



HAL
open science

A Survey on Recent Progress in the Theory of Evolutionary Algorithms for Discrete Optimization

Benjamin Doerr, Frank Neumann

► **To cite this version:**

Benjamin Doerr, Frank Neumann. A Survey on Recent Progress in the Theory of Evolutionary Algorithms for Discrete Optimization. *ACM Transactions on Evolutionary Learning and Optimization*, 2021, 1 (4), pp.1-43. 10.1145/3472304 . hal-03408953

HAL Id: hal-03408953

<https://hal.science/hal-03408953v1>

Submitted on 29 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Survey on Recent Progress in the Theory of Evolutionary Algorithms for Discrete Optimization

Benjamin Doerr* Frank Neumann†

July 9, 2021

Abstract

The theory of evolutionary computation for discrete search spaces has made significant progress in the last ten years. This survey summarizes some of the most important recent results in this research area. It discusses fine-grained models of runtime analysis of evolutionary algorithms, highlights recent theoretical insights on parameter tuning and parameter control, and summarizes the latest advances for stochastic and dynamic problems. We regard how evolutionary algorithms optimize submodular functions and we give an overview over the large body of recent results on estimation of distribution algorithms. Finally, we present the state of the art of drift analysis, one of the most powerful analysis technique developed in this field.

1 Introduction

Evolutionary computing techniques have been applied in a large variety of different settings ranging from classical optimization problems in the context of supply chain management and renewable energy [BM16, TWD⁺13, NAW20] over to the creation of music and art [Dos13, Lew08, NAN20]. The easy applicability of evolutionary algorithms makes them attractive also to users from outside of computer science disciplines and is one of the major reason for their success in a wide range of engineering applications such as

*Laboratoire d'Informatique (LIX), CNRS, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

†Optimisation and Logistics, School of Computer Science, The University of Adelaide, Adelaide, South Australia, Australia

the design of water networks [BDM15] or processing and planning in mining [MD10, OWBM13, RNRN21].

The theoretical understanding and analysis of evolutionary algorithms is key to further improve the applicability and performance of evolutionary computing methods in a wide range of settings. The area of runtime analysis has played a predominant role during the last 25 years in the area of theory of evolutionary computation when considering discrete optimization problems. This area provides rigorous insights based on mathematical analyses of the search behaviour of evolutionary algorithms. Such theorem-proof style results do not depend on the design of experimental settings which often comes with problems regarding interpretability. While the mathematical results obtained here are often for more restricted settings than in experimental studies, their proofs not only guarantee the result to be correct, but also show why a certain algorithm shows a certain optimization behavior. This way, theoretical research can provide a deeper way of understanding how evolutionary algorithms work.

The goal of this survey is to point out important research directions and their main results of the last 5-10 years. Due to the size of this research area, this survey cannot give a complete description of the progress in this field. We therefore focus on six areas that from our (subjective) point of view are currently the most active and promising ones.

While the area of rigorous runtime analysis provides theoretical insights based on mathematical proofs, it should be noted that there are other approaches contributing to the theoretical understanding of evolutionary algorithms. In particular, important insights into the working behaviour of evolutionary algorithms can be gained through sound experimental studies which investigate important modules and parameters of evolutionary algorithms [BP14, SFAM18]. Furthermore, the area of fitness landscape analysis [PA12, KP19] has contributed significantly to the theoretical understanding of evolutionary computing methods in relation to characteristics of the landscape of optimization problems such as the classical Traveling Salesperson problem [WSOC14, OV18]. Feature-based analysis is another technique which tries to capture features of an optimization problem that make it hard or easy to be optimised by a given algorithm [SvH11, MBT⁺13, NWN15]. All these areas make important contributions to the theoretical understanding of evolutionary computing and set the basis for the novel design of high performing evolutionary algorithms. We concentrate in this survey on results obtained in the area of runtime analysis and recommend the reader to consult the original articles for the previously named areas.

