Why Is Optimization Difficult?

Thomas Weise, Michael Zapf, Raymond Chiong, and Antonio J. Nebro

Abstract This chapter aims to address some of the fundamental issues that
are often encountered in optimization problems, making them difficult to
solve. These issues include premature convergence, ruggedness, causality, de-
ceptiveness, neutrality, epistasis, robustness, overfitting, oversimplification,
multi-objectivity, dynamic fitness, the No Free Lunch Theorem, etc. We ex-
plain why these issues make optimization problems hard to solve and present
some possible countermeasures for dealing with them. By doing this, we hope
to help both practitioners and fellow researchers to create more efficient op-
timization applications and novel algorithms.

1 Introduction

Optimization, in general, is concerned with finding the best solutions for a
given problem. Its applicability in many different disciplines makes it hard

Thomas Weise and Michael Zapf
University of Kassel, Distributed Systems, Wilhelmshöher Allee 73, 34121 Kassel, e-mail:
`weise@vs.uni-kassel.de` and e-mail: `zapf@vs.uni-kassel.de`

Raymond Chiong
School of Computing & Design, Swinburne University of Technology (Sarawak Campus),
93576 Kuching, Sarawak, Malaysia, e-mail: `rchiong@swinburne.edu.my`

Antonio J. Nebro
Dept. Lenguajes y Ciencias de la Computación, ETSI Informática, University of Málaga,
Campus de Teatinos, 29071 Málaga, Spain, e-mail: `antonio@lcc.uma.es`

to give an exact definition. Mathematicians, for instance, are interested in finding the maxima and minima of a real function from within an allowable set of variables. In computing and engineering, the goal is to maximize the performance of a system or application with minimal runtime and resources. In the business industry, people aim to optimize the efficiency of a production process or the quality and desirability of their current products.

All these examples show that optimization is indeed part of our everyday life. We often try to maximize our gain by minimizing the cost we need to bear. However, are we really able to achieve an "optimal" condition? Frankly, whatever problems we are dealing with, it is rare that the optimization process will produce a solution that is truly optimal. It may be optimal for one audience or for a particular application, but definitely not in all cases.

As such, various techniques have emerged for tackling different kinds of optimization problems. In the broadest sense, these techniques can be classified into exact and stochastic algorithms. Exact algorithms, such as branch and bound, $A^\star$ search, or dynamic programming can be highly effective for small-size problems. When the problems are large and complex, especially if they are either NP-complete or NP-hard, i.e., have no known polynomial-time solutions, the use of stochastic algorithms becomes mandatory. These stochastic algorithms do not guarantee an optimal solution, but they are able to find quasi-optimal solutions within a reasonable amount of time.

In recent years, metaheuristics, a family of stochastic techniques, has become an active research area. They can be defined as higher level frameworks aimed at efficiently and effectively exploring a search space [25]. The initial work in this area was started about half a century ago (see [175, 78, 24], and [37]). Subsequently, a lot of diverse methods have been proposed, and today, this family comprises many well-known techniques such as Evolutionary Algorithms, Tabu Search, Simulated Annealing, Ant Colony Optimization, Particle Swarm Optimization, etc.

There are different ways of classifying and describing metaheuristic algorithms. The widely accepted classification would be the view of nature-inspired vs. non nature-inspired, i.e., whether or not the algorithm somehow emulates a process found in nature. Evolutionary Algorithms, the most widely used metaheuristics, belong to the nature-inspired class. Other techniques with increasing popularity in this class include Ant Colony Optimization, Particle Swarm Optimization, Artificial Immune Systems, and so on. Scatter search, Tabu Search, and Iterated Local Search are examples of non nature-inspired metaheuristics. Unified models of metaheuristic optimization procedures have been proposed by Vaessens et al [220, 221], Rayward-Smith [169], Osman [158], and Taillard et al [210].

In this chapter, our main objective is to address some fundamental issues that make optimization problems difficult based on the nature-inspired class of metaheuristics. Apart from the reasons of being large, complex, and dynamic, we present a list of problem features that are often encountered and explain why some optimization problems are hard to solve. Some of the is-

sues that will be discussed, such as multi-modality and overfitting, concern global optimization in general. We will also elaborate on other issues which are often linked to Evolutionary Algorithms, e.g., epistasis and neutrality, but can occur in virtually all metaheuristic optimization processes.

These concepts are important, as neglecting any one of them during the design of the search space and operations or the configuration of the optimization algorithms can render the entire invested effort worthless, even if highly efficient optimization methods are applied. To the best of our knowledge, to date there is not a single document in the literature comprising all such problematic features. By giving clear definitions and comprehensive introductions on them, we hope to create awareness among fellow scientists as well as practitioners in the industry so that they could perform optimization tasks more efficiently.

The rest of this chapter is organized as follows: In the next section, premature convergence to local minima is introduced as one of the major symptoms of failed optimization processes. Ruggedness (Section 3), deceptiveness (Section 4), too much neutrality (Section 5), and epistasis (Section 6), some of which have been illustrated in Fig. 1[1], are the main causes which may lead to this situation. Robustness, correctness, and generality instead are features which we expect from valid solutions. They are challenged by different types of noise discussed in Section 7 and the affinity of overfitting or overgeneralization (see Section 8). Some optimization tasks become further complicated because they involve multiple, conflicting objectives (Section 9) or dynamically changing ones (Section 10). In Section 11, we give a short introduction into the No Free Lunch Theorem, from which we can follow that no panacea, no magic bullet can exist against all of these problematic features. We will conclude our outline of the hardships of optimization with a summary in Section 12.

1.1 Basic Terminology

In the following text, we will utilize a terminology commonly used in Evolutionary Algorithms community and sketched in Fig. 2 based on the example of a simple Genetic Algorithm. The possible solutions x of an optimization problem are elements of the problem space $\mathbb{X}$. Their utility as solutions is evaluated by a set $\mathbf{f}$ of objective functions f which, without loss of generality, are assumed to be subject to minimization. The set of search operations utilized by the optimizers to explore this space does not directly work on

---

[1] We include in Fig. 1 different examples of fitness landscapes, which relate solution candidates (or genotypes) to their objective values. The small bubbles in Fig. 1 represent solution candidates under investigation. An arrow from one bubble to another means that the second individual is found by applying one search operation to the first one. The objective values here are subject to minimization.

Fig. 1.a: Best Case

Fig. 1.b: Low Total Variation

Fig. 1.c: Multimodal

Fig. 1.d: Rugged

Fig. 1.e: Deceptive

Fig. 1.f: Neutral

Fig. 1.g: Needle-In-A-Haystack

Fig. 1.h: Nightmare

Fig. 1: Different possible properties of fitness landscapes (minimization).
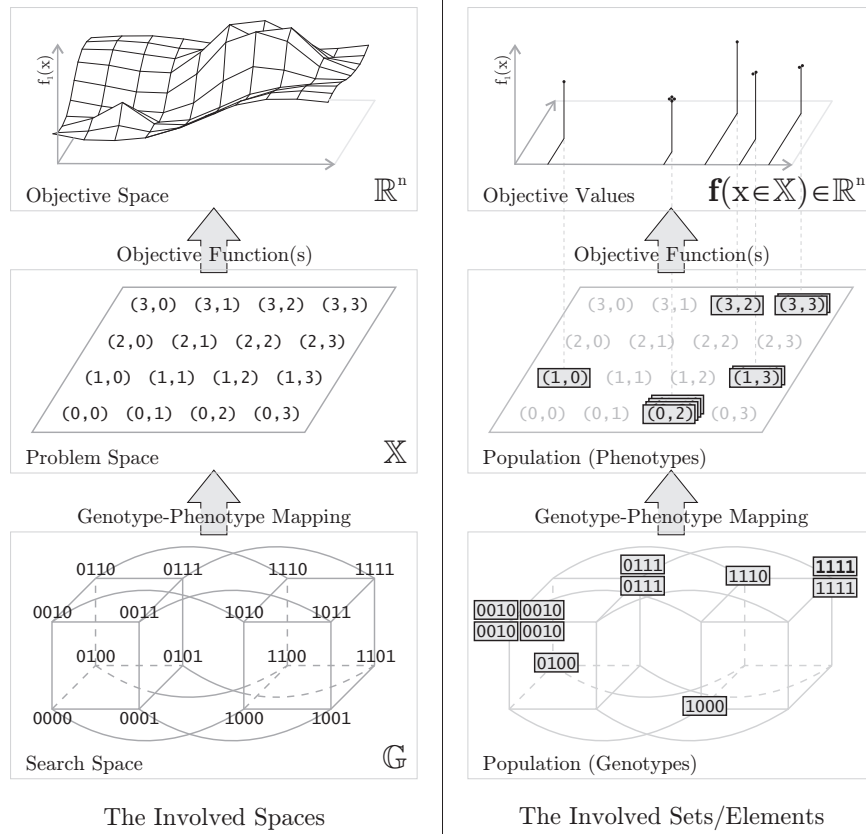
Fig. 2: The involved spaces and sets in optimization.

them. Instead, they are applied to the elements (the genotypes) of the search space $\mathbb{G}$ (the genome). They are mapped to the solution candidates by a genotype-phenotype mapping gpm : $\mathbb{G} \mapsto \mathbb{X}$. The term *individual* is used for both, solution candidates and genotypes.

1.2 The Term "Difficult"

Before we go more into detail about what makes these landscapes *difficult*, we should establish the term in the context of optimization. The degree of difficulty of solving a certain problem with a dedicated algorithm is closely related to its *computational complexity*, i.e., the amount of resources such as

time and memory required to do so. The computational complexity depends on the number of input elements needed for applying the algorithm. This dependency is often expressed in the form of approximate boundaries with the Big-O-family notations introduced by Bachmann [10] and made popular by Landau [122]. Problems can be further divided into *complexity classes*. One of the most difficult complexity classes owning to its resource requirements is NP, the set of all decision problems which are solvable in polynomial time by non-deterministic Turing machines [79]. Although many attempts have been made, no algorithm has been found which is able to solve an NP-complete [79] problem in polynomial time on a deterministic computer. One approach to obtaining near-optimal solutions for problems in NP in reasonable time is to apply metaheuristic, randomized optimization procedures.

As already stated, optimization algorithms are guided by objective functions. A function is *difficult* from a mathematical perspective in this context if it is not continuous, not differentiable, or if it has multiple maxima and minima. This understanding of difficulty comes very close to the intuitive sketches in Fig. 1.

In many real world applications of metaheuristic optimization, the characteristics of the objective functions are not known in advance. The problems are usually NP or have unknown complexity. It is therefore only rarely possible to derive boundaries for the performance or the runtime of optimizers in advance, let alone exact estimates with mathematical precision.

Most often, experience, rules of thumb, and empirical results based on the models obtained from related research areas such as biology are the only guides available. In this chapter we discuss many such models and rules, providing a better understanding of when the application of a metaheuristic is feasible and when not, as well as with indicators on how to avoid defining problems in a way that makes them *difficult*.

2 Premature Convergence

2.1 Introduction

An optimization algorithm has *converged* if it cannot reach new solution candidates anymore or if it keeps on producing solution candidates from a "small"[2] subset of the problem space. Global optimization algorithms will usually converge at some point in time. One of the problems in global optimization is that it is often not possible to determine whether the best solution currently known is situated on a local or a global optimum and thus, if convergence is acceptable. In other words, it is usually not clear whether the

---

[2] according to a suitable metric like numbers of modifications or mutations which need to be applied to a given solution in order to leave this subset

optimization process can be stopped, whether it should concentrate on refining the current optimum, or whether it should examine other parts of the search space instead. This can, of course, only become cumbersome if there are multiple (local) optima, i.e., the problem is multimodal as depicted in Fig. 1.c.

A mathematical function is multimodal if it has multiple maxima or minima [195, 247]. A set of objective functions (or a vector function) $\mathbf{f}$ is multimodal if it has multiple (local or global) optima – depending on the definition of "optimum" in the context of the corresponding optimization problem.

2.2 The Problem

An optimization process has *prematurely converged* to a local optimum if it is no longer able to explore other parts of the search space than the area currently being examined *and* there exists another region that contains a superior solution [192, 219]. Fig. 3 illustrates examples for prematurely convergence.
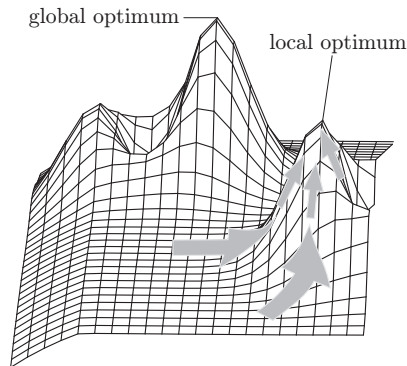
Fig. 3.a: Example 1: Maximization   Fig. 3.b: Example 2: Minimization

Fig. 3: Premature convergence in the objective space.

The existence of multiple global optima itself is not problematic and the discovery of only a subset of them can still be considered as successful in many cases (see Section 9). The occurrence of numerous local optima, however, is more complicated.

The phenomenon of *domino convergence* has been brought to attention by Rudnick [184] who studied it in the context of his BinInt problem [184, 213]. In principle, domino convergence occurs when the solution candidates have features which contribute significantly to different degrees of the total fitness.

If these features are encoded in separate genes (or building blocks) in the genotypes, they are likely to be treated with different priorities, at least in randomized or heuristic optimization methods.

Building blocks with a very strong positive influence on the objective values, for instance, will quickly be adopted by the optimization process (i.e., "converge"). During this time, the alleles of genes with a smaller contribution are ignored. They do not come into play until the optimal alleles of the more "important" blocks have been accumulated. Rudnick [184] called this sequential convergence phenomenon *domino convergence* due to its resemblance to a row of falling domino stones [213].

In the worst case, the contributions of the less salient genes may almost look like noise and they are not optimized at all. Such a situation is also an instance of premature convergence, since the global optimum which would involve optimal configurations of all blocks will not be discovered. In this situation, restarting the optimization process will not help because it will always turn out the same way. Example problems which are often likely to exhibit domino convergence are the Royal Road [139] and the aforementioned BinInt problem [184].

2.3 One Cause: Loss of Diversity

In biology, diversity is *the variety and abundance of organisms at a given place and time* [159, 133]. Much of the beauty and efficiency of natural ecosystems is based on a dazzling array of species interacting in manifold ways. Diversification is also a good investment strategy utilized by investors in the economy in order to increase their profit.

In population-based global optimization algorithms as well, maintaining a set of diverse solution candidates is very important. Losing diversity means approaching a state where all the solution candidates under investigation are similar to each other. Another term for this state is convergence. Discussions about how diversity can be measured have been provided by Routledge [183], Cousins [49], Magurran [133], Morrison and De Jong [148], and Paenke et al [159].

