

Evolution Strategies — Simple ‘Models’ of Natural Processes?*

Dipl.–Inform. Frank Kursawe
Department of Computer Science
Systems Analysis Group
44221 Dortmund, Germany

Contents

1	Introduction	2
2	The $(\mu \uparrow \lambda)$ Evolution Strategy	6
3	Transferring the Observations To ...	9
3.1	...Biology	9
3.2	...Organizational Psychology	12
3.3	...Economics	13

*This article has already been published in [12].

Abstract

During the last years, (global) optimization algorithms imitating certain principles of nature have proved their usefulness in various domains of applications. Especially those principles are worth copying where nature has found ‘stable islands’ in a ‘turbulent ocean’ of solution possibilities. Such phenomena can be found in annealing processes, central nervous systems and biological evolution which in turn have lead to the following optimization methods: simulated annealing, (artificial) neural networks and the field of evolutionary computation comprising genetic algorithms, genetic programming, evolutionary programming and *evolution strategies* on which this paper will focus.

Because these algorithms can serve as simple models of the underlying natural processes, why not forget about their problem solving capabilities for a moment and put the emphasis on their self-adapting behaviour? Why not take the opportunity and study e. g. the role of mutation, recombination and selection pressure and their relation to each other with simple objective function(s) serving as the environment in the computer’s artificial world? And why not even go one step further and translate these results to other domains like biology, organizational psychology and economics, thus stressing the similarity of the notions ‘self-organization’ and ‘evolution’ and their usefulness as a common descriptive language across the scientific disciplines?

1 Introduction

Suppose we have to solve the following problem: According to some quality criterion (objective function, model) we are looking for such a setting of the parameters (inputs, regulators in Figure 1) that the output reaches its (global) optimum.




Figure 1: The ‘black box’-model of optimization

More formally: Given a quality function (model)

$$f : M \subseteq \mathbb{R}^n \rightarrow \mathbb{R} \quad ,$$

find a parameter setting $\vec{x}^* \in M$ so that

$$f(\vec{x}) \leq f(\vec{x}^*)$$

holds for all $\vec{x} \in M$. The above is often abbreviated to $f(\vec{x}) \rightarrow \max$. $f(\vec{x}^*)$ is called *global maximum*. Because of $\max\{f(\vec{x})\} = -\min\{-f(\vec{x})\}$ the minimization task is equivalent.

What can we do in order to solve this problem?

1. We can either try to gain more knowledge or exploit what we already know about the interior of the black box in Figure 1. If the objective function turns out to be smooth and differentiable (see Figure 2, left) analytical methods will produce the exact solution.
2. If this turns out to be impossible, we might resort to the brute force method of enumerating the entire search space. But with the number of possibilities growing exponentially in n — the number of dimensions (inputs) — this method becomes infeasible even for low-dimensional spaces.
3. Consequently, mathematicians have developed theories for certain kinds of problems leading to specialized optimization procedures. These algorithms perform well if the black box fulfils their respective prerequisites. For example, Dantzig's *simplex algorithm* [6] probably represents the best known multidimensional method capable of efficiently finding the global optimum of a linear, hence convex, objective function in a search space limited by linear constraints.

Gradient strategies are no longer tied to these linear worlds, but they ‘smooth’ their world by exploiting the objective function’s first partial derivatives one has to supply in advance. Therefore, these algorithms rely on a locally linear internal model of the black box. *Newton strategies* additionally require the second partial derivatives, thus building a quadratic internal model. *Quasi-Newton*, *conjugate gradient* and *variable metric strategies* approximate this information during the search.

The deterministic strategies mentioned so far cannot cope with deteriorations, so the search will stop if anticipated improvements no longer occur. In a multimodal environment (see Figure 2, right), these algorithms move ‘uphill’ from their respective starting points. Hence, they can only converge to the next local optimum.