We now give an overview on the different areas covered in this survey. We start with approaches that provide a *fine-grained runtime analysis* of evolutionary computing in Section 2. This area investigates the runtime with respect to the given input size and additional parameters such as a fixed target value or results achievable within a fixed budget. Another important area that provides a fine-grained view is the area of parameterized runtime analysis which takes into account important structural parameters of a given problem instance. These approaches allow one to give a more fine-grained view on the runtime behavior and reveals how structural parameters or computational budgets influence the results. The parameterized analyses carried out in this area are focused on classical combinatorial optimization problems such as minimum vertex cover and the Euclidean Traveling Salesperson problems and we will summarize the main results for them.

Setting the *parameters of evolutionary algorithms* is a key challenge in the profitable use of these heuristics. In Section 3, we discuss how recent theoretical works suggest to set the parameters. We also discuss different ways to let the algorithm optimize its parameters itself, which currently appears as a very powerful, easy-to-use approach.

Dynamic and stochastic problems play a key role in many real-world applications [RP13, RNRN21, XNN21] and evolutionary algorithms have been shown to be very successful in dynamic and stochastic environments [NY12, RKD17]. The theoretical investigations in terms of runtime analysis for such problems have been started by Droste [Dro02, Dro03, Dro04] in the early 2000s and a wide range of results have been obtained during the last 10 years. We will summarize such results in Section 4.

Many important problems can be formulated in terms of a *submodular functions* with a given set of constraints. The analysis and the design of evolutionary algorithms for submodular optimization problems has gained a lot of attention during the last 5 years. Various types of constraints as well as dynamic and stochastic settings have been investigated and provably efficient evolutionary algorithms outperforming previous state-of-the-art greedy approaches have been designed. The most important results and the different areas investigated are presented in Section 5.

Estimation-of-distribution algorithms (EDAs) are evolutionary algorithms which do not evolve a population of good solution candidates, but a probability distribution on the search space that allows one to sample good solutions. Due to the complicated nature of the underlying mathematical objects (a random process taking probability distributions as states), for a long time the theoretical understanding of these algorithms was very limited. The last few years, however, have seen great progress in this topic, both showing

new advantages of EDAs such as robustness to noise and giving advice in how to set their parameters. We review some of these results in Section 6.

We end this survey with a discussion of an important analysis tools developed in the EA community. The area of *drift analysis* has provided a wide range of analytical methods that allow to translate an understanding of the typical progress of an EA into information about its runtime. We will summarize the most important drift theorems and their applications together with the challenges involved when using drift analysis in Section 7.

2 Fine-grained Runtime Analysis of Randomized Search Heuristics

Traditional runtime analysis investigates the runtime of an evolutionary computing technique on a given problem with respect to the size of problem instance. Usually, a worst-case view is taken and upper bounds on the runtime of all problem instances of a given size are proven.

Many important optimization problems can be stated as binary problems. A binary problem is defined on the search space $\{0, 1\}^n$ and a problem instance is given as a pseudo-Boolean functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$. We call n the size of the instance. A common goal is to analyse the expected number of solutions that an evolutionary algorithm generates until an optimal solution with respect to f has been produced for the first time. In general, the runtime is measured in terms of the number of solutions that are produced and evaluated before the considered algorithm has obtained a given goal for the first time. The *expected optimization time* refers to the expected number of constructed solutions until an optimal solution has been produced for the first time.

Studies focus on simplified algorithms such as the (1+1) EA which is a special case of the $(\mu+1)$ EA shown in Algorithm 1 for maximizing a pseudo-Boolean function f . Randomized local search (RLS) is an even simpler variant and differs from the (1+1) EA by flipping in each iteration exactly 1 bit chosen uniformly at random. For problems where 1-bit flips leads to obvious local optima a combination of 1-bit and 2-bit flips is considered [NW07].

