2 the optimal lifetime is larger than 1000 as indicated in (a). The lifetime evolves toward large values, i.e., potentially toward the optimum.




Figure 8.4: Evolution of lifetime in Environment 3 at a change interval of $T = 10$ in presence of a reproduction/lifetime trade-off. (a) shows the mean lifetime in the population over time and (b) the standard deviation of the lifetime present in the population. Error bars indicate the standard-error over 30 independent simulation runs. According to Section 7.5.1 the optimal lifetime is 1, as marked as a gray line in (a). The lifetime evolves toward a near-optimal value.

In the analysis model of evolution, individual lifetime L is the crucial parameter to distribute adaptation effort between the level of evolutionary adaptation and learning. Simply encoding L in the individual genotype and allowing it to evolve in a mutation-selection cycle results in the evolution of an infinitely increasing lifetime. This can be explained by the absence of a negative effect of lifetime on reproductive success per time. Incorporating a trade-off between lifetime and reproduction per time that can be found similarly in natural organisms, disables the bias toward long lifetimes and a near optimal balance of evolution and learning emerges from a self-adaptation process.

Undoubtedly, there may be other factors and constraints in nature which determine the average individual lifetime in a species. However, this and the preceding chapter provide a purely adaptational argument for the evolution of a certain balance between evolution and learning.

The results are even more interesting from a computational intelligence point of view. The preceding chapter has shown that the balance between evolution and learning influences the exploration/exploitation balance of the overall adaptation process. Hence, the extended model presented in this chapter, provides a means for self-adaptation of the exploration/exploitation balance under changing environmental conditions.

Conclusion and Outlook

“Properly speaking, such a work is never finished; one must declare it so when, according to time and circumstances, one has done one’s best.”

Johann Wolfgang von Goethe, Italian Journey

Evolution and learning are the two major mechanisms in natural adaptation. The interplay between these two mechanisms allows populations of biological organisms to adapt to various changing environmental conditions, a capability that is also demanded in today’s and tomorrow’s large-scale digital processing systems. Many facets of the dynamics that arise from the interplay between evolution and learning have not been understood, yet. This thesis has studied some of these aspects and developed models that may serve as a basis for the design of computational systems that employ nature-inspired adaptation mechanisms. This chapter summarizes the conclusions (Section 9.1) and contributions (Section 9.2) from the various studies of this thesis, and suggests future research steps (Section 9.3).

9.1 Conclusion

The core of this thesis is the development of the gain function framework which provides a general explanation under what conditions learning accelerates or decelerates evolutionary change. In its simple form, the gain function is formulated in terms of the relative fitness gain of an individual with respect to the absence of learning. In its extended formulation, the influence of an increase in a learning parameter is taken into account. The gain function considers the influence of learning on selection pressure. The basic idea is that learning accelerates evolutionary change if genetically strong individuals benefit proportionally more from learning than weak individuals. This case is indicated by a positive gain function derivative. The acceleration effect of learning is interpreted as the occurrence of the *Baldwin*

effect by many authors. Correspondingly, a negative gain function derivative indicates a scenario in which genetically weak individuals benefit more from learning than the strong ones. This leads to a deceleration of evolutionary change, an effect that has become known as *hiding effect* in recent years. The extended gain function framework considers the influence of a parameter on the mapping from genotype to fitness. Originally, this parameter was interpreted as a learning parameter. However, the mathematical treatment does not limit the interpretation of the parameter. It may as well be interpreted as a parameter that influences development (ontogenesis).

Although the formal derivation of the gain function naturally imposes some simplifying assumptions it has been employed successfully in a variety of contexts in this thesis. It explains in what situations learning accelerates or decelerates evolution. For example, if learning shifts each individual by a positive constant distance in phenotype space toward higher fitness, this type of learning will accelerate evolutionary change where the logarithm of the fitness function (mapping from genotype to fitness) is convex, and decelerate it where the logarithm of the fitness function is concave. With a similar analysis it has been proven that noise in the genotype-phenotype-mapping can actually accelerate evolutionary change - a somewhat non-intuitive result.

The model of Hinton and Nowlan [64] is perhaps the most prominent simulation model of evolution and learning. In the literature their results have been explained with different arguments. However, the gain function analysis has shown that a selection pressure argument is sufficient to explain their main result that learning accelerates evolution in the simulation model.

The gain function perspective allowed an interesting reconsideration of Mery and Kawecki's [107] biological experiment with fruit flies. The evolutionary dynamics as recorded in the experimental data allowed to draw some conclusions on the effect of individual learning on fitness. If the starting point of learning is far away from the learning goal, learning is not very beneficial. Learning is not very beneficial either, if it starts close to the learning target. The maximum benefit is achieved for a starting point with intermediate distance to the target. Interestingly, this result corresponds to some theories on animal and insect learning [137, 122].

Beyond the formal derivation of the gain function, the dynamics of evolution and learning have also been studied via simulation in an environment with limited availability of computational resources. Under these conditions the rate of evolutionary adaptation and the intensity of individual learning need to be balanced. The proposed model allows to specify the distribution of the computational resource consumption between evolution and learning. It turned out that balancing evolution and learning is a means to adjust the exploration/exploitation behavior of the overall adaptation process. The optimal balance is influenced by the type and rate of environmental change.

In nature, the optimal balance can not be set externally. To some extent it may be constrained by natural laws and may thus be an emergent property of evolution. Similarly, in evolutionary computation the optimal balance between evolution and learning may not be known in advance. Therefore, it is desired that the optimal balance emerges from a self-adaptation process. It turned out that self-adaptation of the evolution/learning balance can be achieved by genetically encoding an individual's lifetime (learning time) and letting it evolve. However, it is required to incorporate a biologically plausible trade-off between lifetime and reproduction.

Summarizing, this thesis has significantly deepened the understanding of the biological interrelationship between evolution and learning. In fact, parts of this work have been published [131] in the same journal that also published what later became known as the *Baldwin effect* [7]. As shown by the studies on the influence of learning on evolution this contribution to evolutionary biology and computational biomodelling provides an important improvement for the anticipated transfer of these biological principles to the design of information processing systems which need to adapt to their dynamically changing operating environment. Also steps toward the transfer have been taken by employing standard EA techniques and by addressing the issue of computational resource limitations.

9.2 List of Contributions

In the following, the main contributions of this thesis are explicitly formulated.

Explanation of the adaptational disadvantage of Lamarckism in rapidly changing environments

As shown by others, Lamarckism has an adaptational disadvantage in rapidly changing environments. By using a simplified model (Chapter 3), this thesis has provided a simple explanation for this result: The disadvantage in rapidly changing environments is explained by the movement of the mean genotype. With Lamarckian inheritance, genotype movement is faster than with genetic mutation alone. Though this may be helpful in the short run, it can be detrimental in the long run under dynamic environmental conditions. The near-optimal degree of Lamarckism with respect to the rate of environmental change can be produced by an evolutionary self-adaptation process.

Formulation and proof of the gain function as a mathematical framework to predict the influence of learning on the rate of evolution

The gain function is formulated in terms of the effect of learning on the mapping from genotype to fitness (Chapter 4). For the sake of mathematical analysis, genotype and phenotype are represented by a scalar value. In its initial formulation, the gain function can be used to predict the effect of adding individual learning to the evolutionary process. In its extended formulation, it can be used to predict how a change in a learning parameter affects the rate of evolution. In both versions, the gain function analysis looks at the effect of learning and does not require to consider a particular learning scheme or algorithm. All that is needed is to know how learning influences fitness. The gain function makes exact short-term predictions on the evolutionary dynamics. A simulation study has demonstrated that the gain function is also useful to approximately describe the long term dynamics of the population, e.g., in the case that an acceleration phase is followed by a deceleration phase (Section 5.4). Despite its generality, the gain function has some limitations: It is expectation based and does not account for unlikely stochastic events (cf. Section 6.5).

Identification of the conditions for learning-induced acceleration or deceleration for typical forms of learning

The gain function framework has been applied to the identification of conditions under which typical forms of learning accelerate or decelerate evolution, in particular: a) directional learning accelerates evolution, if the logarithm of the function that maps phenotype to fitness is convex, and decelerates it, if the logarithm is concave (Section 5.1.1); b) it has been mathematically proven that noise in the genotype-phenotype-mapping can accelerate the evolutionary process - a somewhat non-intuitive result (Section 5.1.2); c) the decomposition of individual fitness into an innate and a learning component revealed that learning accelerates evolution only if the fitness attributable to the learning component increase faster than the fitness attributable to the innate component (toward the optimum) (Section 5.2); d) continual fitness assessment may revert the influence of learning on evolution compared to the case of posthumous fitness assessment, if learning curves have different shapes for innately weak and strong individuals (Section 5.3). In order to focus on learning, the innate phenotype value is directly specified by the genotype in these models.

Theoretical underpinning of various studies of coupled evolution and learning

Gain function analyses have been used to produce a theoretical underpinning of several studies of coupled evolution and learning (Chapter 6), namely a) Hinton and Nowlan's simulation study [64], b) Papaj's computational biology experiment [133], and c) Cavallie and Feldmann's [18], Anderson's [5], and Ancel's [3] analytical treatment of the influence of developmental noise on evolution. The gain function has also been utilized to shed some light on evolutionary data of a biological experiment with fruit flies [107] and in particular to connect the evolutionary results to theories on animal and insect learning.