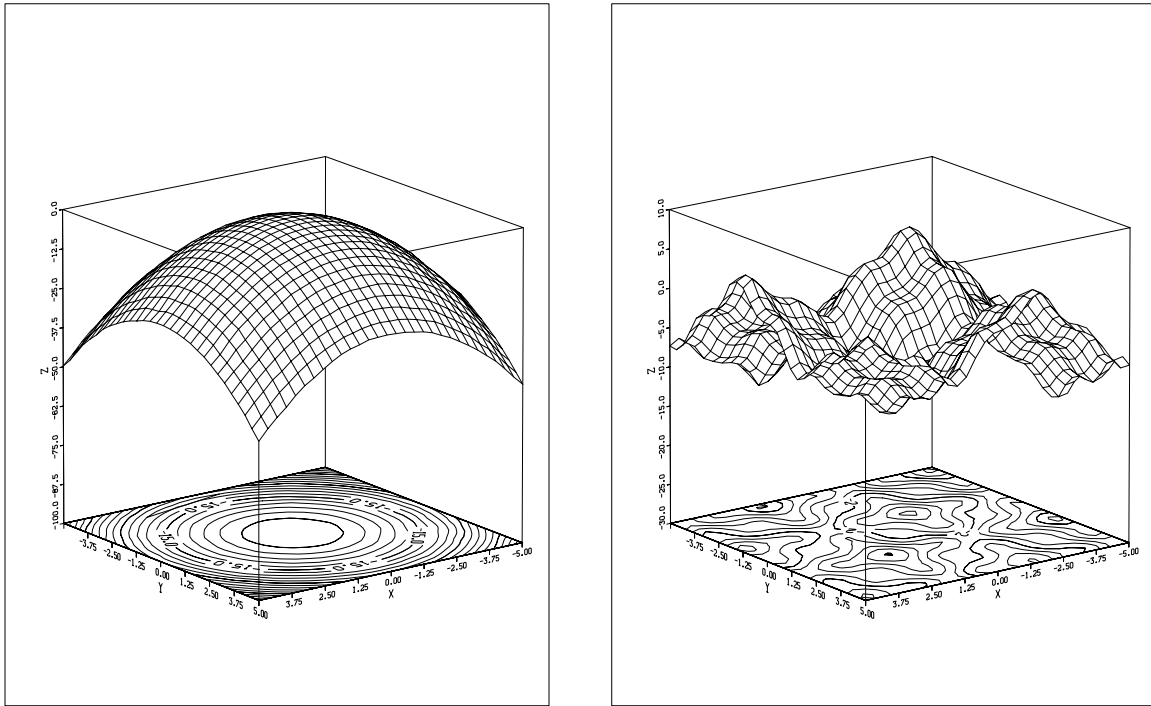


Figure 2: Easy and difficult topologies of objective functions

Newton–Raphson–methods might even diverge if a discrepancy between their internal assumptions and reality occurs. But of course, these methods turn out to be superior if a given task matches their requirements. Not relying on derivatives, *polyeder strategy*, *pattern search* and *rotating coordinate search* should also be mentioned here because they represent robust non–linear optimization algorithms [17].

4. Dealing with technical optimization problems, one will rarely be able to write down the objective function in a closed form $\vec{x} \rightarrow f(\vec{x})$. We often need a *simulation model* in order to capture reality. In general, one cannot even expect these models to behave smoothly. Consequently, derivatives do not exist. That is why optimization algorithms that can successfully deal with black box–type situations have been developed. The increasing applicability is of course paid for by a loss of convergence velocity — compared to algorithms specially designed for the given problem. Furthermore, the guarantee to find the global optimum no longer exists.

But why turn to nature when looking for more powerful algorithms?

In the attempt to create tools for various purposes, mankind has copied — often instinctively — solutions ‘invented’ by nature. Nowadays, one can prove in some cases

that certain forms or structures are not only well adapted to their environment but have even reached the optimum [14]. This is due to the fact that the laws of nature have remained stable during the last 3.5 billion years. For instance, at branching points the measured ratio of the diameters in a system of blood–vessels comes close to the theoretical optimum provided by the laws of fluid dynamics ($2^{-1/3}$). This, of course, only represents a limited, engineering point of view on nature. In general, nature performs *adaptation*, not optimization.