Early analyses reveal that the expected time until RLS and the classical (1+1) EA has obtained an optimal solution for the maximization of the benchmark function

$$OneMax(x) = \sum_{i=1}^n x_i$$

Algorithm 1: $(\mu + 1)$ -EA for maximization problems

- 1 Choose a population P consisting of μ solutions chosen from $\{0, 1\}^n$ uniformly at random;
 - 2 **repeat** *forever*
 - 3 Choose $x \in P$ randomly;
 - 4 Create x' by flipping each bit of x independently with probability $1/n$;
 - 5 Add $\{x'\}$ to P .;
 - 6 Remove an element y from P for which $y = \arg \min_{x \in P} f(x)$ holds;
-

is $\Theta(n \log n)$. Similarly, the expected runtime of RLS and the (1+1) EA for the maximization of the benchmark function

$$\text{LeadingOnes}(x) = \sum_{i=1}^n \prod_{j=1}^i x_j$$

is $\Theta(n^2)$. Studies of this type have been frequently carried out in the area of runtime analysis over the last 25 year. Furthermore, the research area discusses the time to obtain an optimal solution or the time to achieve a good approximation in the case of NP-hard problems. One drawback of the described approach is that it is not able to distill other important characteristics that make a problem instance hard or easy to solve.

During the last years, different approaches have been taken to give a more precise view on the runtime behavior of evolutionary algorithms. This includes approaches for evaluating the runtime until a fixed target value (instead of an optimal value) is achieved. Furthermore, the parameterized analysis allows to obtain results in terms of additional structural parameters. We start by summarizing some important results in the area of fixed budget and fixed target analysis and discuss parameterized results afterwards.

2.1 Fixed Budget and Fixed Target Analysis

Fixed budget analysis has been introduced in [JZ12] and considers the quality that an evolutionary algorithm achieves within a given time budget b . Classical runtime analysis often already provides the insights for such fixed budget analysis as it analyzes the progress that an evolutionary algorithm makes at the different stages of the optimization process. Taking a more fine-grained view on such classical runtime results, fixed budget results can often

be obtained. In [JZ12], results for RLS and the (1+1) EA on the benchmark problems OneMax and LeadingOnes are obtained which are based on more fine-grained insights of previous studies on the expected optimization time. An extension of these fixed budget results has been provided in [JZ14a]. Furthermore, a general method for obtaining fixed budget results based on classical runtime analysis results has been presented in [DJWZ13]. Fixed budget analysis can also be used to provide further insights of other types of randomized search heuristics. For example, artificial immune systems have been analyzed using fixed budget analysis in [JZ14b]. The authors show how to translate previous runtime results on the use of hypermutations in artificial immune systems to results when given a fixed budget.

Related to fixed budget analysis is fixed target analysis [BDDV20]. Instead of considering the solution quality obtainable within a given time budget, fixed target analysis considers the time needed to achieve a given target quality k . The runtime analysis is therefore parameterized by k . Results have been obtained for well studied benchmark problems such as OneMax and LeadingOnes. Furthermore, fixed target results for the classical minimum spanning tree problem have been obtained. In [ABD20b], the dual question of how long an EA needs to find the optimum when already starting with a good solution was investigated. Surprisingly, different types of EAs profit to very different extents from such a favorable initialization. Also different parameter values become optimal when starting with a good instead of a random solution.

2.2 Parameterized Analysis

Parameterized analysis of algorithms [DF99] allows to investigate algorithms not just with respect to the worst-case behaviour regarding the length of the given input, usually denoted by n , but also with respect to some additional parameter(s) that characterize the problem. A problem is defined to be fixed parameter tractable (FPT) with respect to a parameter k iff there is an algorithm that runs in time $O(\text{poly}(n) \cdot f(k))$, where $f(k)$ is a function only depending on k . We call an algorithm an FPT algorithm with respect to a parameter k iff it runs in time $O(\text{poly}(n) \cdot f(k))$. This implies that an FPT algorithm runs in polynomial time if k is constant.

The approach of analyzing evolutionary algorithms in the context of parameterized complexity has been introduced by Kratsch and Neumann [KN13] although there are earlier analyses that investigate RLS and the (1+1) EA for the maximum clique problem on planar graphs [Sto06] and population-based EAs with respect to the runtime dependent on the size of the cliques obtained [Sto07]. We call an evolutionary algorithm a

Algorithm 2: GSEMO for minimization problems

```
1 Choose an initial solution  $x \in \{0, 1\}^n$  uniformly at random;
2 Determine  $f(x)$  and initialize  $P \leftarrow \{x\}$ ;
3 repeat forever
4   | Choose  $x \in P$  randomly;
5   | Create  $x'$  by flipping each bit of  $x$  independently with probability
   |    $1/n$ ;
6   | Determine  $f(x')$ ;
7   | if  $\exists x'' \in P, f(x'') \leq f(x')$  and  $f(x'') \neq f(x')$  then
8   |   |  $P$  is unchanged
9   | else
10  |   | exclude all  $x''$  where  $f(x') \leq f(x'')$  from  $P$  and add  $x'$  to  $P$ 
```

fixed-parameter evolutionary algorithm with respect to a parameter k iff its expected optimization time is $O(\text{poly}(n) \cdot f(k))$.