Discovery of a new type of adaptational advantage in presence of a resource-conflict between evolution and learning

Under computational resource limitations, the rate of evolutionary adaptation and the intensity of individual learning need to be balanced. A model has been proposed that allows to specify the distribution of the computational resource consumption between evolution and learning. To a certain extent a similar trade-off can be found in nature. It turned out that an increase in the degree of learning, a) allowed the population to maintain a higher degree of diversity and at the same time b) reduced the generation turnover. The interplay between a) and b) affects the exploration/exploitation behavior of the overall adaptation process. Hence, the adjustment of the evolution/learning balance indirectly influences the exploration/exploitation balance of the entire adaptation process. Finally, the optimal balance is influenced by the type and rate of environmental change. Examples have shown that only for a certain evolution/learning balance the population can cope with the environmental dynamics.

Demonstration that biologically-plausible reproduction constraints allow successful self-adaptation of the evolution/learning balance

Simply encoding an individual's lifetime in its genotype and evolving it by means of mutation and selection leads to the undesired infinite growth of lifetime in the population. Incorporating an individual trade-off between reproduction probability and lifetime creates the conditions for successful self-adaptation of the lifetime (learning time) toward the optimal overall adaptation behavior.

9.3 Outlook

Despite of the contributions of this thesis, many facets of the dynamics that arise from the interplay between the two adaptation mechanisms remain only partially understood.

The gain function is designed to predict the influence of learning on selection pressure. This might be extended toward other aspects.

For example, in all setups in the simulation study in Chapter 7, an increased degree of individual learning caused a larger diversity which was partly due to the change in the effective fitness landscape. In all examples where an increased diversity has been observed, the gain function has a negative derivative. Since a decreasing gain function reflects a learning-induced reduction of selection pressure, it seems intuitively clear that it also leads to an increased diversity. However, unless proven mathematically, there is no guarantee that this intuition is correct. Thus, a mathematical formulation and proof of a "diversity gain function" is a natural extension of this thesis. The original gain function considers the change in the mean genotype of the population. A starting point for the development of the diversity gain function could be to analyze the change in the population's variance of the genotype. The variance can be considered as a first approximation of diversity. However, as a final diversity measure the variance is inappropriate, since a decrease in the population's genotype richness does not imply an increase in variance.

Since the gain function is based on the analysis of the expectation population dynamics and does not account for the variance of the population movement it does not allow to make predictions on the influence of learning on the time needed to cross a fitness valley toward a region with higher fitness. Such a prediction cannot be made expectation-based since fitness valley crossing requires an "unlikely" event. A stochastic analysis seems more appropriate to predict the time to cross a fitness valley. It turned out that during the work on this PhD dissertation, a first step in this direction has been taken. In his PhD thesis, Elhanan Borenstein [12] presents a heuristic analysis tool for the estimation of the fitness valley crossing time. Combining stochastic analysis with the basic idea of the gain function seems a promising approach for future research.

Another issue of future research should be the further study of non-monotonic gain functions. Although not yet shown mathematically it is safe to say that learning probably accelerates evolution if the vast majority of the individuals of a population is located in a fitness landscape region with increasing gain function. The vector of gain functions derivatives (one entry corresponds to one individual) may provide valuable information for both cases a monotonic and a non-monotonic gain function.

Certainly, future work should consider the further translation of biological adaptation processes to digital processing systems. For most real-world scenarios, it may be appropriate to employ different adaptation techniques on the level of evolutionary adaptation and on the level of individual learning. This takes up the idea of memetic or hybrid evolutionary algorithms that are largely motivated by the benefits that arise from coupling coarse-grained with fine-grained search. These algorithms are mainly applied to the optimization of stationary functions, yet. The results of this thesis encourage the study of hybrid algorithms in changing environments such as optimization of dynamic objective functions, and control. An important aspect of these applications will be the roles of online and offline evolution and learning. An integration of these aspects into a theoretical framework is a subject of future work.

Geometric Explanation for the Fitness Valley in Experiment 1 of
Chapter 3

Experiment 1 has shown that for a given T , the population fitness over time is minimal for an intermediate λ . A possible explanation is outlined in the following: With a very low mutation rate it is assumed that genotype changes within time T are mainly induced by Lamarckism and that mutation-induced random genetic changes are negligible. Furthermore, it is assumed that the population fitness is well represented by the expected fitness of the population mean genotype. Thus, population fitness can be expressed w.r.t. the population mean distance to the optimal genotype, which is denoted as d . Assume that initially $d = 0.5$ and between two environmental changes (within T), this distance is reduced by a distance of D , where D depends on the level of Lamarckism λ and the learning parameter a , i.e., $D(\lambda, a)$. In the simplified model of this chapter, it is known that $\partial D/\partial a \geq 0$ and more importantly for this analysis $\partial D/\partial \lambda \geq 0$, i.e., D is increasing with λ . Let us first consider the case where $(0 < D \leq 0.5)$ such that the population never reaches the optimum within T or just immediately before the environmental change at T , e.g., because λ is too small: At the time, just before an environmental change occurs, the population has a distance of $d = 0.5 - D$ to the optimum. Immediately after the environmental change, this distance becomes $d = 0.5 + D$ since the optimal genotype has changed (from 0 to 1 or from 1 to 0). Since the population always moves back and forth between these two states, the expected population fitness over time is approximately

$$\bar{f}(D, a) = \frac{1}{D} \int_{0.5-D}^{0.5+D} f_{\text{exp}}(d, a) \, dd, \quad (\text{A.1})$$

where the expected fitness of d is $f_{\text{exp}}(d, a) = 2 - \phi(d, a)$ (cf. equations 3.3 and 3.4). This assumes that the fit phenotype's fitness is twice the unfit phenotype's fitness. Equation A.1




Figure A.1: Geometrical explanation for the population fitness valley for intermediate λ at intermediate T encountered in Experiment 1 (cf. Fig 3.3). The figure shows Equation A.2 with $a = 0.5$. A population fitness minimum occurs at $D = 0.5$ (cf. text).

can be reformulated with straight-forward calculations. Substituting n for $(1/(1 - a))$, we obtain

$$\bar{f}(D, n) = \begin{cases} 2 + \frac{(0.5-D)^{n+1}}{2D(n+1)} - (0.5 + D)^{n+1} & \text{if } 0 < D \leq 0.5 \\ 2 + \frac{1}{2D} \frac{2n+1}{n+1} - \frac{1}{D} & \text{if } 0.5 < D \leq 1 \\ 2 - 0.5^n & \text{if } D = 0. \end{cases} \quad (\text{A.2})$$

The first case ($0 < D \leq 0.5$) corresponds to the above described scenario, where the population never reaches the optimum within T . In the second case ($0.5 < D \leq 1$), the population reaches the optimal genotype within T and stays there until the next environmental change occurs (having the maximum fitness of 2 during this time). Thus, for ($0.5 < D \leq 1$), we obtain $(0.5/D) \cdot \bar{f}(0.5, n) + ((D - 0.5)/D) \cdot 2$, which produces the second case of Equation A.2 after some straightforward calculations. The third case ($D = 0$) corresponds to $\lambda = 0$ (no Lamarckism). Here, the population fitness over time is simply the expected fitness of $d = 0.5$, i.e., the population does not move. Figure A.1 illustrates Equation A.2 for $L = 0.5$. It shows a minimum at $D = 0.5$. For a given constant L , D only depends on λ and we know that D is increasing with λ . Thus, the population fitness \bar{f} is decreasing for small λ and increasing for large λ , producing a minimum for intermediate λ . This provides a possible explanation for the occurrence of the fitness valley for intermediate λ at intermediate T in experiments 1 and 2.

To summarize the main argument of this geometrical explanation: With a low mutation rate, the population's mean genotype movement mainly depends on the level of Lamarckism, i.e., Lamarckism allows quick genotype movement. A (Lamarckism-induced) quickly moving population may be less fit than a population that is not or hardly moving (without Lamarckism): While a quickly moving population has the advantage of approaching a recently changed fitness optimum, it potentially has an adaptational disadvantage when the next environmental change occurs, since it is farther away from the new optimum than the population that has moved less. In the suggested model this disadvantage indeed occurs, and the disadvantage is even larger than the adaptational advantage of approaching a new optimum. Thus, the population fitness is decreasing for increasing level of Lamarckism. However, if the level

of Lamarckism further increases and exceeds a certain threshold, the population can very quickly move to the new optimum and stay there at a high fitness level (until the next environmental change occurs). Thus, at intermediate levels of Lamarckism, the population fitness is increasing with the level of Lamarckism.