Preserving diversity is directly linked with maintaining a good balance between exploitation and exploration [159] and has been studied by researchers from many domains, such as

- Genetic Algorithms [156, 176, 177],
- Evolutionary Algorithms [28, 29, 123, 149, 200, 206],
- Genetic Programming [30, 38, 39, 40, 53, 93, 94],
- Tabu Search [81, 82], and
- Particle Swarm Optimization [239].

The operations which create new solutions from existing ones have a very large impact on the speed of convergence and the diversity of the populations [69, 203]. The step size in Evolution Strategy is a good example of this issue: setting it properly is very important and leads to the "exploration versus exploitation" problem [102] which can be observed in other areas of global optimization as well.[3]

In the context of optimization, *exploration* means finding new points in areas of the search space which have not been investigated before. Since computers have only limited memory, already evaluated solution candidates usually have to be discarded. Exploration is a metaphor for the procedure which allows search operations to find novel and maybe better solution structures. Such operators (like mutation in Evolutionary Algorithms) have a high chance of creating inferior solutions by destroying good building blocks but also a small chance of finding totally new, superior traits (which, however, is not guaranteed at all).

Exploitation, on the other hand, is the process of improving and combining the traits of the currently known solution(s), as done by the crossover operator in Evolutionary Algorithms, for instance. Exploitation operations often incorporate small changes into already tested individuals leading to new, very similar solution candidates or try to merge building blocks of different, promising individuals. They usually have the disadvantage that other, possibly better, solutions located in distant areas of the problem space will not be discovered.

Almost all components of optimization strategies can either be used for increasing exploitation or in favor of exploration. Unary search operations that improve an existing solution in small steps can be built, hence being exploitation operators (as is done in Memetic Algorithms, for instance). They can also be implemented in a way that introduces much randomness into the individuals, effectively making them exploration operators. Selection operations in Evolutionary Computation choose a set of the most promising solution candidates which will be investigated in the next iteration of the optimizers. They can either return a small group of best individuals (exploitation) or a wide range of existing solution candidates (exploration).

Optimization algorithms that favor exploitation over exploration have higher convergence speed but run the risk of not finding the optimal solution and may get stuck at a local optimum. Then again, algorithms which perform excessive exploration may never improve their solution candidates well enough to find the global optimum or it may take them very long to discover it "by accident". A good example for this dilemma is the Simulated Annealing algorithm [117]. It is often modified to a form called *simulated quenching* which focuses on exploitation but loses the guaranteed convergence to the optimum [110]. Generally, optimization algorithms should employ at least one search operation of explorative character and at least one which is able

---

[3] More or less synonymously to exploitation and exploration, the terms *intensifications* and *diversification* have been introduced by Glover [81, 82] in the context of Tabu Search.

to exploit good solutions further. There exists a vast body of research on the trade-off between exploration and exploitation that optimization algorithms have to face [7, 57, 66, 70, 103, 152].

2.4 Countermeasures

As we have seen, global optimization algorithms are optimization methods for finding the best possible solution(s) of an optimization problem instead of prematurely converging to a local optimum. Still, there is no general approach to ensure their success. The probability that an optimization process prematurely converges depends on the characteristics of the problem to be solved and the parameter settings and features of the optimization algorithms applied [215].

A very crude and yet, sometimes effective measure is restarting the optimization process at randomly chosen points in time. One example for this method is *GRASP*s, *Greedy Randomized Adaptive Search Procedures* [71, 72], which continuously restart the process of creating an initial solution and refining it with local search. Still, such approaches are likely to fail in domino convergence situations.

In order to extend the duration of the evolution in Evolutionary Algorithms, many methods have been devised for steering the search away from areas which have already been frequently sampled. This can be achieved by integrating density metrics into the fitness assignment process. The most popular of such approaches are sharing and niching based on the Euclidean distance of the solution candidates in objective space [55, 85, 104, 138]. Using low selection pressure furthermore decreases the chance of premature convergence but also decreases the speed with which good solutions are exploited.

Another approach against premature convergence is to introduce the capability of self-adaptation, allowing the optimization algorithm to change its strategies or to modify its parameters depending on its current state. Such behaviors, however, are often implemented not in order to prevent premature convergence but to speed up the optimization process (which may lead to premature convergence to local optima) [185, 186, 187].

3 Ruggedness and Weak Causality

3.1 The Problem: Ruggedness

Optimization algorithms generally depend on some form of gradient in the objective or fitness space. The objective functions should be continuous and

exhibit low total variation[4], so the optimizer can descend the gradient easily. If the objective functions are unsteady or fluctuating, i.e., going up and down, it becomes more complicated for the optimization process to find the right directions to proceed to. The more rugged a function gets, the harder it becomes to optimize it. From a simplified point of view, ruggedness is multimodality plus steep ascends and descends in the fitness landscape. Examples of rugged landscapes are Kauffman's fitness landscape [113, 115], the p-Spin model [6], Bergman and Feldman's jagged fitness landscape [19], and the sketch in Fig. 1.d.

3.2 One Cause: Weak Causality

During an optimization process, new points in the search space are created by the search operations. Generally we can assume that the genotypes which are the input of the search operations correspond to phenotypes which have previously been selected. Usually, the better or the more promising an individual is, the higher are its chances of being selected for further investigation. Reversing this statement suggests that individuals which are passed to the search operations are likely to have a good fitness. Since the fitness of a solution candidate depends on its properties, it can be assumed that the features of these individuals are not so bad either. It should thus be possible for the optimizer to introduce slight changes to their properties in order to find out whether they can be improved any further[5]. Normally, such modifications should also lead to small changes in the objective values and, hence, in the fitness of the solution candidate.

Definition 1 (Strong Causality).   *Strong causality* (locality) means that small changes in the properties of an object also lead to small changes in its behavior [170, 171, 180].

This principle (proposed by Rechenberg [170, 171]) should not only hold for the search spaces and operations designed for optimization, but applies to natural genomes as well. The offspring resulting from sexual reproduction of two fish, for instance, has a different genotype than its parents. Yet, it is far more probable that these variations manifest in a unique color pattern of the scales, for example, instead of leading to a totally different creature.

Apart from this straightforward, informal explanation here, causality has been investigated thoroughly in different fields of optimization, such as Evolution Strategy [170, 65], structure evolution [129, 130], Genetic Programming [65, 107, 179, 180], genotype-phenotype mappings [193], search operators [65], and Evolutionary Algorithms in general [65, 182, 207].

---

[4] `http://en.wikipedia.org/wiki/Total_variation` [accessed 2008-04-23]

[5] We have already mentioned this under the subject of exploitation.

In fitness landscapes with weak (low) causality, small changes in the solution candidates often lead to large changes in the objective values, i.e., ruggedness. It then becomes harder to decide which region of the problem space to explore and the optimizer cannot find reliable gradient information to follow. A small modification of a very bad solution candidate may then lead to a new local optimum and the best solution candidate currently known may be surrounded by points that are inferior to all other tested individuals.

The lower the causality of an optimization problem, the more rugged its fitness landscape is, which leads to a degradation of the performance of the optimizer [120]. This does not necessarily mean that it is impossible to find good solutions, but it may take very long to do so.

3.3 Countermeasures

To our knowledge, no viable method which can directly mitigate the effects of rugged fitness landscapes exists. In population-based approaches, using large population sizes and applying methods to increase the diversity can decrease the influence of ruggedness, but only up to a certain degree. Utilizing the Baldwin effect [13, 100, 101, 234] or Lamarckian evolution [54, 234], i.e., incorporating a local search into the optimization process, may further help to smoothen out the fitness landscape [89].

Weak causality is often a home-made problem: it results from the choice of the solution representation and search operations. Thus, in order to apply Evolutionary Algorithms in an efficient manner, it is necessary to find representations which allow for iterative modifications with bounded influence on the objective values.

4 Deceptiveness

4.1 Introduction

Especially annoying fitness landscapes show *deceptiveness* (or deceptivity). The gradient of deceptive objective functions leads the optimizer away from the optima, as illustrated in Fig. 1.e.

The term deceptiveness is mainly used in the Genetic Algorithm community in the context of the Schema Theorem. Schemas describe certain areas (hyperplanes) in the search space. If an optimization algorithm has discovered an area with a better average fitness compared to other regions, it will focus on exploring this region based on the assumption that highly fit areas are likely to contain the true optimum. Objective functions where this is not

the case are called deceptive [20, 84, 127]. Examples for deceptiveness are the ND fitness landscapes [17], trap functions [1, 59, 112] like the one illustrated in Fig. 4, and the fully deceptive problems given by Goldberg et al [86, 60].
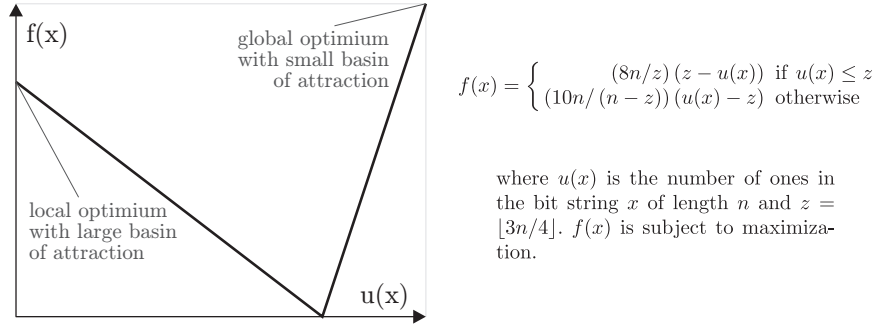

$$f(x) = \begin{cases} (8n/z)\,(z - u(x)) & \text{if } u(x) \leq z \\ (10n/\,(n - z))\,(u(x) - z) & \text{otherwise} \end{cases}$$

where $u(x)$ is the number of ones in the bit string x of length n and $z = \lfloor 3n/4 \rfloor$. $f(x)$ is subject to maximization.

Fig. 4: Ackley's "Trap" function [1, 112].

4.2 The Problem

If the information accumulated by an optimizer actually guides it away from the optimum, search algorithms will perform worse than a random walk or an exhaustive enumeration method. This issue has been known for a long time [228, 140, 141, 212] and has been subsumed under the No Free Lunch Theorem which we will discuss in Section 11.

4.3 Countermeasures

Solving deceptive optimization tasks perfectly involves sampling many individuals with very bad features and low fitness. This contradicts the basic ideas of metaheuristics and thus, there are no efficient countermeasures against deceptivity. Using large population sizes, maintaining a very high diversity, and utilizing linkage learning (see Section 6.3.2) are, maybe, the only approaches which can provide at least a small chance of finding good solutions.

5 Neutrality and Redundancy

5.1 The Problem: Neutrality

We consider the outcome of the application of a search operation to an element of the search space as neutral if it yields no change in the objective values [15, 172]. It is challenging for optimization algorithms if the best solution candidate currently known is situated on a plane of the fitness landscape, i.e., all adjacent solution candidates have the same objective values. As illustrated in Fig. 1.f, an optimizer then cannot find any gradient information and thus, no direction in which to proceed in a systematic manner. From its point of view, each search operation will yield identical individuals. Furthermore, optimization algorithms usually maintain a list of the best individuals found, which will then overflow eventually or require pruning.

The degree of neutrality ν is defined as the fraction of neutral results among all possible products of the search operations Op applied to a specific genotype [15]. We can generalize this measure to areas G in the search space $\mathbb{G}$ by averaging over all their elements. Regions where ν is close to one are considered as *neutral*.

$$\forall g_1 \in \mathbb{G} \Rightarrow \nu(g_1) = \frac{|\{g_2|P(g_2=Op(g_1))>0 \wedge \mathbf{f}(\mathrm{gpm}(g_2))=\mathbf{f}(\mathrm{gpm}(g_1))\}|}{|\{g_2|P(g_2 = Op(g_1)) > 0\}|} \quad (1)$$

$$\forall G \subseteq \mathbb{G} \Rightarrow \nu(G) = \frac{1}{|G|} \sum_{g \in G} \nu(g) \quad (2)$$

5.2 Evolvability

Another metaphor in global optimization borrowed from biological systems is evolvability [52]. Wagner [225, 226] points out that this word has two uses in biology: According to Kirschner and Gerhart [118], a biological system is evolvable if it is able to generate heritable, selectable phenotypic variations. Such properties can then be evolved and changed by natural selection. In its second sense, a system is evolvable if it can acquire new characteristics via genetic change that help the organism(s) to survive and to reproduce. Theories about how the ability of generating adaptive variants has evolved have been proposed by Riedl [174], Altenberg [3], Wagner and Altenberg [227], and Bonner [26], amongst others. The idea of evolvability can be adopted for global optimization as follows:

Definition 2 (Evolvability). The evolvability of an optimization process in its current state defines how likely the search operations will lead to solution candidates with new (and eventually, better) objectives values.

The direct *probability of success* [170, 22], i.e., the chance that search operators produce offspring fitter than their parents, is also sometimes referred to as *evolvability* in the context of Evolutionary Algorithms [2, 5].

5.3 Neutrality: Problematic and Beneficial

The link between evolvability and neutrality has been discussed by many researchers. The evolvability of neutral parts of a fitness landscape depends on the optimization algorithm used. It is especially low for Hill Climbing and similar approaches, since the search operations cannot directly provide improvements or even changes. The optimization process then degenerates to a random walk, as illustrated in Fig. 1.f. The work of Beaudoin et al [17] on the ND fitness landscapes shows that neutrality may "destroy" useful information such as correlation.

Researchers in molecular evolution, on the other hand, found indications that the majority of mutations have no selective influence [77, 106] and that the transformation from genotypes to phenotypes is a many-to-one mapping. Wagner [226] states that neutrality in natural genomes is beneficial if it concerns only a subset of the properties peculiar to the offspring of a solution candidate while allowing meaningful modifications of the others. Toussaint and Igel [214] even go as far as declaring it a necessity for self-adaptation.

The theory of *punctuated equilibria* in biology introduced by Eldredge and Gould [67, 68] states that species experience long periods of evolutionary inactivity which are interrupted by sudden, localized, and rapid phenotypic evolutions [47, 134, 12]. It is assumed that the populations explore neutral layers during the time of stasis until, suddenly, a relevant change in a genotype leads to a better adapted phenotype [224] which then reproduces quickly.