2.2.1 Results for Minimum Vertex Cover

The minimum vertex cover problem is the classical problem in the area of parameterized complexity and several FPT algorithms are available. The input of the minimum vertex cover problem is an undirected graph $G = (V, E)$ and the goal is to find a minimum set of vertices $V' \subseteq V$ such that each edge is covered by at least one node of V' , i.e. $e \cap V' \neq \emptyset$ holds for all $e \in E$.

Kratsch and Neumann [KN13] showed that a simple evolutionary multi-objective algorithm called GSEMO (see Algorithm 2) frequently used in the area of runtime analysis [NW06, GL10, FHH⁺10] is able to compute a kernelization for the problem. A kernelization is a reduced problem where the decision for some nodes whether or not to include them has already been made in an optimal way. It is shown in [KN13] that such a kernelization can be obtained by using two different types of helper objectives as a second objective. The first one considered in the article is the number of uncovered edges of a given solution x . The second approach uses the optimal value of the linear programming relaxation of the graph consisting only of the uncovered edges of a given solution x . Note that both helper objectives estimate the degree of infeasibility of a solution x which is quite common when using multi-objective models for single-objective optimization problems in the context of evolutionary computing [NW06, FHH⁺10].

Having obtained such a solution an alternative mutation operator flipping bits corresponding to nodes that are adjacent to so far uncovered edges can obtain an optimal solution in time $O(f(k) \cdot \text{poly}(n))$ which leads to the result that the examined evolutionary algorithms are fixed parameter evolutionary algorithms. It has also been shown that a factor $(1 + \epsilon)$ -approximation, $0 \leq \epsilon \leq 1$, can be obtained in expected time $O(n^2 \cdot \log n + OPT \cdot n^2 + n \cdot 4^{(1-\epsilon) \cdot OPT})$, where OPT is the value of an optimal solution, when using the LP relaxation as the second objective. This gives a trade-off between approximation quality and runtime. Setting $\epsilon = 1$, it shows that the approach computes a factor 2-approximation in expected polynomial time.

In the weighted vertex cover problem, each node has a positive weight and the goal is to minimize the sum of the weights of the chosen nodes under the condition that all edges are covered. The use of the dual formulation of the vertex cover in form of edge sets has been investigated by Pourhassan et al. [PSN19]. They have generalized the edge based presentation by Jansen et al. [JOZ13] to the weighted case and shown that their evolutionary multi-objective algorithm is a fixed parameter evolutionary algorithm for the weighted vertex cover problem. The authors have shown that a 2-approximation for the weighted vertex cover problem is obtained by the algorithm in expected polynomial time and presented a population-based approach which achieves a $(1 + \epsilon)$ -approximation in expected time $O(n \cdot 2^{\min\{n, 2^{(1-\epsilon)OPT}\}} + n^3)$. Setting $\epsilon = 1$, it shows that the approach computes a factor 2-approximation in expected polynomial time for the weighted vertex cover problem.

2.2.2 Results for the Euclidean Traveling Salesperson Problem

The traveling salesperson problem is another very prominent problem in the area of combinatorial optimization. Given a set of n cities $i = 1, \dots, n$, and distances $d(i, j)$ between them the goal is to compute a tour of minimal cost visiting each city exactly once and returning to the origin. A possible solution for the TSP is usually given by a permutation $\pi = (\pi(1), \dots, \pi(n))$ of the given n cities and the goal is to find a tour π that minimizes

$$c(\pi) = d(\pi(n), \pi(1)) + \sum_{i=1}^{n-1} d(\pi(i), \pi(i+1)).$$

In the context of parameterized analysis of evolutionary algorithms, the Euclidean TSP has been investigated by Sutton et al. [SNN14]. Here each city i is given as coordinates (x_i, y_i) and the distance between city i and j is given as $d(i, j) = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2}$. The Euclidean TSP is still NP-hard but admits a polynomial time approximation scheme (PTAS) [Aro98].