APPENDIX B

Proof of Equation 5.16

This appendix proves Equation 5.16 which is rewritten here as two equations

$$\forall x : f(x) > 0 \wedge f'(x) > 0 \wedge f''(x) > 0 \wedge f'''(x) \leq 0 \Rightarrow g'(x) < 0 \quad (\text{B.1})$$

and

$$\forall x : f(x) > 0 \wedge f'(x) > 0 \wedge f''(x) < 0 \wedge f'''(x) \geq 0 \Rightarrow g'(x) > 0 \quad (\text{B.2})$$

Recalling Equation 5.9 the expected fitness of an individual can be written as (and then reformulated)

$$\begin{aligned} & \bar{f}(l_\varepsilon(x)) \\ &= \int_{-\varepsilon_{\max}}^{+\varepsilon_{\max}} p(\varepsilon) f(x + \varepsilon) d\varepsilon \\ &= \int_0^{+\varepsilon_{\max}} p(\varepsilon) (f(x + \varepsilon) + f(x - \varepsilon)) d\varepsilon \\ &= f(x) - \int_0^{+\varepsilon_{\max}} 2p(\varepsilon) f(x) d\varepsilon + \int_0^{+\varepsilon_{\max}} p(\varepsilon) (f(x + \varepsilon) + f(x - \varepsilon)) d\varepsilon \\ &= f(x) + \int_0^{+\varepsilon_{\max}} p(\varepsilon) h(x, \varepsilon) d\varepsilon . \end{aligned} \quad (\text{B.3})$$

with

$$h(x, \varepsilon) = f(x + \varepsilon) + f(x - \varepsilon) - 2f(x) . \quad (\text{B.4})$$

With this reformulation,

$$\bar{f}(l_\varepsilon(x)) - f(x) = \int_0^{+\varepsilon_{\max}} p(\varepsilon) h(x, \varepsilon) d\varepsilon , \quad (\text{B.5})$$

it is first shown that $\text{sign}(h(x))$ depends on $\text{sign}(f''(x))$, in particular

$$\forall x : f''(x) \begin{cases} > 0 \Rightarrow h(x) > 0 \\ < 0 \Rightarrow h(x) < 0 \\ = 0 \Rightarrow h(x) = 0 . \end{cases} \quad (\text{B.6})$$

Consider $f'' > 0$ first: For all functions with $f''(x) > 0$ for all x (convex functions) it is known that

$$\forall \lambda \in]0, 1[: x_1 < x_2 \Rightarrow f(\lambda x_1 + (1 - \lambda)x_2) - (\lambda f(x_1) + (1 - \lambda)f(x_2)) < 0 . \quad (\text{B.7})$$

With $\epsilon > 0$, substituting $x_1 = x - \epsilon$, $x_2 = x + \epsilon$, and for the special case of $\lambda = 0.5$,

$$\begin{aligned} \forall x : f''(x) > 0 &\Rightarrow f(x) - (0.5f(x - \epsilon) + 0.5f(x + \epsilon)) < 0 \\ &\Leftrightarrow f(x + \epsilon) + f(x - \epsilon) - 2f(x) > 0 \\ &\Leftrightarrow h(x) > 0 , \end{aligned} \quad (\text{B.8})$$

which proves the first case of Equation B.6. The second case can be proven in an analogous way. To proof the third case of Equation B.6 ($\forall x : f''(x) = 0$), $f(x)$ can be rewritten as

$$f(x) = ax + b \Rightarrow h(x) = 0 . \quad (\text{B.9})$$

Thus, Equation B.6 is proven.

With equations B.5 and B.6 and because for positive (negative, zero) $h(x, \epsilon)$ the corresponding $\int_0^{+\epsilon_{\max}} p(\epsilon)h(x, \epsilon) d\epsilon$ is positive (negative, zero), too, we obtain

$$\forall x : f''(x) \begin{cases} > 0 \Rightarrow \bar{f}(l_\epsilon(x)) - f(x) > 0 \\ < 0 \Rightarrow \bar{f}(l_\epsilon(x)) - f(x) < 0 \\ = 0 \Rightarrow \bar{f}(l_\epsilon(x)) - f(x) = 0 , \end{cases} \quad (\text{B.10})$$

which which will be used in the final step of the proof.

Using the convexity equation (Equation B.7) in an analogous way as in Equation B.8 for $f'(x)$ and $f'''(x)$, it can be derived

$$\forall x : f'''(x) \begin{cases} > 0 \Rightarrow h'(x) > 0 \\ < 0 \Rightarrow h'(x) < 0 \\ = 0 \Rightarrow h'(x) = 0 . \end{cases} \quad (\text{B.11})$$

Note that for the case of $f''' = 0$, f can be written in the form $f(x) = ax^2 + bx + c$ and we obtain $h(x, \epsilon) = 2a\epsilon^2$ with $\frac{\partial h(x, \epsilon)}{\partial x} = 0$.

Thus, $h(x)$ is monotonic in x and

$$\forall x : f'''(x) \begin{cases} > 0 \Rightarrow \frac{\partial}{\partial x} \int_0^{+\epsilon_{\max}} p(\epsilon)h(x, \epsilon) d\epsilon > 0 \\ < 0 \Rightarrow \frac{\partial}{\partial x} \int_0^{+\epsilon_{\max}} p(\epsilon)h(x, \epsilon) d\epsilon < 0 \\ = 0 \Rightarrow \frac{\partial}{\partial x} \int_0^{+\epsilon_{\max}} p(\epsilon)h(x, \epsilon) d\epsilon = 0 . \end{cases} \quad (\text{B.12})$$

Since with Equation B.5,

$$\frac{\partial}{\partial x} \int_0^{+\varepsilon_{\max}} p(\varepsilon)h(x, \varepsilon) d\varepsilon = \frac{\partial}{\partial x} (\bar{f}(l_\varepsilon(x)) - f(x)) = (\bar{f}(l_\varepsilon(x)))' - f'(x) \quad (\text{B.13})$$

we obtain

$$\forall x : f'''(x) \begin{cases} > 0 \Rightarrow (\bar{f}(l_\varepsilon(x)))' - f'(x) > 0 \\ < 0 \Rightarrow (\bar{f}(l_\varepsilon(x)))' - f'(x) < 0 \\ = 0 \Rightarrow (\bar{f}(l_\varepsilon(x)))' - f'(x) = 0, \end{cases} \quad (\text{B.14})$$

which will also be used in the final step of the proof.

The preceding equations, in particular equations B.10 and B.14 are now used to proof Equation 5.16. Combining

$$\begin{aligned} \text{I} : f > 0 & \text{ (Assumption)} \\ \text{II} : f' > 0 & \text{ (Assumption)} \\ \text{III} : f'' > 0 & \Rightarrow f(x) < \bar{f}(l_\varepsilon(x)) \quad (\text{cf. Equation B.10}) \\ \text{IV} : f''' \leq 0 & \Rightarrow (\bar{f}(l_\varepsilon(x)))' \leq f'(x) \quad (\text{cf. Equation B.14}) \end{aligned}$$

implies

$$\begin{aligned} \text{I} \wedge \text{II} \wedge \text{III} \wedge \text{IV} & \Rightarrow (\bar{f}(l_\varepsilon(x)))' f(x) < f'(x) \bar{f}(l_\varepsilon(x)) \\ & \Rightarrow \left(\frac{\bar{f}(l_\varepsilon(x))}{f(x)} \right)' < 0 \\ & \Rightarrow g'(x) < 0, \end{aligned}$$

which proofs Equation 5.16(a). Similarly, combining

$$\begin{aligned} \text{I} : f > 0 & \text{ (Assumption)} \\ \text{II} : f' > 0 & \text{ (Assumption)} \\ \text{III} : f'' < 0 & \Rightarrow f(x) > \bar{f}(l_\varepsilon(x)) \quad (\text{cf. Equation B.10}) \\ \text{IV} : f''' \geq 0 & \Rightarrow (\bar{f}(l_\varepsilon(x)))' \geq f'(x) \quad (\text{cf. Equation B.14}) \end{aligned}$$

implies

$$\begin{aligned} \text{I} \wedge \text{II} \wedge \text{III} \wedge \text{IV} & \Rightarrow (\bar{f}(l_\varepsilon(x)))' f(x) > f'(x) \bar{f}(l_\varepsilon(x)) \\ & \Rightarrow \left(\frac{\bar{f}(l_\varepsilon(x))}{f(x)} \right)' > 0 \\ & \Rightarrow g'(x) > 0, \end{aligned}$$

which proofs Equation 5.16(b). □

Calculation of the Derivative of Equation 6.21

In this appendix the step-by-step calculation of the derivative of Equation 6.21 is shown.