The idea to imitate basic principles of natural processes for optimum seeking procedures emerged more than three decades ago [4, 13, 10, 16]. Although these algorithms have proved to be robust and direct optimization tools, it is only in the last five years that they have caught the researchers’ attention. This is due to the fact that many people still look at organic evolution as a huge game of dice, thus ignoring the fact that this ‘model’ of evolution cannot have worked: a human germ–cell comprises approximately 50.000 genes, each of which consists of about 300 triplets of nucleic bases. Although the four existing bases only encode 20 different amino acids, $20^{15.000.000} \approx 10^{19.500.000}$ different genotypes had to be tested in only $\approx 10^{17}$ seconds — the age of our planet. So, simply ‘rolling the dice’ could not have produced the diversity of today’s complex living systems. Accordingly, taking random samples from the high–dimensional parameter space of an objective function in order to hit the global optimum *must* fail (*Monte–Carlo search*). But by looking at organic evolution as a cumulative, highly parallel sieving process, the results of which pass on slightly modified into the next sieve, the amazing diversity and efficiency on earth no longer appears miraculous. When building a model, the point is to isolate the main mechanisms which have led to today’s world and which have been subjected to evolution themselves. Inevitably, nature has come up with a mechanism allowing individuals of one species to exchange parts of their genetic information (recombination or crossing–over), thus being able to meet changing environmental conditions in a better way.

Today, one may distinguish four types of *evolutionary algorithms*. For an introduction and a comparison of the first three, see [2].

- *evolution strategies*
- evolutionary programming
- genetic algorithms
- genetic programming

This paper will focus on multimembered *evolution strategies* as proposed in [16, 17].

2 The $(\mu \mp \lambda)$ Evolution Strategy

Rechenberg and Schwefel [13, 16] developed the *evolution strategies* when they wanted to optimize technical objects like e. g. a flashing nozzle. No closed form analytical objective function was available, and hence, no applicable optimization method existed at that time. Their first attempt to imitate principles of organic evolution on a computer still resembled those iterative optimization methods known up to that time. In a two-membered or $(1 + 1)$ evolution strategy, one ‘parent’ generates one offspring per generation by applying normally distributed mutations, i. e. smaller steps occur more likely than big jumps, until a ‘child’ performs better than its ancestor and takes its place. Because of this simple structure, theoretical results for stepsize control and convergence velocity could be derived. The ratio between successful and all mutations should come to $1/5$. This first algorithm has then been enhanced to a $(\mu + 1)$ strategy which incorporated recombination for the first time with several parents being available. The mutation scheme and the exogenous stepsize control were taken across unchanged.

Schwefel [16, 17] generalized these strategies to the multimembered evolution strategy now denoted by $(\mu \mp \lambda)$ which imitates the following basic principles of organic evolution:

- population, leading to the possibility of
- recombination with random mating,
- mutation and
- selection

An ‘individual’ consists of the following ‘genes’ representing a point in the search space:




Figure 3: Structure of an individual

- Real-valued *object variables* x_i have to be tuned by recombination and mutation in such a way that an objective function reaches its global optimum. Referring to the metaphor mentioned previously, the x_i of Figure 3 represent the regulators of the black box in Figure 1.
- Real-valued *strategy variables* or mean ‘stepsizes’ σ_i determine the mutability of the x_i . They represent the standard deviation of a $(0, \sigma_i)$ Gaussian distribution being added to each x_i as an undirected mutation. With an expectancy value of 0, the parents will produce offsprings similar to themselves on average. In order to make a doubling and a halving of a stepsize equally probable, the σ_i mutate log-normally distributed from generation to generation. These stepsizes hide the internal ‘model’ the population has made of its environment so far. In other words, a self-adaptation of the stepsizes has taken over from the exogenous control of the $(1 + 1)$ strategy. This concept works because selection sooner or later prefers those individuals having built a good model of the objective function, thus producing better offsprings. Hence, learning takes place on two levels (see Figure 4).
- Depending on an individual’s x_i , the resulting objective function value $f(\vec{x})$ serves as the ‘phenotype’ (fitness) in the selection step. In a plus strategy, the μ best of all $(\mu + \lambda)$ individuals survive to become the parents of the next generation. Using the comma variant, selection takes place only among the λ offsprings. The second scheme is more realistic and therefore more successful, because no individual may survive forever — which could at least theoretically occur using the plus variant. Untypical for conventional optimization algorithms and lavish at first sight, a comma strategy allowing intermediate deterioration performs better. Only by ‘forgetting’ individuals with a good phenotype, which may have been achieved with an internal model being no longer appropriate for further progress, a permanent adaptation of the stepsizes can take place and avoid long stagnation phases due to misadapted stepsizes.