In terms of parameterized analysis, the impact of the number of inner points has been considered which is given by the number of points that do not lie on the convex hull of the points in 2D. We denote by $n - k$ the points on the convex hull and k the number of inner points.

The Euclidean TSP can be solved by classical algorithms in time $O(\text{poly}(n) \cdot f(k))$ using dynamic programming [DHOW06]. This makes use of the properties that an optimal solution has to visit the points of the convex hull in the order as they appear on the hull. The difficult task is then to "fill in" the inner points such that an optimal solution is obtained.

Investigations in the area of evolutionary algorithms focused on runtime analysis with respect to the number of inner points for the Euclidean TSP. The first part of the analysis carried out in [SN12] analyzes the expected time until the classical (1+1) EA using inversion as the mutation operators has computed a tour that is intersection free. The analysis depends on the progress that can be made by inversion operations removing an intersection and this progress depends on the angle $\epsilon > 0$ between any three points in the given set of cities. It is assumed that any three points are not collinear, i.e. do not lie on a single straight line. For an intersection free tour it is known that the points on the convex hull appear in the permutation in the same order as they appear on the convex hull. The remaining part analyzes the expected time until an optimal solution has been produced from an intersection free tour. This part of the analysis uses that at most $2k$ inversion operations are sufficient to produce from an intersection free tour an optimal tour. This implies that the (1+1) EA obtains an optimal solution in expected time $O(n^3 m^5 + n^{4k} (2k - 1)!)$ when the points are placed on an $m \times m$ grid and no set of three points is collinear. Here the parameter m for the grid directly determines the smallest angle that any set of three non collinear points can have. Note that the runtime bound does not meet the requirement of a fixed parameter evolutionary algorithm.

Afterwards, the ability of evolutionary algorithms to fill in the inner points correctly given that the points on the convex hull are in correct order has been examined. Ant colony optimization [NSN13b] and evolutionary algorithms [NSN13a, SNN14] have been investigated in this context. For ant colony optimization, the crucial aspect to obtain a runtime of $O(n^k)$ is to construct solutions following the order on the convex hull. For evolutionary algorithms, a population-based algorithm building on a previous approach of Theile [The09] and allowing to build an optimal tour following dynamic programming leads to a fixed parameter evolutionary algorithm with respect to the number of inner points. Furthermore, it is shown in [SNN14] that a simple $(\mu + \lambda)$ -EA searching for a permutation of the inner points and con-

necting them to the outer points using the dynamic programming approach given in [DHOW06] leads to a fixed parameter evolutionary algorithm.

2.2.3 Further Results for Combinatorial Optimization Problems

The early studies of Storch [Sto06] for the maximum clique problem in planar graphs investigated the runtime of RLS and the (1+1) EA with respect to the size of the maximum clique. The fitness (to be maximized) of a search point $x \in \{0, 1\}^n$, representing a selection of nodes, is given by the number of selected nodes if x represents a clique and $-\infty$ otherwise. The algorithms investigated start with the initial solution $x = 0^n$ which is a feasible solution. For standard bit mutations an expected optimization time of $\Theta(n^6)$ has been shown for (1+1) EA. However, it should be noted that the size of a maximum clique in a planar graph is at most 4 as the complete graph on 5 vertices is not planar. Improved results have been shown in [Sto06] for restart strategies used in RLS and for variants of the $(\mu + 1)$ -EA always deleting an individual with the worst fitness from the population.

The use of problem-specific mutation operators in the (1+1) EA for the maximum leaf spanning tree problem has been investigated in [KLNO10b]. In this work, it has been pointed out that standard bit flip mutations do not lead to fixed parameter evolutionary algorithms, where the parameter is the value of an optimal solution. Edge exchanges that include an edge currently not present in a spanning tree and that remove an edge from the resulting cycle are frequently used for spanning tree problems as they again lead to spanning trees. Using edge exchanges for mutation where the number of edge exchanges is chosen according to a Poisson distribution with expected value 1, it has been shown in [KLNO10b] that the resulting (1+1) EA is a fixed-parameter evolutionary algorithm when taking the value of an optimal solution OPT as the parameter.