$$\begin{aligned}
 & \frac{\partial^2}{\partial x \partial (LT)} \log f_{\phi_l}(x, L, T) \\
 &= \frac{\partial}{\partial x} \left[\frac{\partial}{\partial (LT)} \log f_{\phi_l}(x, L, T) \right] \\
 &= \frac{\partial}{\partial x} \left[\frac{\partial}{\partial (LT)} \log \left(1 - \frac{(x-1)^2 (e^{-LT} - 1)^2}{(LT)^2} \right) \right] \\
 &= \frac{\partial}{\partial x} - \frac{(x-1)^2}{1 - \frac{(x-1)^2 (e^{-LT} - 1)^2}{(LT)^2}} \frac{\partial}{\partial (LT)} \left(\frac{(e^{-LT} - 1)^2}{(LT)^2} \right) \\
 &= \frac{\partial}{\partial x} - \frac{(x-1)^2 \left(-\frac{2e^{-LT}(-1+e^{-LT})}{(LT)^2} - \frac{2(-1+e^{-LT})^2}{(LT)^3} \right)}{1 - \frac{(x-1)^2 (e^{-LT} - 1)^2}{(LT)^2}},
 \end{aligned}$$

which can be simplified

$$\begin{aligned}
 & [\cdot] \\
 &= \frac{\partial}{\partial x} \frac{(x-1)^2 e^{-2LT} (LT)^{-3} (2e^{2LT}(-1+e^{-LT})^2 + 2LT e^{LT}(-1+e^{-LT}))}{1 - \frac{(x-1)^2(-1+e^{-LT})^2}{(LT)^2}} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 e^{-2LT} (LT)^{-3} (e^{2LT} - 2e^{LT} + 1 - LT e^{LT} + LT)}{1 - \frac{(x-1)^2(-1+e^{-LT})^2}{(LT)^2}} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 e^{-2LT} (LT)^{-3} (-1+e^{LT})(-LT+e^{LT}-1)}{1 - \frac{(x-1)^2(-1+e^{-LT})^2}{(LT)^2}} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 (-1+e^{LT})(-LT+e^{LT}-1)}{e^{2LT} (LT)^3 \left(1 - \frac{(x-1)^2(1-2e^{-LT}+2e^{-2LT})}{(LT)^2}\right)} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 (-1+e^{LT})(-LT+e^{LT}-1)}{e^{2LT} (LT)^3 - (LT e^{2LT} - 2LT e^{LT} + LT)(x-1)^2} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 (-1+e^{LT})(-LT+e^{LT}-1)}{a(e^{2LT}(x-1)^2 - (x-1)^2 - e^{2LT}(x-1)^2 + e^{LT}(LT)^2)} \\
 &= \frac{\partial}{\partial x} \frac{2(x-1)^2 (-1+e^{LT})(-LT+e^{LT}-1)}{a(e^{2LT}(x-1)^2 - (x-1)^2 + e^{2LT}((LT)^2 - (x-1)^2))} ,
 \end{aligned}$$

and the derivative with respect to x becomes

$$\begin{aligned}
 & [\cdot] \\
 &= \frac{2(-1+e^{LT})(-LT+e^{LT}-1)}{LT} \frac{\partial}{\partial x} \frac{(x-1)^2}{2e^{LT}(x-1)^2 - (x-1)^2 + e^{2LT}((LT)^2 - (x-1)^2)} , \\
 &= [\cdot] \\
 &= \frac{4LT e^{2LT}(-1+e^{LT})(-LT+e^{LT}-1)(x-1)}{(2e^{LT}(x-1)^2 - (x-1)^2 + e^{2LT}((LT)^2 - (x-1)^2))^2} ,
 \end{aligned}$$

which equals the right-hand side of Equation 6.21. □

Basins of Attraction in Environments 2 and 4 of Chapter 7

The environments as defined in Equations 7.9 and 7.11 have two optima, one at $z = 0$ and the other at $z = 1$. The corresponding basins of attraction have equal size within the interval $[0, 1]$ if there is a minimum at $z = 0.5$. The following derivations assume $z_{\text{opt}} = 0$. However, the transfer to the case $z_{\text{opt}} = 1$ is trivial. With this assumption, the adaptive value function that corresponds to Environments 2 and 4 can be written as

$$f(z) = h e^{-\left(\frac{z}{\sigma_{\text{opt}}}\right)^2} + e^{-16(z-1)^2} . \quad (\text{D.1})$$

The first derivative of f w.r.t. z is

$$f'(z) = -\frac{2hz}{\sigma_{\text{opt}}^2} e^{-\left(\frac{z}{\sigma_{\text{opt}}}\right)^2} - 32(z-1)e^{-16(z-1)^2} , \quad (\text{D.2})$$

and the second derivative of f w.r.t. z is

$$f''(z) = \left(\frac{4hz^2}{\sigma_{\text{opt}}^4} - \frac{2h}{\sigma_{\text{opt}}^2}\right) e^{-\left(\frac{z}{\sigma_{\text{opt}}}\right)^2} + (1024(z-1)^2 - 32) e^{-\frac{1}{4\sigma_{\text{opt}}^2}} . \quad (\text{D.3})$$

where f' denotes $\frac{\partial f}{\partial z}$ and f'' denotes $\frac{\partial^2 f}{\partial z^2}$. A necessary condition for a minimum at $z = 0.5$ is a zero first derivative at this point. Evaluating f' at $z = 0.5$ and setting it to zero yields

$$\begin{aligned} f'(0.5) &= 0 \\ \Leftrightarrow \frac{h}{\sigma_{\text{opt}}^2} e^{-\frac{1}{4\sigma_{\text{opt}}^2}} - 16e^{-4} &= 0 \\ \Leftrightarrow h &= 16e^{-4} \sigma_{\text{opt}}^2 e^{\frac{1}{4\sigma_{\text{opt}}^2}} . \end{aligned} \quad (\text{D.4})$$

A sufficient condition for a minimum at $z = 0.5$ is a positive second derivative of f at this point, i.e.

$$\begin{aligned} f''(0.5) &> 0 \\ \Leftrightarrow h \left(\frac{1}{\sigma_{\text{opt}}^4} - \frac{2}{\sigma_{\text{opt}}^2}\right) e^{-\frac{1}{4\sigma_{\text{opt}}^2}} + 224e^{-4} &> 0 . \end{aligned} \quad (\text{D.5})$$




Figure D.1: Relationship between h and σ_{opt} in Environments 2 and 4 that needs to be satisfied in order to have equal size of basins of attraction of the two optima within the interval $[0, 1]$.

Including the necessary condition of Equation D.4 yields

$$\begin{aligned}
 & 16e^{-4}\sigma_{\text{opt}}^2 \left(\frac{1}{\sigma_{\text{opt}}^4} - \frac{2}{\sigma_{\text{opt}}^2} \right) + 224e^{-4} > 0 \\
 \Leftrightarrow & \left(\frac{1}{\sigma_{\text{opt}}^2} - 1 \right) + 14 > 0 \tag{D.6} \\
 \Leftrightarrow & \frac{1}{\sigma_{\text{opt}}^2} > -13 ,
 \end{aligned}$$

which is true. Thus for all combinations h and σ_{opt} that satisfy Equation D.4, f has a minimum at $z = 0.5$. The relationship of Equation D.4 is shown in the left graph of Figure D.1. h is a function of σ_{opt} but the inverse function does not exist because there exist two corresponding σ_{opt} for a given h . Since h is a function of σ_{opt} (according to Equation D.4) the fitness landscape f with a minimum at $z = 0.5$ can be drawn w.r.t. z to σ_{opt} as shown in Figure D.2. The global fitness maximum is at $z = 0$ for small σ_{opt} and at $z = 1$ for large σ_{opt} . The transition occurs where the two maxima have equal function value, i.e. $f(0) - f(1) = 0$. The corresponding σ_{opt} can be calculated by solving $f(0) - f(1) = 0$ for σ_{opt} having regard to Equation D.4. Solving this equation shows that the transition occurs at $\sigma_{\text{opt}} = 0.25$ with the corresponding $h = 1$ (according to Equation D.4).

Thus, function f is limited to parameters $\sigma_{\text{opt}} > 0.25$ and $h > 1$. Under this constraints σ_{opt} is a function of h and can be solved numerically. The resulting graph is shown in the right panel of Figure D.1. For each $h > 1$ a unique σ_{opt} can now be determined such that the function f has a minimum at $z = 0$, and global maximum at $z_{\text{opt}} = 0$. Note that this calculations assume that $z_{\text{opt}} = 0$. The calculations for the second case $z_{\text{opt}} = 1$ are analogous and yield the same σ_{opt} for a given h .

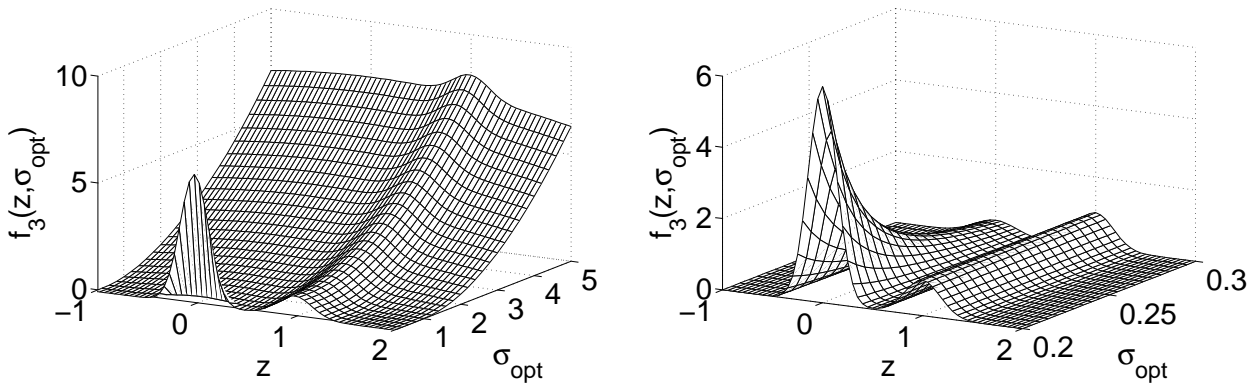


Figure D.2: *Fitness landscape $f(z)$ w.r.t. σ_{opt} such that there exists a minimum at $z = 0.5$. The left panel shows a large σ_{opt} -range, the right panel zooms into a smaller σ_{opt} -range. The global maximum switches from $z = 0$ to $z = 1$ at $\sigma_{opt} = 0.25$. Function f is therefore limited to $\sigma_{opt} < 0.25$*