The key to differentiating between "good" and "bad" neutrality is its degree ν in relation to the number of possible solutions maintained by the optimization algorithms. Smith et al [204] have used illustrative examples similar to Fig. 5 showing that a certain amount of neutral reproductions can foster the progress of optimization. In Fig. 5.a, basically the same scenario of premature convergence as in Fig. 3.a is depicted. The optimizer is drawn to a local optimum from which it cannot escape anymore. Fig. 5.b shows that a little shot of neutrality could form a bridge to the global optimum. The optimizer now has a chance to escape the smaller peak if it is able to find and follow that bridge, i.e., the evolvability of the system has increased. If this bridge gets wider, as sketched in Fig. 5.c, the chance of finding the global optimum increases as well. Of course, if the bridge gets too wide, the optimization process may end up in a scenario like in Fig. 1.f where it cannot find any direction. Furthermore, in this scenario we expect the neutral bridge to lead to somewhere useful, which is not necessarily the case in reality.

Fig. 5.a: Premature Convergence

Fig. 5.b: Small Neutral Bridge
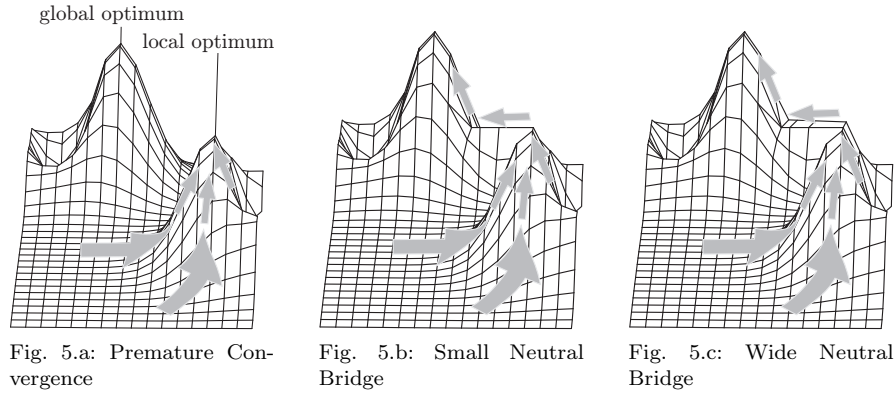
Fig. 5.c: Wide Neutral Bridge

Fig. 5: Possible positive influence of neutrality.

Examples for neutrality in fitness landscapes are the ND family [17], the NKp [15] and NKq [155] models, and the Royal Road [139]. Another common instance of neutrality is *bloat* in Genetic Programming [131].

5.4 Redundancy: Problematic and Beneficial

Redundancy in the context of global optimization is a feature of the genotype-phenotype mapping and means that multiple genotypes map to the same phenotype, i.e., the genotype-phenotype mapping is not injective. The role of redundancy in the genome is as controversial as that of neutrality [230]. There exist many accounts of its positive influence on the optimization process. Shackleton et al [194, 197], for instance, tried to mimic desirable evolutionary properties of RNA folding [106]. They developed redundant genotype-phenotype mappings using voting (both, via uniform redundancy and via a non-trivial approach), Turing machine-like binary instructions, Cellular automata, and random Boolean networks [114]. Except for the trivial voting mechanism based on uniform redundancy, the mappings induced neutral networks which proved beneficial for exploring the problem space. Especially the last approach provided particularly good results [194, 197]. Possibly converse effects like epistasis (see Section 6) arising from the new genotype-phenotype mappings have not been considered in this study.

Redundancy can have a strong impact on the explorability of the problem space. When utilizing a one-to-one mapping, the translation of a slightly modified genotype will always result in a different phenotype. If there exists a many-to-one mapping between genotypes and phenotypes, the search operations can create offspring genotypes different from the parent which

still translate to the same phenotype. The optimizer may now walk along a path through this neutral network. If many genotypes along this path can be modified to different offspring, many new solution candidates can be reached [197]. The experiments of Shipman et al [198, 196] additionally indicate that neutrality in the genotype-phenotype mapping can have positive effects.

Yet, Rothlauf [182] and Shackleton et al [194] show that simple uniform redundancy is not necessarily beneficial for the optimization process and may even slow it down. There is no use in introducing encodings which, for instance, represent each phenotypic bit with two bits in the genotype where `00` and `01` map to `0` and `10` and `11` map to `1`.

5.5 Summary

Different from ruggedness which is always bad for optimization algorithms, neutrality has aspects that may further as well as hinder the process of finding good solutions. Generally we can state that degrees of neutrality ν very close to 1 degenerate optimization processes to random walks. Some forms of neutral networks [14, 15, 27, 105, 208, 222, 223, 238] accompanied by low (nonzero) values of ν can improve the evolvability and hence, increase the chance of finding good solutions.

Adverse forms of neutrality are often caused by bad design of the search space or genotype-phenotype mapping. Uniform redundancy in the genome should be avoided where possible and the amount of neutrality in the search space should generally be limited.

6 Epistasis

6.1 Introduction

In biology, *epistasis* is defined as a form of interaction between different genes [163]. The term was coined by Bateson [16] and originally meant that one gene suppresses the phenotypical expression of another gene. In the context of statistical genetics, epistasis was initially called "epistacy" by Fisher [74]. According to Lush [132], the interaction between genes is epistatic if the effect on the fitness of altering one gene depends on the allelic state of other genes. This understanding of epistasis comes very close to another biological expression: *Pleiotropy*, which means that a single gene influences multiple phenotypic traits [240]. In global optimization, such fine-grained distinctions are usually not made and the two terms are often used more or less synonymously.

Definition 3 (Epistasis). In optimization, Epistasis is the dependency of the contribution of one gene to the value of the objective functions on the allelic state of other genes. [4, 51, 153]

We speak of minimal epistasis when every gene is independent of every other gene. Then, the optimization process equals finding the best value for each gene and can most efficiently be carried out by a simple greedy search [51]. A problem is maximally epistasic when no proper subset of genes is independent of any other gene [205, 153]. Examples of problems with a high degree of epistasis are Kauffman's fitness landscape [113, 115], the p-Spin model [6], and the tunable model of Weise et al [232].

6.2 The Problem

As sketched in Fig. 6, epistasis has a strong influence on many of the previously discussed problematic features. If one gene can "turn off" or affect the expression of many other genes, a modification of this gene will lead to a large change in the features of the phenotype. Hence, the *causality* will be weakened and *ruggedness* ensues in the fitness landscape. On the other hand, subsequent changes to the "deactivated" genes may have no influence on the phenotype at all, which would then increase the degree of *neutrality* in the search space. Epistasis is mainly an aspect of the way in which we define the genome $\mathbb{G}$ and the genotype-phenotype mapping gpm. It should be avoided where possible.
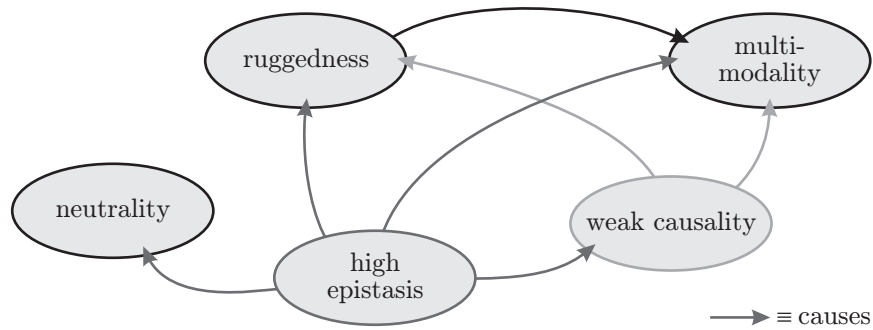

Fig. 6: The influence of epistasis on the fitness landscape.

Generally, epistasis and conflicting objectives in multi-objective optimization should be distinguished from each other. Epistasis as well as pleiotropy is a property of the influence of the elements (the genes) of the genotypes on the phenotypes. Objective functions can conflict *without* the involvement

of any of these phenomena. We can, for example, define two objective functions $f_1(x) = x$ and $f_2(x) = -x$ which are clearly contradicting regardless of whether they are subject to maximization or minimization. Nevertheless, if the solution candidates x as well as the genotypes are simple real numbers and the genotype-phenotype mapping is simply an identity mapping, neither epistatic nor pleiotropic effects can occur.

Naudts and Verschoren [154] have shown for the special case of length-two binary string genomes that deceptiveness does not occur in situations with low epistasis and also that objective functions with high epistasis are not necessarily deceptive. Another discussion about different shapes of fitness landscapes under the influence of epistasis is given by Beerenwinkel et al [18].

6.3 Countermeasures

6.3.1 General

We have shown that epistasis is a root cause for multiple problematic features of optimization tasks. General countermeasures against epistasis can be divided into two groups. The symptoms of epistasis can be mitigated with the same methods which increase the chance of finding good solutions in the presence of ruggedness or neutrality – using larger populations and favoring explorative search operations. Epistasis itself is a feature which results from the choice of the search space structure, the search operations, and the genotype-phenotype mapping. Avoiding epistatic effects should be a major concern during their design. This can lead to a great improvement in the quality of the solutions produced by the optimization process [231]. General advice for good search space design is given in [84, 166, 178] and [229].

6.3.2 Linkage Learning

According to Winter et al [241], *linkage* is "the tendency for alleles of different genes to be passed together from one generation to the next" in genetics. This usually indicates that these genes are closely located in the same chromosome. In the context of Evolutionary Algorithms, this notation is not useful since identifying spatially close elements inside the genotypes is trivial. Instead, we are interested in alleles of different genes which have a joint effect on the fitness [150, 151].

Identifying these linked genes, i.e., learning their epistatic interaction, is very helpful for the optimization process. Such knowledge can be used to protect building blocks from being destroyed by the search operations. Finding approaches for *linkage learning* has become an especially popular discipline in the area of Evolutionary Algorithms with binary [99, 150, 46] and real

[63] genomes. Two important methods from this area are the *messy Genetic Algorithm* (mGA) by Goldberg et al [86] and the *Bayesian Optimization Algorithm* (BOA) [162, 41]. Module acquisition [8] may be considered as a similar effort in the area of Genetic Programming.

Let us take the mGA as an illustrative example for this family of approaches. By explicitly allowing the search operations to rearrange the genes in the genotypes, epistatically linked genes may get located closer to each other by time. As sketched in Fig. 7, the tighter the building blocks are packed, the less likely are they to be destroyed by crossover operations which usually split parent genotypes at randomly chosen points. Hence, the optimization process can strengthen the causality in the search space.

Fig. 7: Two linked genes and their destruction probability under single-point crossover.

7 Noise and Robustness

7.1 Introduction – Noise

In the context of optimization, three types of noise can be distinguished. The first form is noise in the training data used as basis for learning *(i)*. In many applications of machine learning or optimization where a model m for a given system is to be learned, data samples including the input of the system and its measured response are used for training. Some typical examples of situations where training data is the basis for the objective function evaluation are

- the usage of global optimization for building classifiers (for example for predicting buying behavior using data gathered in a customer survey for training),
- the usage of simulations for determining the objective values in Genetic Programming (here, the simulated scenarios correspond to training cases), and
- the fitting of mathematical functions to (x, y)-data samples (with artificial neural networks or symbolic regression, for instance).

Since no measurement device is 100% accurate and there are always random errors, noise is present in such optimization problems.

Besides inexactnesses and fluctuations in the input data of the optimization process, perturbations are also likely to occur during the application of its results. This category subsumes the other two types of noise: perturbations that may arise from inaccuracies in *(ii)* the process of realizing the solutions and *(iii)* environmentally induced perturbations during the applications of the products.

This issue can be illustrated using the process of developing the perfect tire for a car as an example. As input for the optimizer, all sorts of material coefficients and geometric constants measured from all known types of wheels and rubber could be available. Since these constants have been measured or calculated from measurements, they include a certain degree of noise and imprecision *(i)*.

The result of the optimization process will be the best tire construction plan discovered during its course and it will likely incorporate different materials and structures. We would hope that the tires created according to the plan will not fall apart if, accidently, an extra 0.0001% of a specific rubber component is used *(ii)*. During the optimization process, the behavior of many construction plans will be simulated in order to find out about their utility. When actually manufactured, the tires should not behave unexpectedly when used in scenarios different from those simulated *(iii)* and should instead be applicable in all driving scenarios likely to occur.

The effects of noise in optimization have been studied by various researchers; Miller and Goldberg [136, 137], Lee and Wong [125], and Gurin and Rastrigin [92] are some of them. Many global optimization algorithms and theoretical results have been proposed which can deal with noise. Some of them are, for instance, specialized

- Genetic Algorithms [75, 119, 188, 189, 217, 218],
- Evolution Strategies [11, 21, 96], and
- Particle Swarm Optimization [97, 161] approaches.

7.2 The Problem: Need for Robustness

The goal of global optimization is to find the global optima of the objective functions. While this is fully true from a theoretical point of view, it may not suffice in practice. Optimization problems are normally used to find good parameters or designs for components or plans to be put into action by human beings or machines. As we have already pointed out, there will always be noise and perturbations in practical realizations of the results of optimization.

Definition 4 (Robustness). A system in engineering or biology is *robust* if it is able to function properly in the face of genetic or environmental perturbations [225].

Therefore, a local optimum (or even a non-optimal element) for which slight deviations only lead to gentle performance degenerations is usually favored over a global optimum located in a highly rugged area of the fitness landscape [31]. In other words, local optima in regions of the fitness landscape with strong causality are sometimes better than global optima with weak causality. Of course, the level of this acceptability is application-dependent. Fig. 8 illustrates the issue of local optima which are robust vs. global optima which are not. More examples from the real world are:

- When optimizing the control parameters of an airplane or a nuclear power plant, the global optimum is certainly not used if a slight perturbation can have hazardous effects on the system [218].
- Wiesmann et al [235, 236] bring up the topic of manufacturing tolerances in multilayer optical coatings. It is no use to find optimal configurations if they only perform optimal when manufactured to a precision which is either impossible or too hard to achieve on a constant basis.
- The optimization of the decision process on which roads should be precautionary salted for areas with marginal winter climate is an example of the need for dynamic robustness. The global optimum of this problem is likely to depend on the daily (or even current) weather forecast and may therefore be constantly changing. Handa et al [98] point out that it is practically infeasible to let road workers follow a constantly changing plan and circumvent this problem by incorporating multiple road temperature settings in the objective function evaluation.
- Tsutsui et al [218, 217] found a nice analogy in nature: The phenotypic characteristics of an individual are described by its genetic code. During the interpretation of this code, perturbations like abnormal temperature, nutritional imbalances, injuries, illnesses and so on may occur. If the phenotypic features emerging under these influences have low fitness, the organism cannot survive and procreate. Thus, even a species with good genetic material will die out if its phenotypic features become too sensitive to perturbations. Species robust against them, on the other hand, will survive and evolve.