2.3 Future work

There is lot of potential for future studies using fine-grained analyses to provide new and important insights. Fixed budget and fixed target analysis has mainly been applied to classical benchmark problems but the number of results for classical combinatorial optimization problems is very limited. Obtaining fixed budget results for such combinatorial optimization problems which do not follow directly from the corresponding worst case analyses would be very valuable and show the quality obtainable within a given time limit. The same holds for fixed target analysis where such studies would reveal

important insights for combinatorial optimization problems in terms of the time need to reach a desired quality.

The area of parameterized analysis of randomized search heuristics has been explored for some of the classical combinatorial optimization problems and such studies show which types of kernelization can be obtained in expected polynomial time. There are a wide range of kernelization techniques available in the classical parameterized complexity literature and it is unclear which types of kernels can be efficiently computed by evolutionary algorithms or similar search heuristics. Once a kernel is obtained, current studies use mutation operators that sample the reduced search space (almost) uniformly at random. While this leads to fixed parameter evolutionary algorithms, such operators are usually not that effective in practice. The question arises whether there are other useful techniques such as parameter adaptation that can be embedded into such mutation operators and ensure fixed parameter evolutionary algorithms and high performance in practice.

3 Setting the Parameters of Evolutionary Algorithms

The parameters of an evolutionary algorithm allow a user to adjust the EA to the problem to be solved and thus to optimize its performance. This is a great feature of EAs, but, at the same time, a difficult challenge [LLM07]. Missing good parameter values often results in horrible performance. Unfortunately, there is not much general advice on how to set the parameters.

In this section, we show how theoretical works have helped to understand how the parameters of EAs influence their performance. Recently, the theory of EAs has also made big progress in understanding and even designing automated ways to find good parameter values.

3.1 Finding Optimal Static Parameter Values (Parameter Tuning)

By *parameter tuning* we understand the problem (or process) of finding suitable parameter values and then running the EA with these parameters. The parameter values are not changed during the run, so we speak also of *static parameter values*. For reasons of space, we cannot discuss the whole literature on theoretical results that help tuning EA parameters, and therefore pick the mutation rate in elitist algorithms using standard bit mutation as the most prominent example. Other parameters that have attracted the-

Algorithm 3: The $(1 + \lambda)$ EA for maximization problems.

```
1 Choose  $x \in \{0, 1\}^n$  uniformly at random;
2 repeat forever
3   repeat  $i = 1$  to  $\lambda$ 
4     Create  $y_i$  by flipping each bit of  $x$  independently with
       probability  $1/n$ ;
5   Choose  $i \in [1..\lambda]$  such that  $y_i$  has maximal fitness among
        $y_1, \dots, y_\lambda$ , breaking ties uniformly at random;
6   if  $f(y) \geq f(x)$  then
7      $x := y$ 
```

oretical research include the parent and offspring population size (see, e.g., [JJW05, Wit06, RS14, DK15, ADFH18]) and the selection pressure (see, e.g., [JS07, Leh10, Leh11, LY12, ADY19]). For a discussion on how to set the parameters of estimation-of-distribution algorithms, we refer to Section 6.2.3.

We will focus our discussion in the following on the mutation rate, both because it is an important parameter and it is a parameter where a deep theoretical understanding exist. It is intuitively clear that we are faced with a trade-off situation when setting the mutation rate. A small mutation rate leads to slow a progress because the radius of exploration is small. A high mutation rate is detrimental because the random choice of the bits to be flipped on average increases the distance from the target solution, and this effect is linear in the mutation rate.

An early established [Müh92, Bäck93] and generally accepted [Bäck96, BFM97] recommendation is to use the mutation rate $p = \frac{1}{n}$ in standard bit mutation, that is, we generate an offspring by flipping each bit independently with probability $\frac{1}{n}$. With this choice, the expected distance between parent and offspring is one, so we inherit principles from local search. Different from local search, this mutation operator can leave local optima by flipping more than one bit.