Simulation Results for Deterministically Changing Environment 4 of
Chapter 7

In Section 7.5.2 a simulation study was presented based on a stochastically changing environment as defined in Equation 7.11. There, large variations of the actual times between environmental changes could be observed. In a preliminary study, it has been assumed that the environment changes deterministically every T time units, i.e. the deterministic Environment 4, \hat{f}_4 , is given by

$$\hat{f}_4(z, t) = h e^{-\left(\frac{z-z_{\text{opt}}(t)}{\sigma_{\text{opt}}}\right)^2} + e^{-\left(\frac{z-(1-z_{\text{opt}}(t))}{0.25}\right)^2}, \text{ with } h > 1,$$

$$\text{and } z_{\text{opt}}(t) = \begin{cases} 0, & \text{if } \lfloor t/T \rfloor \text{ modulo } 2 \equiv 0 \\ 1, & \text{else} \end{cases} \quad (\text{E.1})$$

A simulation study, similar to the one in Section 7.5.2, was done for the deterministic version of Environment 4, with interesting results. Figures E.1 and E.2 show the results. Qualitatively, we obtain the same results as in the corresponding experiment with stochastic changes (cf. Figures 7.14 and 7.15), namely that the optimal adaptation behavior can be observed for an intermediate lifetime. However, different from the results in stochastic Environment 4, we find a peculiar wave curvature in some curves in Figure E.1. This is particularly clear in case of height factor 5 (right panel) and change intervals 50 and 200. Extending the x-axis to larger lifetime values would reveal a similar form for the change interval 400. It turns out that this finding is an emergent statistical property of the fact that environmental changes appear in constant temporal intervals.

The following example explains this curiosity: If the lifetime of an individual is equal to or shorter than the length of one change interval than the mean adaptive value over an individual's lifetime varies strongly depending on its birthdate. In one extreme case, an individual with genotype 0 may be close to the global optimum throughout its whole life while in the other extreme case the global optimum is at 1 throughout its whole life, thus achieving only a low mean adaptive value. Those individuals that are “biased” toward the




Figure E.1: Mean adaptive value for different constant lifetimes in the deterministically changing Environment 4 for change intervals $T \in \{50, 200, 400\}$ and height factors 2 (left panel) and 5 (right panel), respectively. There exists an optimal lifetime that depends environmental dynamics and height differences between local and global optimum. Notice, the peculiar sine-like curvature of some curves.




Figure E.2: Typical evolutionary runs in deterministically changing Environment 4. The thick gray line shows the global optimum, the thick black dots shows the genotype values, and the smaller gray dots show the phenotype values present in the population at a time. With $L = 1$ (pure population adaptation) the population only occasionally discovers a new global optimum. For long lifetimes $L = 200$ the population is not flexible enough to move the majority of individuals to the current global optimum before the next environmental change. Only in the intermediate case of $L = 30$, a good balance between exploration and exploitation is achieved, and as a consequence, the population follows the environmental dynamics.




Figure E.3: *Explanation for the wave curve in deterministic Environment 4*

global optimum have a higher number of expected offsprings. These offsprings in turn are more likely to be biased toward the local optimum. Thus, one would expect that in this case on average the majority of individuals is located on the (non-global) local optimum. If the lifetime of an individual is twice the length of the environmental change interval the birthdate has no influence on its mean adaptive value since it equally long on the global and on the local optimum hill. Thus, no bias is expected here. However, if the lifetime of an individual is three times the length of the environmental change interval the birthdate becomes important again. In general, if $L = (2n - 1)T$, with $n \in \{1, 2, 3, 4, \dots\}$, the population is biased toward the local optimum, whereas if $L = 2nT$, with $n \in \{1, 2, 3, 4, \dots\}$, the population is not bias toward the local or global optimum.

This explanation is supported by the following experiment. For three different change intervals $T \in \{50, 200, 400\}$ evolution is run for a range of different lifetimes which are multiples of the corresponding T . For each setting the average fraction of individuals that are located on the global optimum hill are measured (depending on the environmental state either $z \leq 0.5$ or $z \geq 0.5$). The results as shown in Figure E.3 indeed support the above given explanation. If $L = (2n - 1)T$ the mean fraction of genotypes in the basin of attraction of the global optimum is smaller than 0.5 and approximately 0.5 if $L = 2nT$.

The peculiar wave-curve of Figure E.1 disappear in environments with stochastic changes.

Bibliography

- [1] D. Ackley and M. Littman. Interactions between learning and evolution. In C. Langton, J. Famer, and S. Rasmussen, editors, *Proceedings of the Second Conference on Artificial Life*, pages 487–509, Redwood City, California, 1991. Addison-Wesley.
- [2] E. Alpaydin. *Introduction to Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, Cambridge, Massachusetts, 2004.
- [3] L. Ancel. Undermining the Baldwin expediting effect: How phenotypic plasticity influences the rate of evolution. *Theoretical Population Biology*, 58(4):307–319, 2000.
- [4] LW. Ancel and J. Bull. Fighting change with change: Adaptive variation in an uncertain world. *Trends in Ecology and Evolution*, 17(12):551–557, 2002.
- [5] R. Anderson. Learning and evolution: A quantitative genetics approach. *Journal of Theoretical Biology*, 175(1):89–101, 1995.
- [6] J. Baker. Reducing bias and inefficiency in the selection algorithm. In J. Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms and their Applications*, pages 14–21, Massachusetts, 1987. Lawrence Erlbaum Associates.
- [7] J. Baldwin. A new factor in evolution. *American Naturalist*, 30:441–451, 1896.
- [8] N. Behera and V. Nanjundiah. Phenotypic plasticity can potentiate rapid evolutionary change. *Journal of Theoretical Biology*, 226:177–184, 2004.
- [9] R. Belew. Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4(1):11–49, 1990.
- [10] Richard K. Belew and Melanie Mitchell, editors. *Adaptive individuals in evolving populations: Models and algorithms*. Addison-Wesley Longman Publishing Co., Inc., Boston, Massachusetts, 1996.
- [11] T. Blikle and L. Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4(4):361–394, 1996.

- [12] E. Borenstein. *Evolutionary Dynamics of Adaptive Populations: The Effect of Phenotypic Plasticity, Imitation and Culture*. PhD thesis, Tel Aviv University, 2006.
- [13] E. Borenstein, I. Meilijson, and E. Ruppin. The effect of phenotypic plasticity on evolution in multi-peaked fitness landscapes. *Journal of Evolutionary Biology*, 19(5):1555–70, 2006.
- [14] J. Branke. *Evolutionary Optimization in Dynamic Environments*. Kluwer, 2001.
- [15] J.J. Bull, L. Ancel-Meyers, and M. Lachmann. Quasispecies made simple. *PLoS Computational Biology*, 1(6), 2005.
- [16] L. Bull. On the Baldwin effect. *Artificial Life*, 5(3):241–246, 1999.
- [17] J.A. Bullinaria. The effect of learning on life history evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007)*, pages 222–229, New York, NY, 2007. ACM Press.
- [18] L. Cavalli-Sforza and M. Feldman. Evolution of continuous variation: Direct approach through joint distribution of genotypes and phenotypes. *Proceedings of the National Academy of Sciences of the United States of America*, 73:1689–1692, 1976.
- [19] F. Cecconi, F. Menczer, and R. Belew. Maturation and the evolution of imitative learning in artificial organisms. *Adaptive Behavior*, 4:29–50, 1996.
- [20] M. Chang, K. Ohkura, K. Ueda, and M. Sugiyama. Group selection and its application to constrained evolutionary optimization. In *The 2003 Congress on Evolutionary Computation (CEC'03)*, volume 1, pages 684–691, Piscataway, New Jersey, 2003. IEEE Press.
- [21] M. Colombetti and M. Dorigo. Evolutionary computation in behavior engineering. In X. Yao, editor, *Evolutionary Computation: Theory and Applications*, chapter 2, pages 37–80. World Scientific Publishing, Singapore, 1999.
- [22] S.H. Cousins. Species diversity measurement: Choosing the right index. *Trends in Ecology and Evolution*, 6(6):190–192, 1991.
- [23] F.H.C. Crick. The biological replication of macromolecules. *Symposia of the Society for Experimental Biology*, 12:138–163, 1958.
- [24] F.H.C. Crick. Central dogma of molecular biology. *Nature*, 227:561–563, 1970.
- [25] J.F. Crow and M. Kimura. The theory of genetic loads. In S.J. Geerts, editor, *Proceedings of the XI'th International Congress of Genetics 3*, volume 2, pages 495–505, Oxford, 1964. Pergamon.
- [26] D. Curran and C. O'Riordan. Measuring diversity in populations employing cultural learning in dynamic environments. In M.S. Capcarrere, A.A. Freitas, P.J. Bentley, C.G. Johnson, and J. Timmis, editors, *Advances in Artificial Life: 8th European Conference, ECAL 2005*, LNCS, Berlin, 2005. Springer.