7.3 Countermeasures

For the special case where the problem space corresponds to the real vectors ($\mathbb{X} \subseteq \mathbb{R}^n$), several approaches for dealing with the problem of robustness

Fig. 8: A robust local optimum vs. a "unstable" global optimum.

have been developed. Inspired by Taguchi methods[6] [209], possible disturbances are represented by a vector $\boldsymbol{\delta} = (\delta_1, \delta_2, .., \delta_n)^T$, $\delta_i \in \mathbb{R}$ in the method of Greiner [87, 88]. If the distribution and influence of the δ_i are known, the objective function $f(\mathbf{x}) : \mathbf{x} \in \mathbb{X}$ can be rewritten as $\tilde{f}(\mathbf{x}, \boldsymbol{\delta})$ [236]. In the special case where $\boldsymbol{\delta}$ is normally distributed, this can be simplified to $\tilde{f}\big((x_1 + \delta_1, x_2 + \delta_2, .., x_n + \delta_n)^T\big)$. It would then make sense to sample the probability distribution of $\boldsymbol{\delta}$ a number of t times and to use the mean values of $\tilde{f}(\mathbf{x}, \boldsymbol{\delta})$ for each objective function evaluation during the optimization process. In cases where the optimal value y of the objective function f is known, Equation 3 can be minimized. This approach is also used in the work of Wiesmann et al [235, 236] and basically turns the optimization algorithm into something like a maximum likelihood estimator.

$$f'(\mathbf{x}) = \frac{1}{t} \sum_{i=1}^{t} \left(y - \tilde{f}(\mathbf{x}, \boldsymbol{\delta}_i) \right)^2 \tag{3}$$

This method corresponds to using multiple, different training scenarios during the objective function evaluation in situations where $\mathbb{X} \nsubseteq \mathbb{R}^n$. By adding random noise and artificial perturbations to the training cases, the chance of obtaining robust solutions which are stable when applied or realized under noisy conditions can be increased.

---

[6] `http://en.wikipedia.org/wiki/Taguchi_methods` [accessed 2008-07-19]

8 Overfitting and Oversimplification

In all scenarios where optimizers evaluate some of the objective values of the solution candidates by using training data, two additional phenomena with negative influence can be observed: overfitting and oversimplification.

8.1 Overfitting

8.1.1 The Problem

Definition 5 (Overfitting). Overfitting is the emergence of an overly complicated model (solution candidate) in an optimization process resulting from the effort to provide the best results for as much of the available training data as possible [64, 80, 190, 202].

A model (solution candidate) $m \in \mathbb{X}$ created with a finite set of training data is considered to be overfitted if a less complicated, alternative model $m' \in \mathbb{X}$ exists which has a smaller error for the set of all possible (maybe even infinitely many), available, or (theoretically) producible data samples. This model m' may, however, have a larger error in the training data.

The phenomenon of overfitting is best known and can often be encountered in the field of artificial neural networks or in curve fitting [124, 128, 181, 191, 211]. The latter means we that have a set A of n training data samples (x_i, y_i) and want to find a function f that represents these samples as well as possible, i.e., $f(x_i) = y_i \ \forall \, (x_i, y_i) \in A$.

There exists exactly one polynomial of the degree $n - 1$ that fits to each such training data and goes through all its points. Hence, when only polynomial regression is performed, there is exactly one perfectly fitting function of minimal degree. Nevertheless, there will also be an infinite number of polynomials with a higher degree than $n - 1$ that also match the sample data perfectly. Such results would be considered as overfitted.

In Fig. 9, we have sketched this problem. The function $f_1(x) = x$ shown in Fig. 9.b has been sampled three times, as sketched in Fig. 9.a. There exists no other polynomial of a degree of two or less that fits to these samples than f_1. Optimizers, however, could also find overfitted polynomials of a higher degree such as f_2 which also match the data, as shown in Fig. 9.c. Here, f_2 plays the role of the overly complicated model m which will perform as good as the simpler model m' when tested with the training sets only, but will fail to deliver good results for all other input data.

A very common cause for overfitting is noise in the sample data. As we have already pointed out, there exists no measurement device for physical processes which delivers perfect results without error. Surveys that represent the opinions of people on a certain topic or randomized simulations will ex-

Fig. 9.a: Three sample points of f_1.

Fig. 9.b: $m' \equiv f_1(x) = x$.

Fig. 9.c: $m \equiv f_2(x)$.

Fig. 9: Overfitting due to complexity.

hibit variations from the true interdependencies of the observed entities, too. Hence, data samples based on measurements will always contain some noise.

In Fig. 10 we have sketched how such noise may lead to overfitted results. Fig. 10.a illustrates a simple physical process obeying some quadratic equation. This process has been measured using some technical equipment and the 100 noisy samples depicted in Fig. 10.b has been obtained. Fig. 10.c shows a function resulting from an optimization that fits the data perfectly. It could, for instance, be a polynomial of degree 99 that goes right through all the points and thus, has an error of zero. Although being a perfect match to the measurements, this complicated model does not accurately represent the physical law that produced the sample data and will not deliver precise results for new, different inputs.


Fig. 10.a: The original physical process.

Fig. 10.b: The measurement/training data.

Fig. 10.c: The overfitted result.

Fig. 10: Fitting noise.

From the examples we can see that the major problem that results from overfitted solutions is the loss of generality.

Definition 6 (Generality). A solution of an optimization process is general if it is not only valid for the sample inputs $a_1, a_2, \ldots, a_n$ which were used

for training during the optimization process, but also for different inputs $a \neq a_i \ \forall i : 0 < i \leq n$ if such inputs a exist.

8.1.2 Countermeasures

There exist multiple techniques that can be utilized in order to prevent overfitting to a certain degree. It is most efficient to apply multiple such techniques together in order to achieve best results.

A very simple approach is to restrict the problem space $\mathbb{X}$ in a way that only solutions up to a given maximum complexity can be found. In terms of function fitting, this could mean limiting the maximum degree of the polynomials to be tested. Furthermore, the functional objective functions which solely concentrate on the error of the solution candidates should be augmented by penalty terms and non-functional objective functions putting pressure in the direction of small and simple models [64, 116].

Large sets of sample data, although slowing down the optimization process, may improve the generalization capabilities of the derived solutions. If arbitrarily many training datasets or training scenarios can be generated, there are two approaches which work against overfitting:

1. The first method is to use a new set of (randomized) scenarios for each evaluation of a solution candidate. The resulting objective values may differ largely even if the same individual is evaluated twice in a row, introducing incoherence and ruggedness into the fitness landscape.
2. At the beginning of each iteration of the optimizer, a new set of (randomized) scenarios is generated which is used for all individual evaluations during that iteration. This method leads to objective values which can be compared without bias.

In both cases it is helpful to use more than one training sample or scenario per evaluation and to set the resulting objective value to the average (or better median) of the outcomes. Otherwise, the fluctuations of the objective values between the iterations will be very large, making it hard for the optimizers to follow a stable gradient for multiple steps.

Another simple method to prevent overfitting is to limit the runtime of the optimizers [190]. It is commonly assumed that learning processes normally first find relatively general solutions which subsequently begin to overfit because the noise "is learned", too.

For the same reason, some algorithms allow to decrease the rate at which the solution candidates are modified by time. Such a decay of the learning rate makes overfitting less likely.

If only one finite set of data samples is available for training/optimization, it is common practice to separate it into a set of training data A_t and a set of test cases A_c. During the optimization process, only the training data is used. The resulting solutions are tested with the test cases afterwards. If their

behavior is significantly worse when applied to A_c than when applied to A_t, they are probably overfitted.

The same approach can be used to detect when the optimization process should be stopped. The best known solution candidates can be checked with the test cases in each iteration without influencing their objective values which solely depend on the training data. If their performance on the test cases begins to decrease, there are no benefits in letting the optimization process continue any further.

8.2 Oversimplification

8.2.1 The Problem

Oversimplification (also called overgeneralization) is the opposite of overfitting. Whereas overfitting denotes the emergence of overly-complicated solution candidates, oversimplified solutions are not complicated enough. Although they represent the training samples used during the optimization process seemingly well, they are rough overgeneralizations which fail to provide good results for cases not part of the training.

A common cause for oversimplification is sketched in Fig. 11: The training sets only represent a fraction of the set of possible inputs. As this is normally the case, one should always be aware that such an incomplete coverage may fail to represent some of the dependencies and characteristics of the data, which then may lead to oversimplified solutions. Another possible reason is that ruggedness, deceptiveness, too much neutrality, or high epistasis in the fitness landscape may lead to premature convergence and prevent the optimizer from surpassing a certain quality of the solution candidates. It then cannot completely adapt them even if the training data perfectly represents the sampled process. A third cause is that a problem space which does not include the correct solution was chosen.

Fig. 11.a shows a cubic function. Since it is a polynomial of degree three, four sample points are needed for its unique identification. Maybe not knowing this, only three samples have been provided in Fig. 11.b. By doing so, some vital characteristics of the function are lost. Fig. 11.c depicts a square function – the polynomial of the lowest degree that fits exactly to these samples. Although it is a perfect match, this function does not touch any other point on the original cubic curve and behaves totally differently at the lower parameter area.

However, even if we had included point P in our training data, it would still be possible that the optimization process would yield Fig. 11.c as a result. Having training data that correctly represents the sampled system does not mean that the optimizer is able to find a correct solution with perfect fitness – the other, previously discussed problematic phenomena can

prevent it from doing so. Furthermore, if it was not known that the system which was to be modeled by the optimization process can best be represented by a polynomial of the third degree, one could have limited the problem space $\mathbb{X}$ to polynomials of degree two and less. Then, the result would likely again be something like Fig. 11.c, regardless of how many training samples are used.


Fig. 11.a: The "real system" and the points describing it.

Fig. 11.b: The sampled training data.

Fig. 11.c: The oversimplified result.

Fig. 11: Oversimplification.

8.2.2 Countermeasures

In order to counter oversimplification, its causes have to be mitigated. Generally, it is not possible to have training scenarios which cover the complete input space of the evolved programs. By using multiple scenarios for each individual evaluation, the chance of missing important aspects is decreased. These scenarios can be replaced with new, randomly created ones in each generation, which will decrease this chance even more. The problem space, i.e., the representation of the solution candidates, should further be chosen in a way which allows constructing a correct solution to the problem defined. Then again, releasing too many constraints on the solution structure increases the risk of overfitting and thus, careful proceeding is recommended.

9 Multi-Objective Optimization

9.1 Introduction

Many optimization problems in the real world have k, possibly contradictory objectives f_i which must be optimized simultaneously. Furthermore, the solutions must satisfy m inequality constraints g and p equality constraints h.

A solution candidate x is *feasible*, if and only if $g_i(x) \geq 0 \; \forall i = 1, 2, .., m$ and $h_i(x) = 0 \; \forall i = 1, 2, .., p$ holds. A *multi-objective optimization problem* (MOP) can then be formally defined as follows:

Definition 7 (MOP). Find a solution candidate $x^\star$ in $\mathbb{X}$ which minimizes (or maximizes) the vector function $\mathbf{f}(x^\star) = (f_i(x^\star), f_2(x^\star), .., f_k(x^\star))^T$ and is feasible, (i.e., satisfies the m inequality constraints $g_i(x^\star) \geq 0 \; \forall i = 1, 2, .., m$, the p equality constraints $h_i(x^\star) = 0 \; \forall i = 1, 2, .., p$).

As in single-objective optimization, nature-inspired algorithms are popular techniques to solve such problems. The fact that there are two or more objective functions implies additional difficulties. Due to the contradictory feature of the functions in a MOP and the fact that there exists no total order in $\mathbb{R}^n$ for $n > 1$, the notions of "better than" and "optimum" have to be redefined. When comparing any two solutions x_1 and x_2, solution x_1 can have a better value in objective f_i, i.e., $f_i(x_1) < f_i(x_2)$, while solution x_2 can have a better value in objective f_j. The concepts commonly used here are *Pareto dominance* and *Pareto optimality*.

Definition 8 (Pareto Dominance). In the context of multi-objective global optimization, a solution candidate x_1 is said to *dominate* another solution candidate x_2 (denoted by $x_1 \preccurlyeq x_2$) if and only if $\mathbf{f}(x_1)$ is partially less than $\mathbf{f}(x_2)$, i.e., $\forall i \in \{1, .., k\} \; f_i(x_1) \leq f_i(x_2) \; \wedge \; \exists j \in \{1, .., k\} : \; f_j(x_1) < f_j(x_2)$.


Fig. 12: Some examples for the dominance relation.

The dominance notion allows us to assume that if solution x_1 dominates solution x_2, then x_1 is preferable to x_2. If both solution are non-dominated (such as candidate ① and ② in Fig. 12), some additional criteria have to be used to choose one of them.

Definition 9 (Pareto Optimality). A feasible point $x^\star \in \mathbb{X}$ is *Pareto-optimal* if and only if there is no feasible $x_b \in \mathbb{X}$ with $x_b \preccurlyeq x^\star$.

This definition states that $x^\star$ is Pareto-optimal if there is no other feasible solution x_b which would improve some criterion without causing a simultaneous worsening in at least one other criterion. The solution to a MOP, considering Pareto optimality, is the set of feasible, non-dominated solutions which is known as *Pareto-optimal set*:

Definition 10 (Pareto-Optimal Set). For a given MOP $\mathbf{f}(x)$, the Pareto optimal set is defined as $\mathcal{P}^\star = \{x^\star \in \mathbb{X} | \neg \exists x \in \mathbb{X} : x \preccurlyeq x^\star\}$.

When the solutions in the Pareto-optimal set are plotted in the objective space (as sketched in Fig. 12), they are collectively known as the *Pareto front*:

Definition 11 (Pareto Front). For a given MOP $\mathbf{f}(x)$ and its Pareto-optimal set $\mathcal{P}^\star$, the Pareto front is defined as $\mathcal{PF}^\star = \{\mathbf{f}(x) | x \in \mathcal{P}^\star\}$.

Obtaining the Pareto front of a MOP is the main goal of multi-objective optimization. In a real scenario, the solutions in the Pareto front are sent to an expert in the MOP, the decision maker, who will be responsible for choosing the best tradeoff solution among all of them. Fig. 13 depicts the Pareto front of a bi-objective MOP. In a real problem example, f_1 could represent the time required by a car to cover a given distance, while f_2 could be the fuel consumption.


Fig. 13: Example of Pareto front of a bi-objective MOP.

The Pareto front of a MOP can contain a large (possibly infinite) number of points. Usually, the goal of optimization is to obtain a fixed-size set of solutions called the *Pareto front approximation set*. Population-based algorithms, such as Genetic Algorithms, are very popular to solve MOPs because they can provide an approximation set in a single run.