By choosing a certain ratio μ/λ , one can determine the convergence property of the evolution strategy: If one wants a fast, but local convergence, one should choose a small ratio (e. g. $(5, 100)$), but looking for the global optimum, one should favour a ‘softer’ selection (e. g. $(15, 100)$). Figure 4 displays the self-adapting capabilities of different evolution strategies under varying selection pressure. Using

$$f(\vec{x}) := \sum_{i=1}^{30} (i \cdot x_i^2)$$

as the objective function, $n = 30$ stepsizes have to adapt properly with respect to each other in order to achieve maximum progress in adjusting the $n = 30$ x_i on the first level.




Figure 4: Optimum selection pressure

Curve *A* in Figure 4 represents the performance of a strategy provided with perfect ‘knowledge’ about its environment ($\sigma_i = c/\sqrt{i}$), whereas in *B*, randomly chosen relations of the σ_i were fixed. In curve *C*, the stepsizes could self-adapt by means of mutation and recombination. Not surprisingly, $\mu = 1$ proved to be the best choice for *A* and *B* where only one stepsize had to adapt, because all relations had been fixed. On the other hand, one should choose μ between 12 and 17 if one wants learning to take place (*C*). Two observations from Figure 4 are remarkable: The (15, 100) strategy converges nearly as fast as variant *A*, and it performs better than the (15, 100) strategy with perfect knowledge. One can regard this phenomenon as a *synergetic* effect: 15 ‘fools’ perform better *collectively* than the same number of ‘specialists’.

Self-adaptation within evolution strategies depends on the following ‘agents’ [18]:

- One cannot model mutation as a ‘pure’ random process. This would imply a complete independence of an offspring from its parents.
- The population has to consist of a sufficiently large number of individuals. Not only the ‘current best’ should be allowed to reproduce, but a set of good individuals. Biologists have coined the term ‘requisite variety’ being necessary to prevent a species from becoming poorer and poorer genetically and eventually dying out.
- In order to exploit the effects of a population ($\mu > 1$), the individuals should recombine their knowledge with that of others (cooperate) because one cannot expect the knowledge to accumulate in the best individual only.

- In order to allow better internal models (stepsizes) to provide better progress in the future, one should accept deterioration from one generation to the next. A limited life-span in nature is not a sign of failure, but an important means of preventing a species from ‘freezing’ genetically.

Evolution strategies have proved to be successful when compared to other iterative methods on a large number of test problems [16]. They are adaptable to nearly all sorts of problems in optimization, because they need very little information about the problem — especially no derivatives of the objective function. For a list of more than 260 applications of evolutionary algorithms, see [1]. They are capable of solving high dimensional, multimodal, nonlinear problems subject to linear and / or nonlinear constraints. The objective function can also hide the result of a simulation, it does not have to be given in a closed, analytical form. This also holds for the constraints which may represent the outcome of a finite elements method (FEM). Evolution strategies have been adapted to vector optimization problems [11], and they can also serve as a heuristic for NP-complete problems like the *travelling salesman problem* or problems with a noisy or changing response surface.

But apart from focusing on their problem solving capability, one may also concentrate on the conditions under which self-adaptation takes place in evolution strategies. And why not transfer these results — although originating from an ideal and artificial world — to other, real-life domains?

3 Transferring the Observations To ...

3.1 ... Biology

Since evolution strategies rely on the collective learning paradigm of natural populations introduced by Darwin and embedded in today’s synthetic theory of evolution, using these strategies as a model of biological processes seems obvious.

- As stated earlier, a limited life-span enables a species to cope with an environment undergoing change. Accordingly, only a comma strategy allowing each parental generation to reproduce exactly once is able to follow an optimum wandering over time. Holding on to good, but outdated phenotypes the plus variant fails to perform this task.