-
- [27] D. Curran and C. O’Riordan. Increasing population diversity through cultural learning. *Adaptive Behavior*, 14(4):315–338, 2006.
- [28] C. Darwin. *The Origin of Species*. John Murray, 1859.
- [29] D. Depew. Baldwin and his many effects. In BH. Weber and D. Depew, editors, *Evolution and Learning - The Baldwin effect reconsidered*, pages 3–31. MIT Press, Cambridge, Massachusetts, 2003.
- [30] H. Dopazo, MB. Gordon, R. Perazzo, and S. Risau-Gusman. A model for the interaction of learning and evolution. *Bulletin of Mathematical Biology*, 63:117–134, 2001.
- [31] H. Dopazo, MB. Gordon, R. Perazzo, and S. Rissau. A model for the emergence of adaptive subsystems. *Bulletin of Mathematical Biology*, 65:27–56, 2003.
- [32] K. Doya. What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks*, 12(7-8):961–974, 1999.
- [33] K. Doya. Recurrent networks: supervised learning. In M. Arbib, editor, *The handbook of brain theory and neural networks*, pages 955–960. The MIT Press, Cambridge, Massachusetts, 2nd edition, 2002.
- [34] AE. Eiben, EHL. Aarts, and KM. van Hee. Global convergence of genetic algorithms: a markov chain analysis. In HP. Schwefel and R. Manner, editors, *Parallel Problem Solving from Nature*, pages 4–12, Berlin, 1991. Springer.
- [35] AE. Eiben and CA. Schippers. On evolutionary exploration and exploitation. *Fundamenta Informaticae*, 35(1-4):35–50, 1998.
- [36] AE. Eiben and JE. Smith. *Introduction to Evolutionary Computation*. Springer, Berlin, 1st edition, 2003.
- [37] M. Eigen. Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58:465–523, 1971.
- [38] M. Eigen and P. Schuster. *The Hypercycle: A Principle of Natural Self-Organization*. Springer, Berlin, 1979.
- [39] D. Floreano, P. Husbands, and S. Nolfi. Evolutionary Robotics. In *Handbook of Robotics*. Springer, Berlin, 2008.
- [40] D. Floreano and F. Mondada. Evolution of plastic neurocontrollers for situated agents. In P. Maes, M. Mataric, JA. Meyer, J. Pollack, H. Roitblat, and S. Wilson, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 402–410, Cambridge, Massachusetts, 1996. The MIT Press.
- [41] D. Floreano and J. Urzelai. Evolutionary robots with self-organization and behavioral fitness. *Neural Networks*, 13:431–443, 2000.

- [42] D. Floreano and J. Urzelai. Neural morphogenesis, synaptic plasticity, and evolution. *Neural Networks*, 120(3-4):225–240, 2001.
- [43] DB. Fogel, editor. *Evolutionary Computation - The Fossil Record*. John Wiley & Sons, Inc., New York, NY, 1998.
- [44] J. Fontanari and F. Meir. The effect of learning on the evolution of asexual populations. *Complex Systems*, 4:401–414, 1990.
- [45] P. Foster. Adaptive mutation: Has the unicorn landed? *Genetics*, 148:1453–1459, 1998.
- [46] R. French and A. Messinger. Genes, phenes and the Baldwin effect. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 277–282, Cambridge, Massachusetts, 1994. The MIT Press.
- [47] DJ. Futuyma. *Evolution*. Sinauer Associates, Sunderland, MA, 2005.
- [48] DE. Goldberg and P. Segrest. Finite markov chain analysis of genetic algorithms. In JJ. Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms*, pages 1–8, Cambridge, MA, 1987. Lawrence Erlbaum Associates.
- [49] D. Gordon. Phenotypic plasticity. In E. Lloyd and E. Kell, editors, *Keywords in Evolutionary Biology*, pages 255–262. Harvard University Press, Cambridge, Massachusetts, 1992.
- [50] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, New York, 2001.
- [51] F. Gruau and D. Whitley. Adding learning to the cellular development of neural networks: Evolution and Baldwin effect. *Evolutionary Computation*, 1(3):213–233, 1993.
- [52] B. Güler. Ein populationsbasiertes Markov-Ketten-Modell zur Analyse des Einflusses von Lernen auf Evolution. Master’s thesis, University of Karlsruhe, 2007. English title: A population-based Markov-chain-model for the analysis of the influence of learning on evolution.
- [53] JBS. Haldane. A mathematical theory of natural and artificial selection. part 1. *Transactions of the Cambridge Philosophical Society*, 23:19–41, 1924.
- [54] JBS. Haldane. The effect of variation on fitness. *American Naturalist*, 71:337–349, 1937.
- [55] SA. Harp, T. Samad, and A. Guha. Toward the genetic synthesis of neural networks. In JD. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms and Their Applications*, pages 360–369, San Mateo, California, 1989. Morgan Kaufmann.
- [56] SA. Harp, T. Samad, and A. Guha. Designing application-specific neural networks using genetic algorithms. In DS. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 447–454, San Francisco, 1990. Morgan Kaufmann Publishers.

-
- [57] WE. Hart. *Adaptive Global Optimization with Local Search*. PhD thesis, University of California, San Diego, 1994.
- [58] WE. Hart, E. William, N. Krasnogor, and JE. Smith. Memetic evolutionary algorithms. In WE. Hart, E. William, N. Krasnogor, and JE. Smith, editors, *Recent Advances in Memetic Algorithms*, pages 3–27. Springer, Berlin, 2005.
- [59] I. Harvey. The puzzle of the persistent question marks: A case study of genetic drift. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 15–22, San Francisco, 1993. Morgan Kaufmann.
- [60] I. Harvey. Is there another new factor in evolution? *Evolutionary Computation*, 4(3):311–327, 1997. Special Issue on Evolution, Learning and Instinct.
- [61] I. Harvey, E. Di Paolo, R. Wood, M. Quinn, and E. Tuci. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79–98, 2005.
- [62] S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall PTR, Upper Saddle River, New Jersey, 1994.
- [63] J. He and X. Yao. From an individual to a population: An analysis of the first hitting time of population-based evolutionary algorithm. *IEEE Transactions on Evolutionary Computation*, 6:495–511, 2003.
- [64] GE. Hinton and SJ. Nowlan. How learning can guide evolution. *Complex Systems*, 1:495–502, 1987.
- [65] GE. Hinton and TJ. Sejnowski, editors. *Unsupervised Learning - Foundations of Neural Computation*, chapter Backcover. MIT Press, Cambridge, Massachusetts, 1999.
- [66] J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.
- [67] J. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, Massachusetts, 2nd edition, 1992.
- [68] R. Holliday and JE. Pugh. DNA modification mechanisms and gene activity during development. *Science*, 187:226–232, 1975.
- [69] DE. Holmes and LC. Jain, editors. *Innovations in Machine Learning: Theory and Applications*. Springer, Berlin, 2006.
- [70] CR. Houck, JA. Joines, MG. Kay, and JR. Wilson. Empirical investigation of the benefits of partial lamarckianism. *Evolutionary Computation*, 5(1):31–60, 1997.
- [71] SH. Hurlbert. The nonconcept of species diversity: A critique and alternative parameters. *Ecology*, 52(4):577–586, 1971.

- [72] M. Hüskén and C. Igel. Balancing learning and evolution. In WB. Langdon, E. Cant-Paz, KE. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, VG. Honavar, G. Rudolph, J. Wegener, L. Bull MA. Potter, AC. Schultz, JF. Miller, E. Burke, and N. Jonoska, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002)*, pages 391–398, San Francisco, 2002. Morgan Kaufmann.
- [73] E. Jablonka and M. Lamb. *Evolution in Four Dimensions - Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. MIT Press, Cambridge, Massachusetts, 2005.
- [74] E. Jablonka, B. Oborny, I. Molnar, E. Kisdi, J. Hofbauer, and T. Czaran. The adaptive advantage of phenotypic memory in changing environments. *Philosophical Transactions of the Royal Society of London. Series B. Biological Sciences*, 29(350):133–141, 1995.
- [75] J. Jaenike and DR. Papaj. Learning and patterns of host use by insects. In M. Isman and BD. Roitberg, editors, *Insect chemical ecology: An evolutionary approach*, pages 245–264. Chapman and Hall, New York, 1992.
- [76] T. Johnston. Selective costs and benefits in the evolution of learning. *Advances in the Study of Behavior*, 12:65–106, 1982.
- [77] TB. Jongeling. Self-organization and competition in evolution: a conceptual problem in the use of fitness landscapes. *Journal of Theoretical Biology*, 178:369–373, 1996.
- [78] BA. Julstrom. Comparing Darwinian, Baldwinian, and Lamarckian search in a genetic algorithm for the 4-cycle problem. *Late Breaking Papers at the 1999 Genetic and Evolutionary Computation Conference*, pages 134–138, 1999.
- [79] S. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
- [80] R. Keesing and D. Stork. Evolution and learning in neural networks: The number and distribution of learning trials affect the rate of evolution. In R. Lippmann, J. Moody, and D. Touretzky, editors, *Proceedings of Neural Information Processing Systems*, pages 804–810, 1991.
- [81] R. Kicinger, T. Arciszewski, and KA. De Jong. Evolutionary computation and structural design: A survey of the state of the art. *Computers and Structures*, 83(23-24):1943–1978, 2005.
- [82] M. Kimura. On the change of population fitness by natural selection. *Heredity*, 12:145–167, 1958.
- [83] H. Kitano. Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4(4):461–476, 1990.
- [84] DE. Koshland Jr. Nature, nurture, and behavior. *Science*, 235(4795):1445–, 1987.