Given that the goal is to find a Pareto front approximation set, two issues arise. First, the optimization process should *converge* to the true Pareto front and return solutions as close to it as possible. Second, they should be uniformly spread along this front.


Fig. 14.a: Bad Convergence and Good Spread

Fig. 14.b: Good Convergence and Bad Spread

Fig. 14.c: Good Convergence and Spread

Fig. 14: Pareto front approximation sets.

Let us examine the three fronts included in Fig. 14. The first picture (Fig. 14.a) shows an approximation set having a very good spread[7] of solutions, but the points are far away from the true Pareto front. Such results are not attractive because they do not provide Pareto-optimal solutions. The second example (Fig. 14.b) contains a set of solutions which are very close to the true Pareto front but cover it only partially, so the decision maker could lose important trade-off solutions. Finally, the front depicted in Fig. 14.c has the two desirable properties of good convergence and spread.

---

[7] In MO optimization, this property is usually called *diversity*. In order to avoid confusion with the (related) diversity property from Section 2.3, we here use the term *spread* instead.

9.2 The Problem

Features such as multi-modality, deceptiveness, or epistasis found in single-objective optimization also affect MOPs, making them more difficult to solve. However, there are some characteristics that are particular to MOPs. Here we comment on two of them: geometry and dimensionality.

The Pareto front in Fig. 13 has a convex geometry, but there are other different shapes as well. In Fig. 15 we show some examples, including non-convex (concave), disconnected, linear, and non-uniformly distributed Pareto fronts. Besides Pareto optimization, there is a wide variety of other concepts for defining what optima are in the presence of multiple objective functions [45]. The simplest approach is maybe to use a weighted sum of all objective values and set $v(x) = \sum_{i=1}^{k} f_i(x)$. Then the optima would be the element(s) $x^\star$ with $\neg \exists x \in \mathbb{X} : v(x) < v(x^\star)$. However, an optimization process driven by such a linear aggregating function will not find portions of Pareto fronts with non-convex geometry as shown by Das and Dennis [50].



Fig. 15.a: Non-Convex (Concave)


Fig. 15.b: Disconnected

Fig. 15.c: linear

Fig. 15.d: Non-Uniformly Distributed

Fig. 15: Examples of Pareto fronts.

Many studies in the literature consider mainly bi-objective MOPs. As a consequence, many algorithms are designed to deal with that kind of problems. However, MOPs having a higher number of objective functions are common in practice, leading to the so-called *many-objective optimization* [165], which is currently a hot research topic. Most of the optimization algorithms applied today utilize the Pareto dominance relation. When the dimension of the MOPs increases, the majority of solution candidates are non-dominated. As a consequence, traditional nature-inspired algorithms have to be redesigned.

9.3 Countermeasures

In order to obtain an accurate approximation to the true Pareto front, many nature-inspired multi-objective algorithms apply a fitness assignment scheme based on the concept of Pareto dominance, as commented before. For example, NSGA-II [61, 62], the most well-known multi-objective technique, assigns to each solution a rank depending on the number of solutions dominating it. Thus, solutions with rank 1 are non-dominated, solutions with rank 2 are dominated by one solution, and so on. Other algorithms, such as SPEA2 [248, 249] introduce the concept of strength, which is similar to the ranking but also considers the number of dominated solutions.

While the use of Pareto-based ranking methods allows the techniques to search in the direction of finding approximations with good convergence, additional strategies are needed to promote spread. The most commonly adopted approach is to include a kind of density estimator in order to select those solutions which are in the less crowded regions of the objective space. Thus, NSGA-II employs the crowding distance [61] and SPEA2 the distance to the k-nearest neighbor [62].

9.4 Constraint Handling

How the constraints mentioned in Definition 7 are handled is a whole research area in itself with roots in single-objective optimization. Maybe one of the most popular approach for dealing with constraints goes back to Courant [48] who introduced the idea of *penalty functions* [73, 44, 201] in 1943: Consider, for instance, the term $f'(x) = f(x) + v\,[h(x)]^2$ where f is the original objective function, h is an equality constraint, and $v > 0$. If f' is minimized, an infeasible individual will always have a worse fitness than a feasible one with the same objective values.

Besides such static penalty functions, *dynamic* terms incorporating the generation counter [111, 157] or *adaptive* approaches utilizing additional

population statistics [95, 199] have been proposed. Rigorous discussions on penalty functions have been contributed by Fiacco and McCormick [73] and Smith and Coit [201].

During the last fifteen years, many approaches have been developed which incorporate constraint handling and multi-objectivity. Instead of using penalty terms, Pareto ranking can also be extended by additionally comparing individuals according to their feasibility, for instance. Examples for this approach are the Method of Inequalities (MOI) Zakian [246] as used by Pohlheim [164] and the Goal Attainment method defined in [76]. Deb [56, 58] even suggested to simply turn constraints into objective functions in his MOEA version of Goal Programming.

10 Dynamically Changing Fitness Landscape

It should also be mentioned that there exist problems with dynamically changing fitness landscapes [33, 32, 36, 147, 173]. The task of an optimization algorithm is, then, to provide solution candidates with momentarily optimal objective values for each point in time. Here we have the problem that an optimum in iteration t will possibly not be an optimum in iteration $t + 1$ anymore.

The moving peaks benchmarks by Branke [33, 32] and Morrison and De Jong [147] are good examples for dynamically changing fitness landscapes. Such problems with dynamic characteristics can, for example, be tackled with special forms [245] of

- Evolutionary Algorithms [9, 34, 35, 145, 146, 216, 237],
- Genetic Algorithms [83, 119, 142, 143, 144],
- Particle Swarm Optimization [23, 42, 43, 126, 160],
- Differential Evolution [135, 244], and
- Ant Colony Optimization [90, 91]

11 The No Free Lunch Theorem

By now, we know the most important problems that can be encountered when applying an optimization algorithm to a given problem. Furthermore, we have seen that it is arguable what actually an optimum is if multiple criteria are optimized at once. The fact that there is most likely no optimization method that can outperform all others on all problems can, thus, easily be accepted. Instead, there exist a variety of optimization methods specialized in solving different types of problems. There are also algorithms which deliver good results for many different problem classes, but may be outperformed by highly specialized methods in each of them.

These facts have been formalized by Wolpert and Macready [242, 243] in their *No Free Lunch Theorems* (NFL) for search and optimization algorithms. Wolpert and Macready [243] focus on single-objective optimization and prove that the sum of the values of any performance measure (such as the objective value of the best solution candidate discovered until a time step m) over all possible objective functions f is always identical for all optimization algorithms.

From this theorem, we can immediately follow that, in order to outperform the optimization method a_1 in one optimization problem, the algorithm a_2 will necessarily perform worse in another. Fig. 16 visualizes this issue. The higher the value of the performance measure illustrated there, the faster will the corresponding problem be solved. The figure shows that general optimization approaches (like Evolutionary Algorithms) can solve a variety of problem classes with reasonable performance. Hill Climbing approaches, for instance, will be much faster than Evolutionary Algorithms if the objective functions are steady and monotonous, that is, in a smaller set of optimization tasks. Greedy search methods will perform fast on all problems with matroid structure. Evolutionary Algorithms will most often still be able to solve these problems, it just takes them longer to do so. The performance of Hill Climbing and greedy approaches degenerates in other classes of optimization tasks as a trade-off for their high utility in their "area of expertise".

|	
....	random walk or exhaustive enumeration or ...
— —	general optimization algorithm - an EA, for instance
- - - - -	specialized optimization algorithm 1; a hill climber, for instance
————	specialized optimization algorithm 2; a depth-first search, for instance

Fig. 16: A visualization of the No Free Lunch Theorem.

One interpretation of the No Free Lunch Theorem is that it is impossible for any optimization algorithm to outperform random walks or exhaustive enumerations on all possible problems. For every problem where a given method leads to good results, we can construct a problem where the same method has exactly the opposite effect (see Section 4). As a matter of fact, doing so is even a common practice to find weaknesses of optimization algorithms and to compare them with each other.

Another interpretation is that every useful optimization algorithm utilizes some form of problem-specific knowledge. Radcliffe [167] states that without such knowledge, search algorithms cannot exceed the performance of simple enumerations. Incorporating knowledge starts with relying on simple assumptions like "if x is a good solution candidate, than we can expect other good solution candidates in its vicinity", i.e., strong causality. The more (correct) problem specific knowledge is integrated (correctly) into the algorithm structure, the better will the algorithm perform. On the other hand, knowledge correct for one class of problems is, quite possibly, misleading for another class. In reality, we use optimizers to solve a given set of problems and are not interested in their performance when (wrongly) applied to other classes.

Today, there exists a wide range of work on No Free Lunch Theorems for many different aspects of machine learning. The website `http://www.no-free-lunch.org/`[8] gives a good overview about them. Further summaries and extensions have been provided by Köppen et al [121] and Igel and Toussaint [108, 109]. Radcliffe and Surry [168] discuss the NFL in the context of Evolutionary Algorithms and the representations used as search spaces. The No Free Lunch Theorem is furthermore closely related to the Ugly Duckling Theorem proposed by Watanabe [228] for classification and pattern recognition.

12 Concluding Remarks

The subject of this introductory chapter was the question about what makes optimization problems hard, especially for metaheuristic approaches. We have discussed numerous different phenomena which can affect the optimization process and lead to disappointing results. If an optimization process has converged prematurely, it has been trapped in a non-optimal region of the search space from which it cannot "escape" anymore (Section 2). Ruggedness (Section 3) and deceptiveness (Section 4) in the fitness landscape, often caused by epistatic effects (Section 6), can misguide the search into such a region. Neutrality and redundancy (Section 5) can either slow down optimization because the application of the search information does not lead to a gain in information or may also contribute positively by creating neutral networks

---

[8] accessed: 2008-03-28

from which the search space can be explored and local optima can be escaped from. The solutions that are derived, even in the presence of noise, should be robust (Section 7). Also, they should neither be too general (oversimplification, Section 8.2) nor too specifically aligned only to the training data (overfitting, Section 8.1). Furthermore, many practical problems are multi-objective, i.e., involve the optimization of more than one criterion at once (Section 9), or concern objectives which may change over time (Section 10).

In the previous section, we discussed the No Free Lunch Theorem and argued that it is not possible to develop the *one* optimization algorithm, the problem-solving machine which can provide us with near-optimal solutions in short time for every possible optimization task. This must sound very depressing for everybody new to this subject.
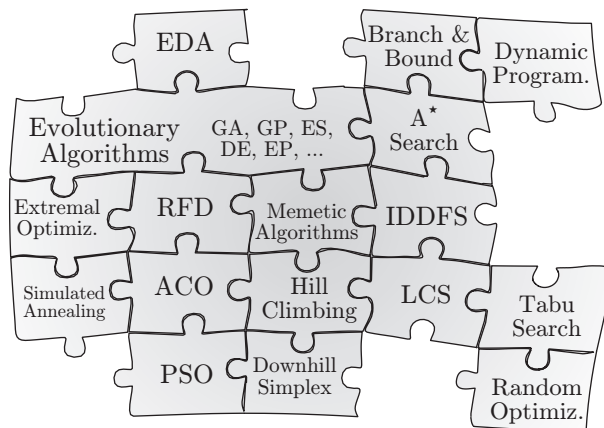

Fig. 17: The puzzle of optimization algorithms.

Actually, quite the opposite is the case, at least from the point of view of a researcher. The No Free Lunch Theorem means that there will always be new ideas, new approaches which will lead to better optimization algorithms to solve a given problem. Instead of being doomed to obsolescence, it is far more likely that most of the currently known optimization methods have at least one niche, one area where they are excellent. It also means that it is very likely that the "puzzle of optimization algorithms" will never be completed. There will always be a chance that an inspiring moment, an observation in nature, for instance, may lead to the invention of a new optimization algorithm which performs better in some problem areas than all currently known ones.

References

[1] Ackley DH (1987) A connectionist machine for genetic hillclimbing, The Springer International Series in Engineering and Computer Science, vol 28. Kluwer Academic Publishers

[2] Altenberg L (1994) The schema theorem and price's theorem. In: Foundations of Genetic Algorithms 3, pp 23–49

[3] Altenberg L (1995) Genome growth and the evolution of the genotype-phenotype map. In: Evolution and Biocomputation – Computational Models of Evolution, Springer-Verlag, pp 205–259

[4] Altenberg L (1996) Nk fitness landscapes. In: Handbook of Evolutionary Computation, Oxford University Press, chap B2.7.2

[5] Altenberg L (1997) Fitness distance correlation analysis: An instructive counterexample. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 57–64

[6] Amitrano C, Peliti L, Saber M (1989) Population dynamics in a spin-glass model of chemical evolution. Journal of Molecular Evolution 29(6):513–525, doi:10.1007/BF02602923

[7] Amor HB, Rettinger A (2005) Intelligent exploration for genetic algorithms: Using self-organizing maps in evolutionary computation. In: Genetic and Evolutionary Computation Conference, GECCO, pp 1531–1538, doi:10.1145/1068009.1068250

[8] Angeline PJ, Pollack J (1993) Evolutionary module acquisition. In: The Second Annual Conference on Evolutionary Programming, Evolutionary Programming Society, pp 154–163

[9] Aragón VS, Esquivel SC (2004) An evolutionary algorithm to track changes of optimum value locations in dynamic environments. Journal of Computer Science & Technology (JCS&T) 4(3):127–133, invited paper

[10] Bachmann PGH (1894) Die Analytische Zahlentheorie / Dargestellt von Paul Bachmann, Zahlentheorie: Versuch einer Gesamtdarstellung dieser Wissenschaft in ihren Haupttheilen, vol Zweiter Theil. B. G. Teubner, Leipzig, Germany

[11] Bäck T, Hammel U (1994) Evolution strategies applied to perturbed objective functions. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 1, pp 40–45, doi:10.1109/ICEC.1994.350045

[12] Bak P, Sneppen K (1993) Punctuated equilibrium and criticality in a simple model of evolution. Physical Review Letters 71:4083–4086, doi:10.1103/PhysRevLett.71.4083

[13] Baldwin JM (1896) A new factor in evolution. The American Naturalist 30:441–451

[14] Barnett L (1997) Tangled webs: Evolutionary dynamics on fitness landscapes with neutrality. Master's thesis, School of Cognitive Science, University of East Sussex, Brighton, UK

[15] Barnett L (1998) Ruggedness and neutrality – the nkp family of fitness landscapes. In: Artificial Life VI: Proceedings of the sixth international conference on Artificial life, pp 18–27