- So far, evolution strategies only use haploid individuals, thus abandoning the effects of dominance and recessivity. Diploid individuals are not worth the additional computing time in an environment remaining stable over time, i.e. in the case of one objective function. But when evolution strategies try to compute the Pareto set of a vector optimization task, the selection step has to be modified in such a way that each of the k (> 1) objective functions becomes the selection criterion according to a probability vector. From an individual's point of view, the environment changes rather drastically which is why diploid individuals turn out to be necessary in the multiple criteria case. And again, one can find a biological analogy: For certain plants the relation

$$\text{degree of polyploidy in percent} = \text{latitude}$$

holds. Measuring time on a geological scale, in Greenland, the past-glacial environment may be regarded as relatively young compared to the environment in south of Greece which has not changed during the last ice-age.

- The basic mechanisms of organic evolution have undergone an evolution themselves, too. Nature has 'invented' (sexual) recombination as a means of exchanging genetic material between the individuals of one species in order to prevent parasites from adapting to a certain genotype too well. In a difficult topology, similar to the one in Figure 2 (right), recombination also turns out to be crucial for finding the global optimum and not just a local one (convergence security):

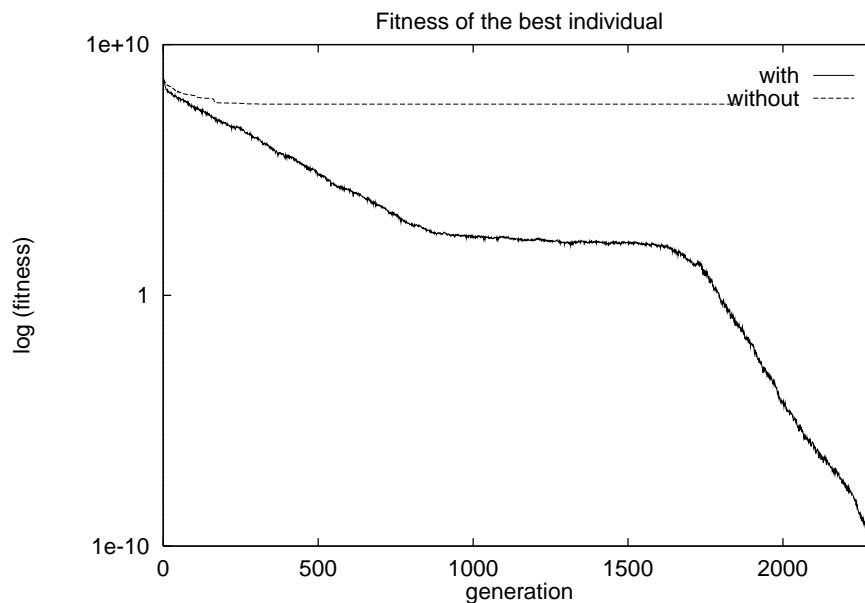


Figure 5: A (15 , 100) strategy with and without recombination

- Schwefel [15] developed a variant of evolution strategies being capable of solving discrete (here: binary) problems by modelling somatic mutations. The observation of an ambiguous relation between an individual's genotype and phenotype forms the biological background of this idea. One may interpret and model this 'disturbed' relation as the result of errors happening during ontogenesis. If this hypothesis of a connection between the genetic mutation rate and the somatic error rate is correct, this model may help to explain the non-genetic variance of morphological attributes (e.g. body length, tail length, weight, ...) observed within populations of genetically *identical* mammals.

Again, within the algorithm, learning takes place on two levels. On the first level, 100 bits have to set correctly. This is achieved by attaching to each bit its private mutation probability. For those bits which have already been set correctly, this rate remains near its starting value ($\approx 10^{-6}$), whereas the wrong bits have to increase their mutability to such a value ($\approx 10^{-2}$) that the desired mutation finally occurs. Afterwards, this value has to decrease again in order to prevent the correct bit from flipping once more. In contrast to the previous one, Figure 6 now demonstrates the influence of recombination on convergence velocity:




Figure 6: A modified (15, 100) strategy using different types of recombination

Setting all mutation probabilities to 10^{-6} , one can expect to perform $\approx 1.800.000$ trials until the last three bits have been set correctly preserving the correct bits.

The modified evolution strategy only needs ≈ 9.000 mutations using recombination, thus emphasizing again that one cannot regard nature as a pure random walk.