- [85] SG. Krantz. *Handbook of Complex Variables*, chapter 2.1.5 - The Fundamental Theorem of Calculus along Curves., page 22. Birkhäuser, Boston, Massachusetts, 1999.
- [86] N. Krasnogor and J. Smith. A tutorial for competent memetic algorithms: model, taxonomy, and design issues. *IEEE Transactions on Evolutionary Computation*, 9(5):474–488, 2005.
- [87] CB. Krimbas. On fitness. *Biology and Philosophy*, 19(2):185–203, 2004.
- [88] L. Krubitzer and DM. Kahn. Nature versus nurture revisited: an old idea with a new twist. *Progress in Neurobiology*, 70(1):33–52, 2003.
- [89] KWC. Ku and MW. Mak. Exploring the effects of Lamarckian and Baldwinian learning in evolving recurrent neural networks. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, pages 617–621, Piscataway, New Jersey, 1997. IEEE press.
- [90] KWC. Ku, MW. Mak, and WC. Siu. Adding learning to cellular genetic algorithms for training recurrent neural networks. *IEEE Transactions on Neural Networks*, 10(2):239–252, 1999.
- [91] KWC. Ku, MW. Mak, and WC. Siu. Approaches to combining local and evolutionary search for neural networks: A review and some new results. In A. Ghosh and S. Tsutsui, editors, *Advances in Evolutionary Computing*, pages 615–642. Springer, Berlin, 2003.
- [92] S. Kumar and PJ. Benley, editors. *On Growth, Form and Computers*. Elsevier, Amsterdam, 2003.
- [93] JB. Lamarck. *Philosophie zoologique ou exposition des considérations relatives l’histoire naturelle des animaux*. UCP (reprinted 1984), 1809.
- [94] R. Lande. Natural selection and random genetic drift in phenotypic evolution. *Evolution*, 30(2):314–334, 1976.
- [95] T. Lenaerts, A. Defaweux, P. van Remortel, J. Reumers, and B. Manderick. Multi-level selection and immune networks: Preliminary discussion of an abstract model. In RK. Standishand MA. Bedau and HA. Abbass, editors, *Proceedings of the eighth international conference on Artificial life (ICAL 2003)*, pages 223–226, Cambridge, Massachusetts, 2003. MIT Press.
- [96] R. Levins. *Evolution in Changing Environments*. Princeton University Press, 1968.
- [97] Z. Lippman and R. Martienssen. The role of RNA interference in heterochromatic silencing. *Nature*, 431:364–370, 1986.
- [98] AE. Magurran. Biological diversity. *Current Biology*, 15(4):116–118, 2005.
- [99] DP. Mandic and J. Chambers. *Recurrent Neural Networks for Prediction: Learning Algorithms, Architectures and Stability*. John Wiley & Sons, Inc., New York, NY, 2001.

- [100] C. Mattiussi, M. Waibel, and D. Floreano. Measures of diversity for populations and distances between individuals with highly reorganizable genomes. *Evolutionary Computation*, 12(4):495–515, 2004.
- [101] G. Mayley. The evolutionary cost of learning. In *From Animals to Animats: From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 458–467, 1996.
- [102] G. Mayley. Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4(3):213–234, 1996.
- [103] G. Mayley. Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life 97*, pages 135–144, Cambridge, Massachusetts, 1997. The MIT Press.
- [104] J. Maynard-Smith. Group selection and kin selection. *Nature*, 201:1145–1147, 1964.
- [105] J. Maynard-Smith. When learning guides evolution. *Nature*, 329(6142):761–762, 1987.
- [106] F. Mery and T. Kawecki. Experimental evolution of learning ability in fruit flies. *Proceedings of the National Academy of Sciences*, 99(22):14274–14279, 2002.
- [107] F. Mery and T. Kawecki. The effect of learning on experimental evolution of resource preference in drosophila melanogaster. *Evolution*, 58(4):757–767, 2004.
- [108] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution*. Springer, Berlin, 1996.
- [109] R. Mills and RA. Watson. On crossing fitness valleys with the Baldwin effect. In LM. Rocha, LS. Yaeger, MA. Bedau, D. Floreano, RL. Goldstone, and A. Vespignani, editors, *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 493–499, Cambridge, Massachusetts, 2006. MIT Press.
- [110] M. Mitchell and S. Forrest. Genetic algorithms and artificial life. *Artificial Life*, 1(3):267–289, 1994.
- [111] CCJ. Moey and JE. Rowe. Population aggregation based on fitness. *Natural Computing: An international journal*, 3(1):5–19, 2004.
- [112] CCJ. Moey and JE. Rowe. A reduced markov model of gas without the exact transition matrix. In X. Yao, EK. Burke, JA. Lozano, J. Smith, JJ. Merelo Guervos, JA. Bullinaria, JE. Rowe, P. Tino, A. Kaban, and H. Schwefel, editors, *Parallel Problem Solving from Nature VIII*, number 3242 in LNCS, Berlin, 2004. Springer.
- [113] BR. Moore. The evolution of learning. *Biological Reviews*, 79:301–335, 2004.

-
- [114] RW. Morrison and KA. De Jong. Measurement of population diversity. In P. Collet, C. Fonlupt, JK. Hao, E. Lutton, and M. Schoenauer, editors, *Selected Papers from the 5th European Conference on Artificial Evolution*, volume 2310 of *LNCS*, pages 31–41. Springer, Berlin, 2001.
- [115] P. Moscato. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. Technical Report 826, California Inst. of Technology, 1989.
- [116] A. Mukhopadhyay and HA. Tissenbaum. Reproduction and longevity: secrets revealed by *c. elegans*. *Trends in Cell Biology*, 17(2):65–71, 2007.
- [117] AE. Nix and MD. Vose. Modeling genetic algorithms with markov chains. *Annals of Mathematics and Artificial Intelligence*, 5(1):79–88, 1992.
- [118] S. Nolfi. How learning and evolution interact: The case of a learning task which differs from the evolutionary task. *Adaptive Behavior*, 7(2):231236, 1999.
- [119] S. Nolfi and D. Floreano. *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. The MIT Press, Cambridge, Massachusetts, 2001.
- [120] S. Nolfi, D. Parisi, and JL. Elman. Learning and evolution in neural networks. *Adaptive Behavior*, 3(1):5–28, 1994.
- [121] S. Noskowitz and I. Goldhirsch. First passage time distribution in random random walk. *Physical Review A*, 42:2047–2064, 1990.
- [122] A. Ohman and U. Dimberg. Facial expressions as conditioned stimuli for electrodermal responses: a case of "preparedness"?. *Journal of Personality and Social Psychology*, 36(11):1251–1258, 1978.
- [123] M. Olhofer, T. Arima, T. Sonoda, M. Fischer, and B. Sendhoff. Aerodynamic shape optimisation using evolutionary strategies. In IC. Parmee and P. Hajela, editors, *Optimisation in Industry III*, pages 83–94, Berlin, 2001. Springer.
- [124] I. Paenke. Efficient search for robust solutions by means of evolutionary algorithms and fitness approximation. Master's thesis, University of Karlsruhe, 2004.
- [125] I. Paenke, J. Branke, and Y. Jin. Efficient search for robust solutions by means of evolutionary algorithms and fitness approximation. *IEEE Transactions on Evolutionary Computation*, 10(4):405–420, 2006.
- [126] I. Paenke, J. Branke, and Y. Jin. On the influence of phenotype plasticity on genotype diversity. In *IEEE Symposium on Foundations of Computational Intelligence*, pages 33–41, Piscataway, New Jersey, 2007. IEEE Press. Best Student's Paper.
- [127] I. Paenke, Y. Jin, and J. Branke. Balancing population and individual level adaptation in changing environments. *Adaptive Behavior*, 2008. submitted.