[16] Bateson W (1909) Mendel's Principles of Heredity. Cambridge University Press

[17] Beaudoin W, Verel S, Collard P, Escazut C (2006) Deceptiveness and neutrality the nd family of fitness landscapes. In: Genetic and Evolutionary Computation Conference, GECCO, pp 507–514, doi:10.1145/1143997.1144091

[18] Beerenwinkel N, Pachter L, Sturmfels B (2006) Epistasis and shapes of fitness landscapes. Eprint arXiv:q-bio/0603034 (Quantitative Biology, Populations and Evolution). `http://arxiv.org/abs/q-bio.PE/0603034` [accessed 2007-08-05]

[19] Bergman A, Feldman MW (1992) Recombination dynamics and the fitness landscape. Physica D: Nonlinear Phenomena 56:57–67, doi:10.1016/0167-2789(92)90050-W

[20] Bethke AD (1980) Genetic algorithms as function optimizers. PhD thesis, University of Michigan, Ann Arbor, MI, USA

[21] Beyer HG (1993) Toward a theory of evolution strategies: Some asymptotical results from the $(1, +\lambda)$-theory. Evolutionary Computation 1(2 (Summer)):165–188

[22] Beyer HG (1994) Toward a theory of evolution strategies: The (μ, λ)-theory. Evolutionary Computation 2(4):381–407

[23] Blackwell T (2007) Particle swarm optimization in dynamic environments. In: Evolutionary Computation in Dynamic and Uncertain Environments, Springer, chap 2, pp 29–52

[24] Bledsoe WW, Browning I (1959) Pattern recognition and reading by machine. In: Proceedings of the Eastern Joint Computer Conference (EJCC) – Papers and Discussions Presented at the Joint IRE - AIEE - ACM Computer Conference, pp 225–232

[25] Blum C, Roli A (2003) Metaheuristics in combinatorial optimization: Overview and conceptual comparison. ACM Computing Surveys 35(3):268–308

[26] Bonner JT (1974) On Development: The Biology of Form, new ed edn. Commonwealth Fund Publications, Harvard University Press

[27] Bornberg-Bauer E, Chan HS (1999) Modeling evolutionary landscapes: Mutational stability, topology, and superfunnels in sequence space. Proceedings of the National Academy of Science of the United States of Americs (PNAS) – Biophysics 96(19):10,689–10,694

[28] Bosman PAN, Thierens D (2002) Multi-objective optimization with diversity preserving mixture-based iterated density estimation evolutionary algorithms. International Journal Approximate Reasoning 31(3):259–289, doi:10.1016/S0888-613X(02)00090-7

[29] Bosman PAN, Thierens D (2002) A thorough documentation of obtained results on real-valued continuous and combinatorial multi-

objective optimization problems using diversity preserving mixture-based iterated density estimation evolutionary algorithms. Tech. Rep. UU-CS-2002-052, Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

[30] Brameier MF, Banzhaf W (2002) Explicit control of diversity and effective variation distance in linear genetic programming. In: EuroGP'02: Proceedings of the 5th European Conference on Genetic Programming, pp 37–49

[31] Branke J (1998) Creating robust solutions by means of evolutionary algorithms. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 119–128, doi:10.1007/BFb0056855

[32] Branke J (1999) Memory enhanced evolutionary algorithms for changing optimization problems. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 3, pp 1875–1882, doi:10.1109/CEC.1999.785502

[33] Branke J (1999) The moving peaks benchmark. Tech. rep., Institute AIFB, University of Karlsruhe, Germany, `http://www.aifb.uni-karlsruhe.de/~jbr/MovPeaks/` [accessed 2007-08-19]. Presented in [32]

[34] Branke J (2000) Evolutionary optimization in dynamic environments. PhD thesis, Universität Karlsruhe (TH), Fakultät für Wirtschaftswissenschaften

[35] Branke J (2001) Evolutionary Optimization in Dynamic Environments, Genetic Algorithms and Evolutionary Computation, vol 3. Kluwer Academic Publishers

[36] Branke J, Salihoğlu E, Uyar Ş (2005) Towards an analysis of dynamic environments. In: Genetic and Evolutionary Computation Conference, GECCO, pp 1433–1440

[37] Bremermann HJ (1962) Optimization through evolution and recombination. Self-Organizing systems pp 93–10

[38] Burke EK, Gustafson SM, Kendall G (2002) Survey and analysis of diversity measures in genetic programming. In: Genetic and Evolutionary Computation Conference, GECCO, pp 716–723

[39] Burke EK, Gustafson SM, Kendall G, Krasnogor N (2003) Is increasing diversity in genetic programming beneficial? an analysis of the effects on fitness. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 1398–1405

[40] Burke EK, Gustafson SM, Kendall G (2004) Diversity in genetic programming: An analysis of measures and correlation with fitness. IEEE Transactions on Evolutionary Computation 8(1):47–62, doi:10.1109/TEVC.2003.819263

[41] Cantú-Paz E, Pelikan M, Goldberg DE (2000) Linkage problem, distribution estimation, and bayesian networks. Evolutionary Computation 8(3):311–340, doi:10.1162/106365600750078808

[42] Carlisle AJ (2002) Applying the particle swarm optimizer to nonstationary environments. PhD thesis, Graduate Faculty of Auburn University

[43] Carlisle AJ, Dozier GV (2002) Tracking changing extrema with adaptive particle swarm optimizer. In: Proceedings of the 5th Biannual World Automation Congress, WAC 2002, vol 13, pp 265–270, doi:10.1109/WAC.2002.1049555

[44] Carroll CW (1959) An operations research approach to the economic optimization of a kraft pulping process. PhD thesis, Institute of Paper Chemistry, Appleton, Wisconsin, USA

[45] Ceollo Coello CA, Lamont GB, van Veldhuizen DA (1st ed: 2002, 2nd ed: 2007) Evolutionary Algorithms for Solving Multi-Objective Problems, Genetic and Evolutionary Computation, vol 5. Kluwer Academic Publishers / Springer, doi:10.1007/978-0-387-36797-2

[46] ping Chen Y (2006) Extending the Scalability of Linkage Learning Genetic Algorithms – Theory & Practice, Studies in Fuzziness and Soft Computing, vol 190. Springer, doi:10.1007/b102053

[47] Cohoon JP, Hegde SU, Martin WN, Richards D (1987) Punctuated equilibria: a parallel genetic algorithm. In: Proceedings of the Second International Conference on Genetic algorithms and their Application, pp 148–154

[48] Courant R (1943) Variational methods for the solution of problems of equilibrium and vibrations. Bulletin of the American Mathematical Society 49(1):1–23, doi:10.1090/S0002-9904-1943-07818-4

[49] Cousins SH (1991) Species diversity measurement: Choosing the right index. Trends in Ecology and Evolution (TREE) 6(6):190–192, doi:10.1016/0169-5347(91)90212-G

[50] Das I, Dennis JE (1997) A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. Structural optimization 14(1):63–69, doi:10.1007/BF01197559

[51] Davidor Y (1990) Epistasis variance: A viewpoint on GA-hardness. In: Proceedings of the First Workshop on Foundations of Genetic Algorithms, pp 23–35

[52] Dawkins R (1987) The evolution of evolvability. In: ALIFE – Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, pp 201–220

[53] de Jong ED, Watson RA, Pollack JB (2001) Reducing bloat and promoting diversity using multi-objective methods. In: Genetic and Evolutionary Computation Conference, GECCO, pp 11–18

[54] de Lamarck JBPAdC (1809) Philosophie zoologique – ou Exposition des considérations relatives à l'histoire naturelle des Animaux. Dentu / G. Baillière, Paris, France / Harvard University

[55] Deb K (1989) Genetic algorithms in multimodal function optimization. Master's thesis, The Clearinghouse for Genetic algorithms, University of Alabama, Tuscaloosa, tCGA Report No. 89002

[56] Deb K (1999) Solving goal programming problems using multi-objective genetic algorithms. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 77–84, doi:10.1109/CEC.1999.781910

[57] Deb K (2001) Genetic algorithms for optimization. KanGAL Report 2001002, Kanpur Genetic Algorithms Laboratory (KanGAL), Kanpur, PIN 208 016, India

[58] Deb K (2001) Nonlinear goal programming using multi-objective genetic algorithms. Journal of the Operational Research Society 52(3):291–302

[59] Deb K, Goldberg DE (1993) Analyzing deception in trap functions. In: Foundations of Genetic Algorithms 2, pp 93–108

[60] Deb K, Goldberg DE (1994) Sufficient conditions for deceptive and easy binary functions. Annals of Mathematics and Artificial Intelligence 10(4):385–408

[61] Deb K, Agrawal S, Pratab A, Meyarivan T (2000) A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 849–858, also: KanGAL Report No. 200001

[62] Deb K, Pratab A, Agrawal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6(2):182–197, doi:10.1109/4235.996017

[63] Deb K, Sinha A, Kukkonen S (2006) Multi-objective test problems, linkages, and evolutionary methodologies. In: Genetic and Evolutionary Computation Conference, GECCO, ACM, New York, NY, USA, pp 1141–1148

[64] Dietterich T (1995) Overfitting and undercomputing in machine learning. ACM Computing Surveys (CSUR) 27(3):326–327, doi:10.1145/212094.212114

[65] Droste S, Wiesmann D (1998) On representation and genetic operators in evolutionary algorithms. Tech. Rep. CI–41/98, Fachbereich Informatik, Universität Dortmund

[66] Eiben ÁE, Schippers CA (1998) On evolutionary exploration and exploitation. Fundamenta Informaticae 35(1-4):35–50

[67] Eldredge N, Gould SJ (1972) Punctuated equilibria: an alternative to phyletic gradualism. In: Schopf TJM (ed) Models in Paleobiology, Freeman, Cooper & Co / DoubleDay, chap 5, pp 82–115

[68] Eldredge N, Gould SJ (1977) Punctuated equilibria: The tempo and mode of evolution reconsidered. Paleobiology 3(2):115–151

[69] Eshelman LJ, Schaffer JD (1991) Preventing premature convergence in genetic algorithms by preventing incest. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 115–122

[70] Eshelman LJ, Caruana RA, Schaffer JD (1989) Biases in the crossover landscape. In: Proceedings of the third international conference on Genetic algorithms, pp 10–19

[71] Feo TA, Resende MGC (1995) Greedy randomized adaptive search procedures. Journal of Global Optimization 6(2):109–133, doi:10.1007/BF01096763

[72] Festa P, Resende MG (2004) An annotated bibliography of grasp. AT&T Labs Research Technical Report TD-5WYSEW, AT&T Labs

[73] Fiacco AV, McCormick GP (1968) Nonlinear Programming: Sequential Unconstrained Minimization Techniques. John Wiley & Sons Inc.

[74] Fisher SRA (1918) The correlations between relatives on the supposition of mendelian inheritance. Philosophical Transactions of the Royal Society of Edinburgh 52:399–433

[75] Fitzpatrick JM, Grefenstette JJ (1988) Genetic algorithms in noisy environments. Machine Learning 3(2–3):101–120, doi:10.1007/BF00113893

[76] Fonseca CM, Fleming PJ (1993) Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In: Proceedings of the 5th International Conference on Genetic Algorithms, pp 416–423

[77] Forst CV, Reidys C, Weber J (1995) Evolutionary dynamics and optimization: Neutral networks as model-landscapes for RNA secondary-structure folding-landscapes. In: European Conference on Artificial Life, pp 128–147

[78] Friedberg RM (1958) A learning machine: Part i. IBM Journal of Research and Development 2:2–13

[79] Garey MR, Johnson DS (1979) Computers and Intractability: A Guide to the Theory of NP-Completeness. Series of Books in the Mathematical Sciences, W. H. Freeman & Co., New York, USA

[80] Geman S, Bienenstock E, Doursat R (1992) Neural networks and the bias/variance dilemma. Neural Computation 4(1):1–58

[81] Glover F (1990) Tabu search – part ii. Operations Research Society of America (ORSA) Journal on Computing 2(1):190–206

[82] Glover F, Taillard ÉD, de Werra D (1993) A user's guide to tabu search. Annals of Operations Research 41(1):3–28, doi:10.1007/BF02078647

[83] Gobb HG, Grefenstette JJ (1993) Genetic algorithms for tracking changing environments. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 523–529

[84] Goldberg DE (1989) Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley Longman Publishing Co.

[85] Goldberg DE, Richardson J (1987) Genetic algorithms with sharing for multimodal function optimization. In: Proceedings of the Second

International Conference on Genetic algorithms and their Application, pp 41–49

[86] Goldberg DE, Deb K, Korb B (1989) Messy genetic algorithms: motivation, analysis, and first results. Complex Systems 3:493–530

[87] Greiner H (1994) Robust filter design by stochastic optimization. Proceedings of SPIE (The International Society for Optical Engineering) 2253:150–161, doi:10.1117/12.192107

[88] Greiner H (1996) Robust optical coating design with evolutionary strategies. Applied Optics 35:5477–5483

[89] Gruau F, Whitley LD (1993) Adding learning to the cellular development of neural networks: Evolution and the baldwin effect. Evolutionary Computation 1(3):213–233, doi:10.1162/evco.1993.1.3.213

[90] Guntsch M, Middendorf M (2001) Pheromone modification strategies for ant algorithms applied to dynamic tsp. In: Proceedings of the EvoWorkshops on Applications of Evolutionary Computing, pp 213–222

[91] Guntsch M, Middendorf M, Schmeck H (2001) An ant colony optimization approach to dynamic TSP. In: Genetic and Evolutionary Computation Conference, GECCO, pp 860–867

[92] Gurin LS, Rastrigin LA (1965) Convergence of the random search method in the presence of noise. Automation and Remote Control 26:1505–1511

[93] Gustafson SM (2004) An analysis of diversity in genetic programming. PhD thesis, University of Nottingham, School of Computer Science & IT

[94] Gustafson SM, Ekárt A, Burke EK, Kendall G (2004) Problem difficulty and code growth in genetic programming. Genetic Programming and Evolvable Machines 5(3):271–290, doi:10.1023/B:GENP.0000030194.98244.e3

[95] Hadj-Alouane AB, Bean JC (1992) A genetic algorithm for the multiple-choice integer program. Tech. Rep. 92-50, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbour, MI 48109-2117, USA

[96] Hammel U, Bäck T (1994) Evolution strategies on noisy functions: How to improve convergence properties. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 159–168

[97] Han L, He X (2007) A novel opposition-based particle swarm optimization for noisy problems. In: ICNC'07: Proceedings of the Third International Conference on Natural Computation, vol 3, pp 624–629, doi:10.1109/ICNC.2007.119