3.2 ... Organizational Psychology

Darwin's basic assumption that organisms can only survive if they perform some sort of adaptation within their environment has inspired psychologists who study the organizational structure of companies. Why should *one* organizational structure that handles various situations be sufficient? After all, different environments require different, specially adapted structures. This rather situational or contingency approach has empirically found out that bureaucratic organizations most easily prosper in stationary or settled surroundings. But dynamical environments demand for the internal adaptability of so-called organic organizations relying on well-educated and autonomous, self-supporting members. In this context, knowledge can be regarded as a social product enabled by and permanently expanded and corrected by an exchange of ideas. If practical problems go beyond the limits of one scientific discipline, one does not necessarily need inventions. The required knowledge often already exists, but the respective 'owners' fail to discover its relevance within the given problem's frame of reference.

In evolution strategies, the recombination operator 'exchanges ideas' — the existing knowledge hidden in the stepsizes. As stated earlier, abandoning recombination (cooperation) when solving a simple problem only leads to a loss of efficiency (see Figure 6), but facing a difficult task the solution yielded will be far from the global optimum (see Figure 5).

There are hints that, under certain, but rare conditions, the formation of a group adapts to the task's complexity. In the case of simple problems, centralized network structures develop, whereas complex tasks lead to decentralized structures. Within existing organizations, formal hierarchies normally prevent this process.

The design of a meta-evolution strategy capable of adapting population size, selection pressure and recombination type is currently under development. This algorithm should confirm the following result. In a symmetric, smooth and simple world like the one in Figure 2 (left), a (1, 100) strategy with only one stepsize for all x_i performs far better than a (15, 100) strategy with 30 stepsizes which in turn is more appropriate in difficult, multimodal environments:




Figure 7: A (1, 100) versus a (15, 100) strategy competing in the topologies of Figure 2, $n = 30$

The notions ‘lean management’ and ‘lean production’ may also fit into the frame of this section. They have originated from a 1990 MIT study which tried to convey a Japanese concept — known as the *Toyota production system* since 1973 already — to the U.S. car industry. Often, ‘lean’ is being misunderstood nowadays as a justification of dismissals. Doing so can reduce the unspecific — in an economic sense: unnecessary — resources or the requisite variety in such a way that a change in the outside world becomes more difficult to cope with. But with the true background being the idea of a *lean organization*, one should concentrate on integrating those activities that are not directly involved in the production process, like e.g. surveillance, quality control and maintenance. Shifting these functions to the worker the enhanced responsibility leads to a more interesting job, thus establishing a feedback loop between a worker’s creative abilities and his/her own working conditions. Additionally, regular discussions (‘quality circles’) held across several hierarchical layers promote the distribution of knowledge.

3.3 ... Economics

The main principle of synergetics consists in explaining an observable order from within the system or *endogenously*. In economics, the *invisible hand* (Adam Smith) already conveyed the idea of self-organizing markets in order to explain how individual decisions lead to the formation of structures. With this idea serving as a descriptive ‘theory’ only, the scientific interest shifted towards the analysis of market equilibria using the laws of mechanics. Nowadays, *evolutionary economics* no longer examines the process of converging towards an equilibrium but rather concentrates on the formation and diffusion of innovations. Furthermore, time has become irreversible. One has to take into consideration the history of an economic process when explaining its further development.

With concepts like the ‘bounded rationality’ trying to overcome the idealistic assumption of an agent’s intentional behaviour, why not model individuals stochastically in order to capture the variety of behavioural norms? With a selection step present (externalities) a long-term macroscopic order may emerge which in turn effects or even ‘enslaves’ the ‘fast’ variables on the lower level. A market can remain stable for a certain time, but when approaching one of this stable ‘island’s’ borders, already small variations can lead to a new state of order (bifurcation or path-dependencies, lock-in effects). The victory of the VHS video system over the Beta video system, which was said to be superior technically, depended on a small majority of people preferring VHS in the beginning [5]. Externalities have a strong influence in this example. The more widespread one system is, the better the support from the complementary good markets. Consequently, the decision depends on the perception of the number of systems already sold and the expected utility assigned to it.

Figure 8 displays the performance of an evolution strategy under high selective pressure. One can observe phases of stagnation and phases of rapid change emerging endogenously — a phenomenon biologists refer to as ‘punctuated equilibria’ [7, 8].