- [128] I. Paenke, T.J. Kawecki, and B. Sendhoff. The influence of learning on the rate of evolution. Technical Report 06/04, Honda Research Institute Europe, August 2006.
- [129] I. Paenke, T.J. Kawecki, and B. Sendhoff. On the influence of lifetime learning on selection pressure. In *Artificial Life 10*, pages 500–506, Cambridge, Massachusetts, 2006. MIT Press.
- [130] I. Paenke, T.J. Kawecki, and B. Sendhoff. The influence of learning on evolution - a mathematical framework. *Artificial Life*, 2008. in press.
- [131] I. Paenke, B. Sendhoff, and T.J. Kawecki. Influence of plasticity and learning on evolution under directional selection. *American Naturalist*, 170(2):E47–E58, 2007.
- [132] I. Paenke, B. Sendhoff, J. Rowe, and C. Fernando. On the adaptive disadvantage of Lamarckianism in rapidly changing environments. In F. Almeida e Costa, editor, *Advances in Artificial Life, 9th European Conference on Artificial Life*, pages 355–364, Berlin, 2007. Springer.
- [133] D. Papaj. Optimizing learning and its effect on evolutionary change in behavior. In L. Real, editor, *Behavioral Mechanisms in Evolutionary Ecology.*, pages 133–154. University of Chicago Press, Chicago, Illinois, 1994.
- [134] MR. Papini. Pattern and process in the evolution of learning. *Psychological Review*, 109(1):186–201, 2002.
- [135] D. Parisi, S. Nolfi, and F. Cecconi. Learning, behavior and evolution. In F. Varela and P. Bourguin, editors, *Toward a practice of autonomous systems*, pages 207–216, Cambridge, Massachusetts, 1992. The MIT Press.
- [136] EC. Pielou. Shannon’s formula as a measure of specific diversity: Its use and misuse. *American Naturalist*, 100(914):463–465, 1966.
- [137] D. Potter and D. Held. Absence of food-aversion learning by a polyphagous scarab, *popillia japonica*, following intoxication by geranium, pelargonium x hortorum. *Entomologia Experimentalis et Applicata*, 91(1):83–88, 1999.
- [138] RR. Puentedura. The Baldwin effect in the age of computation. In BH. Weber and DJ. Depew, editors, *Evolution and Learning - The Baldwin Effect Reconsidered*, pages 219–234. MIT Press, Cambridge, Massachusetts, 2003.
- [139] I. Rechenberg. *Evolutionstrategie '94*. Friedrich Frommann Verlag, 1994.
- [140] E. Richards. Inherited epigenetic variation revisiting soft inheritance. *Nature Reviews Genetics*. *Advanced online publication*, 2006.
- [141] GE. Robinson. GENOMICS: Beyond nature and nurture. *Science*, 304(5669):397–399, 2004.

-
- [142] M. Rocha and P. Cortez. The relationship between learning and evolution in static and dynamic environments. In C. Fyfe, editor, *International Symposium on Engineering of Intelligent Systems - proceedings*, pages 377–383. ICSC Academic Press, 2000.
- [143] D. Roff. *Life History Evolution*. Sinauer Associates, Sunderland, MA, 2002.
- [144] A. Rogers and A. Prügel-Bennett. Genetic drift in genetic algorithm selection schemes. *IEEE Transactions on Evolutionary Computation*, 3(4):298–303, 1999.
- [145] RD. Routledge. Diversity indices: Which ones are admissible? *Journal of Theoretical Biology*, 76(4):503–515, 1979.
- [146] G. Rudolph. Finite markov chain results in evolutionary computation: a tour d’horizon. *Fundamenta Informaticae*, 35(1-4):67–89, 1998.
- [147] DE. Rumelhart, GE. Hinton, and RJ. Williams. Learning internal representations by error propagation. In DE. Rumelhart and JL. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1, pages 318–362. The MIT Press, Cambridge, Massachusetts, 1986.
- [148] T. Sasaki and M. Tokoro. Evolving learnable neural networks under changing environments with various rates of inheritance of acquired characters: Comparison of Darwinian and Lamarckian evolution. *Artificial Life*, 5(3):203–223, 1999.
- [149] T. Sasaki and M. Tokoro. Comparison between Lamarckian and Darwinian evolution on a model using neural networks and genetic algorithms. *Knowledge and Information Systems*, 2(2):201–222, 2000.
- [150] HP. Schwefel. *Evolution and Optimum Seeking: The Sixth Generation*. John Wiley & Sons, Inc., New York, 1993.
- [151] R. Selten and R. Stoecker. End behavior in sequences of finite prisoner’s dilemma supergames: A learning theory approach. *Journal of Economic Behavior and Organization*, 7(1):47–70, 1986.
- [152] B. Sendhoff. *Evolution of Structures - Optimization of Artificial Neural Structures for Information Processing*. PhD thesis, Ruhr-Universität Bochum, 1998.
- [153] B. Sendhoff, M. Kreutz, and W. von Seelen. A condition for the genotype-phenotype mapping: Causality. In *Genetic Algorithms: Proceedings of the 7th International Conference (ICGA)*, pages 73–80. Morgan Kaufmann, 1997.
- [154] CE. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423 and 623–656, 1948.
- [155] SHARK EALib (C++ Evolutionary Algorithm library), 2007. <http://shark-project.sourceforge.net>.
- [156] RM. Sibly and P. Calow. *Physiological Ecology of Animals*. Blackwell Scientific Publications, 1984.

- [157] EH. Simpson. Measurement of diversity. *Nature*, 163:688, 1949.
- [158] GG. Simpson. The Baldwin effect. *Evolution*, 7:110–117, 1953.
- [159] R. Skipper. The heuristic role of Sewall Wright’s 1932 adaptive landscape diagram. In *Philosophy of Science (Proceedings)*, volume 71, pages 1176–1188, 2004.
- [160] HB. Slade and SA. Schwartz. Mucosal immunity: The immunology of breastmilk. *Journal of Allergy and Clinical Immunology*, 80:348–356, 1987.
- [161] WM. Spears and KA. De Jong. Analyzing GAs using markov chains with semantically ordered and lumped states. In RK. Belew and MD. Vose, editors, *Proceedings of the 4th Workshop on Foundations of Genetic Algorithms*, pages 85–100. Morgan Kaufmann, 1996.
- [162] H. Spencer. *Principles of biology*, volume 1. Williams and Norgate, 1864.
- [163] F. Spitzer. *Principles of Random Walk*. Springer, Berlin, 2nd edition, 2001.
- [164] PF. Stadler and CR. Stephens. Landscapes and effective fitness. *Comments on Theoretical Biology*, 8:389–431, 2003.
- [165] SC. Stearns. Trade-offs in life-history evolution. *Functional Ecology*, 3:259–268, 1989.
- [166] SC. Stearns. *The evolution of life histories*. Oxford University Press, New York, 1992.
- [167] D. Stephens. Change, regularity and value in the evolution of animal learning. *Behavioral Ecology*, 2(1):77–89, 1991.
- [168] MW. Strickberger. *Evolution*. Jones and Barlett, 1990.
- [169] RS. Sutton and AG. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, Cambridge, Massachusetts, 1998.
- [170] J. Suzuki. A markov chain analysis on a genetic algorithms. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 146–153, San Mateo, CA, 1993. Morgan Kauffman.
- [171] R. Suzuki and T. Arita. How Learning Can Affect the Course of Evolution in Dynamic Environments. In *Proceedings of the Fifth International Symposium on Artificial Life and Robotics*, pages 260–263, 2000.
- [172] R. Suzuki and T. Arita. Repeated occurrences of the Baldwin effect can guide evolution on rugged fitness landscapes. In *IEEE Symposium on Artificial Life*, pages 8–14, Piscataway, New Jersey, 2007. IEEE Press.
- [173] PM. Todd and GG. Miller. Exploring adaptive agency II: Simulating the evolution of associative learning. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 306–315, 1991.

- [174] P. Turney. Myths and Legends of the Baldwin Effect. In T. Fogarty and G. Venturini, editors, *Proceedings of the 13th International Conference on Machine Learning (ICML-96)*, pages 135–142, 1996.
- [175] P. Turney, D. Whitley, and R. Anderson. Evolution, learning and instinct: 100 years of the Baldwin effect. *Evolutionary Computation*, 4(3):iv–viii, 1996. Editorial to the Special Issue: The Baldwin Effect.
- [176] MD. Vose. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press, Cambridge, Massachusetts, 1998.
- [177] CH. Waddington. Genetic assimilation of the bithorax phenotype. *Evolution*, 10(1):1–13, 1956.
- [178] CH. Waddington. Genetic assimilation. *Advances in Genetics*, 10:257–93, 1961.
- [179] GP. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.
- [180] L. Wang, KC. Tan, and CM. Chew. *Evolutionary Robotics: From Algorithms to Implementations*. World Scientific Publishing, Singapore, 2006.
- [181] BH. Weber and DJ. Depew, editors. *Evolution and Learning - The Baldwin effect reconsidered*. MIT Press, Cambridge, Massachusetts, 2003.
- [182] PJ. Werbos. *Beyond regression: New tools for prediction and analysis in the behavioral sciences*. PhD thesis, Harvard University, Cambridge, 1974.
- [183] MJ. West-Eberhard. *Developmental Plasticity and Evolution*. Oxford University Press, New York, 2003.
- [184] D. Whitley. A genetic algorithm tutorial. *Statistics & Computing*, 4(2):65–85, 1994.
- [185] D. Whitley, VS. Gordon, and K. Mathias. Lamarckian evolution, the Baldwin effect and functional optimization. In Y. Davidor, HP. Schwefel, and R. Manner, editors, *Parallel Problem Solving from Nature (PPSN III)*, pages 6–15, Berlin, 1994. Springer.
- [186] DS. Wilson. What is wrong with absolute individual fitness? *TRENDS in Ecology and Evolution*, 19(5), 2004.
- [187] S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D.F. Jones, editor, *Proceedings of the Sixth International Congress of Genetics*, pages 356–366, Menasha, Wisconsin, 1932. Brooklyn botanic garden.
- [188] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.

The book develops a unifying perspective on how learning influences evolutionary dynamics which for the first time explains seemingly contradicting results from the literature. It also reveals formerly unknown facets of the complex interplay of the two major biological adaptation mechanisms. The results apply to algorithmic models from computational intelligence and to quantitative models from evolutionary biology alike, thus contribute to the transfer of principles of biological adaptation to digital processing systems which need to operate under increasingly complex and dynamic conditions.

ISBN: 978-3-86644-247-4

www.uvka.de