[98] Handa H, Lin D, Chapman L, Yao X (2006) Robust solution of salting route optimisation using evolutionary algorithms. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 3098–3105, doi:10.1109/CEC.2006.1688701

[99] Harik GR (1997) Learning gene linkage to efficiently solve problems of bounded difficulty using genetic algorithms. PhD thesis, University of Michigan, Ann Arbor

[100] Hinton GE, Nowlan SJ (1987) How learning can guide evolution. Complex Systems 1:495–502

[101] Hinton GE, Nowlan SJ (1996) How learning can guide evolution. In: Adaptive individuals in evolving populations: models and algorithms, Addison-Wesley Longman Publishing Co., Inc., pp 447–454

[102] Holland JH (1975) Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. The University of Michigan Press, Ann Arbor, reprinted by MIT Press, April 1992, NetLibrary, Inc

[103] Holland JH (1992) Genetic algorithms. Scientific American 267(1):44–50

[104] Horn J, Nafpliotis N, Goldberg DE (1994) A niched pareto genetic algorithm for multiobjective optimization. In: Proceedings of the First IEEE Conference on Evolutionary Computation, vol 1, pp 82–87, doi:10.1109/ICEC.1994.350037

[105] Huynen MA (1996) Exploring phenotype space through neutral evolution. Journal of Molecular Evolution 43(3):165–169, doi:10.1007/PL00006074

[106] Huynen MA, Stadler PF, Fontana W (1996) Smoothness within ruggedness: The role of neutrality in adaptation. Proceedings of the National Academy of Science, USA 93:397–401

[107] Igel C (1998) Causality of hierarchical variable length representations. In: Proceedings of the 1998 IEEE World Congress on Computational Intelligence, pp 324–329

[108] Igel C, Toussaint M (2003) On classes of functions for which no free lunch results hold. Information Processing Letters 86(6):317–321, doi:10.1016/S0020-0190(03)00222-9

[109] Igel C, Toussaint M (2003) Recent results on no-free-lunch theorems for optimization. ArXiv EPrint arXiv:cs/0303032 (Computer Science, Neural and Evolutionary Computing) http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0303032 [accessed 2008-03-28]

[110] Ingber L (1996) Adaptive simulated annealing (asa): Lessons learned. Control and Cybernetics 25(1):33–54

[111] Joines JA, Houck CR (1994) On the use of non-stationary penalty functions to solve nonlinear constrained optimization problems with ga's. In: Proceedings of the First IEEE Conference on Evolutionary Computation, pp 579–584, doi:10.1109/ICEC.1994.349995

[112] Jones T (1995) Evolutionary algorithms, fitness landscapes and search. PhD thesis, The University of New Mexico

[113] Kauffman SA (1988) Adaptation on rugged fitness landscapes. In: Stein DL (ed) Lectures in the Sciences of Complexity: The Proceedings of the 1988 Complex Systems Summer School, Addison Wesley Publishing

Company, Santa Fe Institute Studies in the Sciences of Complexity, vol Lecture I, pp 527–618

[114] Kauffman SA (1993) The Origins of Order: Self-Organization and Selection in Evolution. Oxford University Press

[115] Kauffman SA, Levin SA (1987) Towards a general theory of adaptive walks on rugged landscapes. Journal of Theoretical Biology 128(1):11–45

[116] Kearns MJ, Mansour Y, Ng AY, Ron D (1995) An experimental and theoretical comparison of model selection methods. In: COLT'95: Proceedings of the eighth annual conference on Computational learning theory, ACM Press, pp 21–30, doi:10.1145/225298.225301

[117] Kirkpatrick S, Gelatt, Jr CD, Vecchi MP (1983) Optimization by simulated annealing. Science 220(4598):671–680

[118] Kirschner M, Gerhart J (1998) Evolvability. Proceedings of the National Academy of Science of the USA (PNAS) 95(15):8420–8427

[119] Kita H, Sano Y (2003) Genetic algorithms for optimization of noisy fitness functions and adaptation to changing environments. In: 2003 Joint Workshop of Hayashibara Foundation and 2003 Workshop on Statistical Mechanical Approach to Probabilistic Information Processing (SMAPIP)

[120] Kolarov K (1997) Landscape ruggedness in evolutionary algorithms. In: Proceedings of the IEEE Conference on Evolutionary Computation, pp 19–24

[121] Köppen M, Wolpert DH, Macready WG (2001) Remarks on a recent paper on the "no free lunch" theorems. IEEE Transactions on Evolutionary Computation 5(3):295–296, doi:10.1109/4235.930318

[122] Landau E (1909) Handbuch der Lehre von der Verteilung der Primzahlen. B. G. Teubner, Leipzig, Germany, reprinted by Chelsea, New York, 1953

[123] Laumanns M, Thiele L, Deb K, Zitzler E (2001) On the convergence and diversity-preservation properties of multi-objective evolutionary algorithms. Tech. Rep. 108, Computer Engineering and Networks Laboratory (TIK), Department of Electrical Engineering, Swiss Federal Institute of Technology (ETH) Zurich and Kanpur Genetic Algorithms Laboratory (KanGAL), Department of Mechanical Engineering, Indian Institute of Technology Kanpur

[124] Lawrence S, Giles CL (2000) Overfitting and neural networks: Conjugate gradient and backpropagation. In: Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN'00), IEEE Computer Society, vol 1, pp 1114–1119

[125] Lee JYB, Wong PC (1995) The effect of function noise on gp efficiency. In: Progress in Evolutionary Computation, pp 1–16, doi:10.1007/3-540-60154-6_43

[126] Li X, Jürgen Branke, Blackwell T (2006) Particle swarm with speciation and adaptation in a dynamic environment. In: Genetic

and Evolutionary Computation Conference, GECCO, pp 51–58, doi:10.1145/1143997.1144005

[127] Liepins GE, Vose MD (1991) Deceptiveness and genetic algorithm dynamics. In: Proceedings of the First Workshop on Foundations of Genetic Algorithms (FOGA), pp 36–50

[128] Ling CX (1995) Overfitting and generalization in learning discrete patterns. Neurocomputing 8(3):341–347

[129] Lohmann R (1992) Structure evolution and neural systems. In: Dynamic, Genetic, and Chaotic Programming: The Sixth-Generation, Wiley-Interscience, pp 395–411

[130] Lohmann R (1993) Structure evolution and incomplete induction. Biological Cybernetics 69(4):319–326, doi:10.1007/BF00203128

[131] Luke S, Panait L (2006) A comparison of bloat control methods for genetic programming. Evolutionary Computation 14(3):309–344

[132] Lush JL (1935) Progeny test and individual performance as indicators of an animal's breeding value. Journal of Dairy Science 18(1):1–19

[133] Magurran AE (2005) Biological diversity. Current Biology Magazine 15:R116–R118, doi:10.1016/j.cub.2005.02.006

[134] Martin WN, Lienig J, Cohoon JP (1997) Island (migration) models: Evolutionary algorithms based on punctuated equilibria. In: Handbook of Evolutionary Computation, Oxford University Press, chap 6.3

[135] Mendes R, Mohais AS (2005) Dynde: a differential evolution for dynamic optimization problems. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 3, pp 2808–2815

[136] Miller BL, Goldberg DE (1995) Genetic algorithms, tournament selection, and the effects of noise. IlliGAL Report 95006, Illinois Genetic Algorithms Laboratory, Department of General Engineering, University of Illinois

[137] Miller BL, Goldberg DE (1996) Genetic algorithms, selection schemes, and the varying effects of noise. Evolutionary Computation 4(2):113–131, doi:10.1162/evco.1996.4.2.113

[138] Miller BL, Shaw MJ (1995) Genetic algorithms with dynamic niche sharing for multimodal function optimization. IlliGAL Report 95010, Department of General Engineering, University of Illinois at Urbana-Champaign

[139] Mitchell M, Forrest S, Holland JH (1991) The royal road for genetic algorithms: Fitness landscapes and GA performance. In: Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life, pp 245–254

[140] Mitchell TM (1981) Generalization as search. In: Webber BL, Nilsson NJ (eds) Readings in Artificial Intelligence, 2nd edn, Tioga Pub. Co. Press / Morgan Kaufmann Publishers / Elsevier Science & Technology Books, pp 517–542

[141] Mitchell TM (1982) Generalization as search. Artificial Intelligence 18(2):203–226, doi:10.1016/0004-3702(82)90040-6

[142] Mori N, Kita H, Nishikawa Y (1996) Adaptation to a changing environment by means of the thermodynamical genetic algorithm. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 513–522, doi:10.1007/3-540-61723-X_1015

[143] Mori N, Imanishi S, Kita H, Nishikawa Y (1997) Adaptation to changing environments by means of the memory based thermodynamical genetic algorithm. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 299–306

[144] Mori N, Kita H, Nishikawa Y (1998) Adaptation to a changing environment by means of the feedback thermodynamical genetic algorithm. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 149–158, doi:10.1007/BFb0056858

[145] Morrison RW (2002) Designing evolutionary algorithms for dynamic environments. PhD thesis, George Mason University, USA

[146] Morrison RW (2004) Designing Evolutionary Algorithms for Dynamic Environments, Natural Computing, vol 24(1). Springer

[147] Morrison RW, De Jong KA (1999) A test problem generator for non-stationary environments. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 3, pp 2047–2053, doi:10.1109/CEC.1999.785526

[148] Morrison RW, De Jong KA (2001) Measurement of population diversity. In: Selected Papers of the 5th International Conference on Artificial Evolution, Evolution Artificielle, EA 2001, pp 1047–1074, doi:10.1007/3-540-46033-0_3

[149] Mostaghim S (2004) Multi-objective evolutionary algorithms: Data structures, convergence and, diversity. PhD thesis, Fakultät für Elektrotechnik, Informatik und Mathematik, Universität Paderborn, Deutschland (Germany)

[150] Munetomo M, Goldberg DE (1999) Linkage identification by non-monotonicity detection for overlapping functions. Evolutionary Computation 7(4):377–398, doi:10.1162/evco.1999.7.4.377

[151] Munetomo M, Goldberg DE (1999) Linkage identification by non-monotonicity detection for overlapping functions. IlliGAL Report 99005, Illinois Genetic Algorithms Laboratory (IlliGAL), University of Illinois at Urbana-Champaign

[152] Muttil N, Liong SY (2004) Superior exploration–exploitation balance in shuffled complex evolution. Journal of Hydraulic Engineering 130(12):1202–1205

[153] Naudts B, Verschoren A (1996) Epistasis on finite and infinite spaces. In: Proceedings of the 8th International Conference on Systems Research, Informatics and Cybernetics, pp 19–23

[154] Naudts B, Verschoren A (1999) Epistasis and deceptivity. Bulletin of the Belgian Mathematical Society 6(1):147–154

[155] Newman MEJ, Engelhardt R (1998) Effect of neutral selection on the evolution of molecular species. Proceedings of the Royal Society

of London B (Biological Sciences) 256(1403):1333–1338, doi:10.1098/r-spb.1998.0438

[156] Oei CK, Goldberg DE, Chang SJ (1991) Tournament selection, niching, and the preservation of diversity. IlliGAl Report 91011, Illinois Genetic Algorithms Laboratory (IlliGAL), Department of Computer Science, Department of General Engineering, University of Illinois at Urbana-Champaign

[157] Olsen AL (1994) Penalty functions and the knapsack problem. In: Proceedings of the First IEEE Conference on Evolutionary Computation, vol 2, pp 554–558, doi:10.1109/ICEC.1994.350000

[158] Osman IH (1995) An introduction to metaheuristics. In: Lawrence M, Wilsdon C (eds) Operational Research Tutorial Papers, Stockton Press, Hampshire, UK, pp 92–122, publication of the Operational Research Society, Birmingham, UK

[159] Paenke I, Branke J, Jin Y (2007) On the influence of pheno-type plasticity on genotype diversity. In: First IEEE Symposium on Foundations of Computational Intelligence (FOCI'07), pp 33–40, doi:10.1109/FOCI.2007.372144

[160] Pan G, Dou Q, Liu X (2006) Performance of two improved parti-cle swarm optimization in dynamic optimization environments. In: ISDA'06: Proceedings of the Sixth International Conference on Intel-ligent Systems Design and Applications (ISDA'06), IEEE Computer Society, vol 2, pp 1024–1028

[161] Pan H, Wang L, Liu B (2006) Particle swarm optimization for function optimization in noisy environment. Applied Mathematics and Compu-tation 181(2):908–919, doi:10.1016/j.amc.2006.01.066

[162] Pelikan M, Goldberg DE, Cantú-Paz E (1999) Boa: The bayesian op-timization algorithm. In: Genetic and Evolutionary Computation Con-ference, GECCO, pp 525–532

[163] Phillips PC (1998) The language of gene interaction. Genetics 149(3):1167–1171

[164] Pohlheim H (2005) Geatbx introduction – evolutionary algorithms: Overview, methods and operators. Tech. rep., `http://www.GEATbx.com` [accessed 2007-07-03], documentation for GEATbx version 3.7

[165] Purshouse RC (2003) On the evolutionary optimisation of many ob-jectives. PhD thesis, Department of Automatic Control and Systems Engineering, The University of Sheffield

[166] Radcliffe NJ (1992) Non-linear genetic representations. In: Proceed-ings of the International Conference on Parallel Problem Solving from Nature, PPSN, Elsevier, pp 259–268

[167] Radcliffe NJ (1994) The algebra of genetic algorithms. Annals of Math-ematics and Artificial Intelligence 10(4), doi:10.1007/BF01531276

[168] Radcliffe NJ, Surry PD (1995) Fundamental limitations on search algo-rithms: Evolutionary computing in perspective. In: van Leeuwen J (ed) Computer Science Today – Recent Trends and Developments, Springer-

Verlag, Lecture Notes in Computer Science (LNCS), vol 1000, pp 275–291, doi:10.1007/BFb0015249

[169] Rayward-Smith VJ (1994) A unified approach to tabu search, simulated annealing and genetic algorithms. In: Rayward-Smith VJ (ed) Applications of Modern Heuristic Methods – Proceedings of the UNICOM Seminar on Adaptive Computing and Information Processing, Alfred Waller Ltd / Nelson Thornes Ltd / Unicom Seminars Ltd, Henley-on-Thames, UK, vol I, pp 55–78

[170] Rechenberg I (1973) Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Frommann-Holzboog Verlag, Stuttgart

[171] Rechenberg I (1994) Evolutionsstrategie '94, Werkstatt Bionik und Evolutionstechnik, vol 1. Frommann Holzboog