Figure 8: A (4, 40) strategy adjusting 30 parameters

Maybe, unwanted discontinuities can be overcome by a closer ‘cooperation’ between economies, too. And maybe, a sustainable growth can only be achieved in this way —

if it exists at all. The last worldwide stock market crash (without any serious economic foundation) has demonstrated how sensitively markets with selfish, maximizing agents, rigid rules and a deceptive security respond to small disturbances. But if the assumption of continuity does not hold [3], maybe recessive phases help to ‘forget’ traditional technologies, firms and institutions, thus releasing niches for the entry of innovative elements. This leads to an enhanced diversity resulting in a more robust economy [9].

The assumption of a nonlinear world surrounding us bears consequences. For we do *not* know exactly what we are doing when interfering with any kind of system comprising nonlinearities. Even the best intentions can lead to unbounded consequences.

References

- [1] **T. Bäck, F. Hoffmeister and H.-P. Schwefel**, *Applications of evolutionary algorithms*, Report of the Systems Analysis Research Group SYS-2/92, Universität Dortmund, Fachbereich Informatik, Feb. 1992.
- [2] **T. Bäck and H.-P. Schwefel**, *An overview of evolutionary algorithms for parameter optimization*, *Evolutionary Computation*, 1 1993, pages 1–23.
- [3] **K. E. Boulding**, *What is evolutionary economics?*, *Journal of Evolutionary Economics*, 1 1991, pages 9–17.
- [4] **H. J. Bremermann**, *Optimization through Evolution and Recombination*, in: *Self-Organizing Systems*, M. C. Yovits, G. T. Jacobi and D. G. Goldstein (Eds.), Spartan Books, Washington, D.C., 1962, pages 93–106.
- [5] **J. Briggs and F. D. Peat**, *Turbulent Mirror — An Illustrated Guide to Chaos Theory and the Science of Wholeness*, Harper & Row Publishers, New York, 1989.
- [6] **G. B. Dantzig**, *Lineare Programmierung und Erweiterungen*, Springer, Berlin, 1966.
- [7] **N. Eldredge and S. J. Gould**, *Punctuated equilibria: An alternative to phyletic gradualism*, in: *Models in Paleobiology*, T. J. M. Schopf (Ed.), Freeman, 1972, Chapter 5, pages 82–115.
- [8] **N. Eldredge and S. J. Gould**, *Punctuated Equilibria: The Tempo and Mode of Evolution Reconsidered*, *Paleobiology*, 3 1977, pages 115–151.
- [9] **G. Eliasson**, *Deregulation, innovative entry and structural diversity*, *Journal of Evolutionary Economics*, 1 1991, pages 49–63.
- [10] **J. H. Holland**, *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, Michigan, 1975.
- [11] **F. Kursawe**, *Evolution strategies for vector optimization*, in: *Proceedings of the Tenth International Conference on Multiple Criteria Decision Making, Taipei, 19.–24.7.1992*, G.-H. TZENG and P. L. YU (Eds.), Taipei, 1992, National Chiao Tung University, pages 187–193.
- [12] ———, *Evolution Strategies — Simple “Models” of Natural Processes?*, in: *Revue Internationale de Systémique*, Volume 7, N° 5, November 1993, pages 627–642.
- [13] **I. Rechenberg**, *Evolutionstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, Frommann-Holzboog, Stuttgart, 1973.
- [14] **R. Rosen**, *Optimality Principles in Biology*, Butterworths, London, 1967.
- [15] **H.-P. Schwefel**, *Binäre Optimierung durch somatische Mutation*, Technischer Bericht, Technische Universität Berlin und Medizinische Hochschule Hannover, May 1975.
- [16] **H.-P. Schwefel**, *Numerische Optimierung von Computer-Modellen mittels der Evolutionstrategie*, Birkhäuser, Basel, 1977.
- [17] ———, *Numerical Optimization of Computer Models*, Wiley & Sons, Chichester, 1981.
- [18] ———, *Collective Phenomena in Evolutionary Systems*, in: *31st Annual Meeting of the International Society for General System Research*, Budapest, 1987, pages 1025–1033.