[172] Reidys CM, Stadler PF (2001) Neutrality in fitness landscapes. Applied Mathematics and Computation 117(2–3):321–350, doi:10.1016/S0096-3003(99)00166-6

[173] Richter H (2004) Behavior of evolutionary algorithms in chaotically changing fitness landscapes. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 111–120

[174] Riedl RJ (1977) A systems-analytical approach to macroevolutionary phenomena. Quarterly Review of Biology pp 351–370

[175] Robbins H, Monro S (1951) A stochastic approximation method. Annals of Mathematical Statistics 22(3):400–407, doi:10.1214/aoms/1177729586

[176] Ronald S (1995) Preventing diversity loss in a routing genetic algorithm with hash tagging. Complexity International 2, `http://www.complexity.org.au/ci/vol02/sr_hash/` [accessed 2008-12-07]

[177] Ronald S (1996) Genetic algorithms and permutation-encoded problems. diversity preservation and a study of multimodality. PhD thesis, University Of South Australia. Department of Computer and Information Science

[178] Ronald S (1997) Robust encodings in genetic algorithms: A survey of encoding issues. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 43–48, doi:10.1109/ICEC.1997.592265

[179] Rosca JP (1995) An analysis of hierarchical genetic programming. Tech. Rep. TR566, The University of Rochester, Computer Science Department

[180] Rosca JP, Ballard DH (1995) Causality in genetic programming. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 256–263

[181] Rosin PL, Fierens F (1995) Improving neural network generalisation. In: Proceedings of the International Geoscience and Remote Sensing Symposium, "Quantitative Remote Sensing for Science and Applications", IGARSS'95, IEEE, vol 2, pp 1255–1257, doi:10.1109/IGARSS.1995.521718

[182] Rothlauf F (2002 (1st ed.), 2006 (2nd ed.)) Representations for Genetic and Evolutionary Algorithms, 2nd edn. Physica-Verlag

[183] Routledge RD (1979) Diversity indices: Which ones are admissible? Journal of Theoretical Biology 76:503–515, doi:10.1016/0022-5193(79)90015-8

[184] Rudnick WM (1992) Genetic algorithms and fitness variance with an application to the automated design of artificial neural networks. PhD thesis, Oregon Graduate Institute of Science & Technology

[185] Rudolph G (1999) Self-adaptation and global convergence: A counter-example. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 1, pp 646–651

[186] Rudolph G (2001) Self-adaptive mutations may lead to premature convergence. IEEE Transactions on Evolutionary Computation 5(4):410–414

[187] Rudolph G (2001) Self-adaptive mutations may lead to premature convergence. Tech. Rep. CI–73/99, Fachbereich Informatik, Universität Dortmund

[188] Sano Y, Kita H (2000) Optimization of noisy fitness functions by means of genetic algorithms using history of search. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 571–580, doi:10.1007/3-540-45356-3_56

[189] Sano Y, Kita H (2002) Optimization of noisy fitness functions by means of genetic algorithms using history of search with test of estimation. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 360–365

[190] Sarle W (2007) What is overfitting and how can i avoid it? Usenet FAQs: compaineural-nets FAQ 3: Generalization(3)

[191] Sarle WS (1995) Stopped training and other remedies for overfitting. In: Proceedings of the 27th Symposium on the Interface: Computing Science and Statistics, pp 352–360

[192] Schaffer JD, Eshelman LJ, Offutt D (1990) Spurious correlations and premature convergence in genetic algorithms. In: Proceedings of the First Workshop on Foundations of Genetic Algorithms (FOGA), pp 102–112

[193] Sendhoff B, Kreutz M, von Seelen W (1997) A condition for the genotype-phenotype mapping: Causality. In: Proceedings of the International Conference on Genetic Algorithms, ICGA, pp 73–80

[194] Shackleton M, Shipman R, Ebner M (2000) An investigation of redundant genotype-phenotype mappings and their role in evolutionary search. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 493–500

[195] Shekel J (1971) Test functions for multimodal search techniques. In: Proceedings of the Fifth Annual Princeton Conference on Information Science and Systems, Princeton University Press, pp 354–359

[196] Shipman R (1999) Genetic redundancy: Desirable or problematic for evolutionary adaptation? In: Proceedings of the 4^{th} International Conference on Artificial Neural Nets and Genetic Algorithms, pp 1–11

[197] Shipman R, Shackleton M, Ebner M, Watson R (2000) Neutral search spaces for artificial evolution: a lesson from life. In: Bedau M, McCaskill JS, Packard NH, Rasmussen S, McCaskill J, Packard N (eds) Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life, The MIT Press, Bradford Books, Complex Adaptive Systems

[198] Shipman R, Shackleton M, Harvey I (2000) The use of neutral genotype-phenotype mappings for improved evolutionary search. BT Technology Journal 18(4):103–111

[199] Siedlecki WW, Sklansky J (1989) Constrained genetic optimization via dynamic reward-penalty balancing and its use in pattern recognition. In: Proceedings of the third international conference on Genetic algorithms, pp 141–150

[200] Singh G, Deb K (2006) Comparison of multi-modal optimization algorithms based on evolutionary algorithms. In: Genetic and Evolutionary Computation Conference, GECCO, pp 1305–1312

[201] Smith AE, Coit DW (1997) Penalty functions. In: Handbook of Evolutionary Computation, Oxford University Press in cooperation with the Institute of Physics Publishing, chap C 5.2

[202] Smith M (1993/1996) Neural Networks for Statistical Modeling. John Wiley & Sons, Inc. / International Thomson Computer Press

[203] Smith SSF (2004) Using multiple genetic operators to reduce premature convergence in genetic assembly planning. Computers in Industry 54(1):35–49

[204] Smith T, Husbands P, Layzell P, O'Shea M (2002) Fitness landscapes and evolvability. Evolutionary Computation 10(1):1–34, doi:10.1162/106365602317301754

[205] Spatz BM, Rawlins GJE (eds) (1990) Proceedings of the First Workshop on Foundations of Genetic Algorithms, Morgan Kaufmann Publishers, Inc.

[206] Spieth C, Streichert F, Speer N, Zell A (2004) Utilizing an island model for ea to preserve solution diversity for inferring gene regulatory networks. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 1, pp 146–151

[207] Stagge P, Igel C (2000) Structure optimization and isomorphisms. In: Theoretical Aspects of Evolutionary Computing, Springer, pp 409–422

[208] Stewart T (2001) Extrema selection: accelerated evolution on neutral networks. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, vol 1, doi:10.1109/CEC.2001.934366

[209] Taguchi G (1986) Introduction to Quality Engineering: Designing Quality into Products and Processes. Asian Productivity Organization /

American Supplier Institute Inc. / Quality Resources / Productivity Press Inc., translation of Sekkeisha no tame no hinshitsu kanri

[210] Taillard ÉD, Gambardella LM, Gendrau M, Potvin JY (2001) Adaptive memory programming: A unified view of metaheuristics. European Journal of Operational Research 135(1):1–16, doi:10.1016/S0377-2217(00)00268-X

[211] Tetko IV, Livingstone DJ, Luik AI (1995) Neural network studies, 1. comparison of overfitting and overtraining. Journal of Chemical Information and Computer Sciences 35(5):826–833

[212] Thierens D (1999) On the scalability of simple genetic algorithms. Tech. Rep. UU-CS-1999-48, Department of Information and Computing Sciences, Utrecht University

[213] Thierens D, Goldberg DE, Pereira ÂG (1998) Domino convergence, drift, and the temporal-salience structure of problems. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 535–540, doi:10.1109/ICEC.1998.700085

[214] Toussaint M, Igel C (2002) Neutrality: A necessity for self-adaptation. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC, pp 1354–1359

[215] Trelea IC (2003) The particle swarm optimization algorithm: convergence analysis and parameter selection. Information Processing Letters 85(6):317–325, doi:10.1016/S0020-0190(02)00447-7

[216] Trojanowski K (1994) Evolutionary algorithms with redundant genetic material for non-stationary environments. PhD thesis, Instytut Podstaw Informatyki PAN, Institute of Computer Science, Warsaw, University of Technology, Poland

[217] Tsutsui S, Ghosh A (1997) Genetic algorithms with a robust solution searching scheme. IEEE Transactions on Evolutionary Computation 1:201–208, doi:10.1109/4235.661550

[218] Tsutsui S, Ghosh A, Fujimoto Y (1996) A robust solution searching scheme in genetic search. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 543–552

[219] Ursem RK (2003) Models for evolutionary algorithms and their applications in system identification and control optimization. PhD thesis, Department of Computer Science, University of Aarhus, Denmark

[220] Vaessens RJM, Aarts EHL, Lenstra JK (1992) A local search template. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 67–76

[221] Vaessens RJM, Aarts EHL, Lenstra JK (1998) A local search template. Computers and Operations Research 25(11):969–979, doi:10.1016/S0305-0548(97)00093-2

[222] van Nimwegen E, Crutchfield JP (2001) Optimizing epochal evolutionary search: Population-size dependent theory. Machine Learning 45(1):77–114, doi:10.1023/A:1012497308906

[223] van Nimwegen E, Crutchfield JP, Huynen M (1999) Neutral evolution of mutational robustness. Proceedings of the National Academy of Science of the United States of Americs (PNAS) – Evolution 96(17):9716–9720

[224] van Nimwegen E, Crutchfield JP, Mitchell M (1999) Statistical dynamics of the royal road genetic algorithm. Theoretical Computer Science 229(1–2):41–102, doi:10.1016/S0304-3975(99)00119-X

[225] Wagner A (2005) Robustness and Evolvability in Living Systems. Princeton Studies in Complexity, Princeton University Press

[226] Wagner A (2005) Robustness, evolvability, and neutrality. FEBS Lett 579(8):1772–1778

[227] Wagner GP, Altenberg L (1996) Complex adaptations and the evolution of evolvability. Evolution 50(3):967–976

[228] Watanabe S (1969) Knowing and Guessing: A Quantitative Study of Inference and Information. John Wiley & Sons

[229] Weicker K (2002) Evolutionäre Algorithmen. Leitfäden der Informatik, B. G. Teubner GmbH

[230] Weicker K, Weicker N (2000) Burden and benefits of redundancy. In: Sixth Workshop on Foundations of Genetic Algorithms (FOGA), Morgan Kaufmann, pp 313–333

[231] Weise T, Zapf M, Geihs K (2007) Rule-based Genetic Programming. In: Proceedings of BIONETICS 2007, 2nd International Conference on Bio-Inspired Models of Network, Information, and Computing Systems

[232] Weise T, Niemczyk S, Skubch H, Reichle R, Geihs K (2008) A tunable model for multi-objective, epistatic, rugged, and neutral fitness landscapes. In: Genetic and Evolutionary Computation Conference, GECCO, pp 795–802

[233] Weise T, Zapf M, Chiong R, Nebro Urbaneja AJ (2009) Why is optimization difficult? In: Chiong R (ed) Nature-Inspired Algorithms for Optimisation, Studies in Computational Intelligence, vol 193, Springer-Verlag, Berlin Heidelberg, chap 1, pp 1–50, doi:10.1007/978-3-642-00267-0_1, URL `http://iao.hfuu.edu.cn/images/publications/WZCNU2009WIOD.pdf`

[234] Whitley LD, Gordon VS, Mathias KE (1994) Lamarckian evolution, the baldwin effect and function optimization. In: Proceedings of the International Conference on Parallel Problem Solving from Nature, PPSN, pp 6–15

[235] Wiesmann D, Hammel U, Bäck T (1998) Robust design of multilayer optical coatings by means of evolutionary algorithms. IEEE Transactions on Evolutionary Computation 2:162–167, doi:10.1109/4235.738986

[236] Wiesmann D, Hammel U, Bäck T (1998) Robust design of multilayer optical coatings by means of evolutionary strategies. Sonderforschungsbereich (sfb) 531, Universität Dortmund

[237] Wilke CO (1999) Evolutionary dynamics in time-dependent environments. PhD thesis, Fakultät für Physik und Astronomie, Ruhr-Universität Bochum

[238] Wilke CO (2001) Adaptive evolution on neutral networks. Bulletin of Mathematical Biology 63(4):715–730

[239] Wilke DN, Kok S, Groenwold AA (2007) Comparison of linear and classical velocity update rules in particle swarm optimization: notes on diversity. International Journal for Numerical Methods in Engineering 70(8):962–984

[240] Williams GC (1957) Pleiotropy, natural selection, and the evolution of senescence. Evolution 11(4):398–411, doi:10.2307/2406060

[241] Winter PC, Hickey GI, Fletcher HL (1st: 1998, 2nd ed: 2002, 3rd ed: 2006) Instant Notes in Genetics. Springer, New York / BIOS Scientific Publishers / Taylor & Francis Ltd.

[242] Wolpert DH, Macready WG (1995) No free lunch theorems for search. Tech. Rep. SFI-TR-95-02-010, The Santa Fe Institute

[243] Wolpert DH, Macready WG (1997) No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation 1(1):67–82, doi:10.1109/4235.585893

[244] Wu N (2006) Differential evolution for optimisation in dynamic environments. Tech. rep., School of Computer Science and Information Technology, RMIT University

[245] Yang S, Ong YS, Jin Y (eds) (2007) Evolutionary Computation in Dynamic and Uncertain Environments, Studies in Computational Intelligence, vol 51(XXIII). Springer

[246] Zakian V (1979) New formulation for the method of inequalities. Proceedings of the Institution of Electrical Engineers 126(6):579–584

[247] Žilinskas A (1978) Algorithm as 133: Optimization of one-dimensional multimodal functions. Applied Statistics 27(3):367–375, doi:10.2307/2347182

[248] Zitzler E, Laumanns M, Thiele L (2001) SPEA2: Improving the Strength Pareto Evolutionary Algorithm. Tech. Rep. 103, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich

[249] Zitzler E, Laumanns M, Thiele L (2001) SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In: Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems. Proceedings of the EUROGEN2001 Conference, pp 95–100

```
@incollection{WZCNU2009WIOD,
 author = {Thomas Weise and Michael Zapf and Raymond Chiong and Antonio Jes{\'u}s {Nebro Urbaneja}},
 title = {Why is optimization difficult?},
 booktitle = {Nature-Inspired Algorithms for Optimisation},
 editor = {Raymond Chiong},
 month = apr,
 year = {2009},
 chapter = {1},
 pages = {1--50},
 series = {Studies in Computational Intelligence},
 volume = {193},
 publisher = {Springer-Verlag},
 address = {Berlin Heidelberg},
 isbn = {978-3-642-00266-3},
 doi = {10.1007/978-3-642-00267-0_1},
 url = {http://iao.hfuu.edu.cn/images/publications/WZCNU2009WIOD.pdf}
}
```