A large number of mathematical runtime analyses shows that $p = \frac{1}{n}$ often is optimal and thus complements the experimental support for this recommendation (see, e.g., [Och02] and the references therein). For the performance of the $(1 + 1)$ EA on ONEMAX, a mix of rigorous and heuristic arguments already in [Müh92] and then fully rigorously in [GKS99] shows that $p = \frac{1}{n}$ is asymptotically optimal. For the LEADINGONES benchmark, a rate of $p \approx \frac{1.59}{n}$ was proven to be optimal in [BDN10]. The ONEMAX result was greatly extended in [Wit13] with a proof that $p = \frac{1}{n}$ is the asymptotically

optimal mutation rate for each pseudo-Boolean linear function with non-zero coefficients. In [GW17] it was proven that $p = \frac{1}{n}$ is the asymptotically optimal mutation rate for the $(1 + \lambda)$ EA when the offspring population size λ is not too large. The optimality of $p = \frac{1}{n}$ was also shown for the optimization of long-path functions [Sud13]. For monotone functions, the situation is not fully understood, but again mutation rates around $p = \frac{1}{n}$ appear to be a good choice. For the runtime of the $(1 + 1)$ EA on strictly monotonically increasing functions, a $\Theta(n \log n)$ runtime can easily be shown when the mutation rate is $\frac{c}{n}$ for a constant $0 < c < 1$. That $\frac{c}{n}$ mutation rates for larger c can lead to exponential runtimes was first shown in [DJS⁺13], the best known value for the constant c is 2.13... [LS18]. In the range around $p = \frac{1}{n}$, for a long time only a runtime guarantee of $O(n^{3/2})$ was known for p being exactly $\frac{1}{n}$ [Jan07]. A significant progress on this long-standing problem was only made very recently – in [LMS19] an entropy compression argument was used to show that an $O(n \log^2 n)$ runtime guarantee holds for all mutation rates $p = \frac{c}{n}$, where $c \leq c_0$ for some constant $c_0 > 1$.

We note that the above results give some indication that $p = \frac{1}{n}$ is a good first choice for the mutation rate, but by no means they prove that it always is. Indeed, already in [JW00] an example was constructed such that the $(1 + 1)$ EA with any mutation rate that is not $\Theta(\frac{\log n}{n})$ needs super-polynomial time with high probability to optimize this problem. In [Prü04], the optimal mutation rates for the $(1 + 1)$ EA optimizing hurdle functions with hurdle widths 2 and 3 were shown to be $\frac{2}{n}$ and $\frac{3}{n}$. This result could have led to the following findings, but apparently its broader implications on mutation rates (in a paper primarily discussing crossover) were not detected. So it was only in [DLMN17] that the optimal mutation rate of the $(1 + 1)$ EA on jump functions was shown to be roughly $\frac{k}{n}$, where k is the size of the fitness gap of the jump function. Also, it was shown that a small deviation from the optimal rate, say by a factor of $(1 \pm \varepsilon)$, $\varepsilon > 0$ a constant, leads to a significant increase of the runtime by a factor exponential in k .

This result shows that the optimal mutation rate depends strongly on the input instance, that there is no rate that is universally good for all jump functions, and that the price for missing the right rate is significant. This lead the authors of [DLMN17] suggest to use a random mutation rate, chosen independently for each mutation from a power-law distribution (we remark that random mutation rates were studied earlier to cope with unknown solution lengths [DDK19] and higher-arity representations [DDK18]). This heavy-tailed mutation operator shares with the classic mutation operator the property that a single bit (and more generally, any constant number of bits) is flipped with constant probability. When the power-law exponent is above two, then it also shares the property that an expected constant num-

ber of bits is flipped. Different from the classic recommendation, however, higher numbers of bits are flipped with larger probabilities. This essentially parameterless operator was shown to give on any jump function a performance of the $(1 + 1)$ EA that differs from the one with instance-optimal mutation rate by only a small factor polynomial in k . Heavy-tailed mutation operators proved to be successful in several other discrete optimization problems [FQW18, FGQW18b, FGQW18a, WQT18, ABD20a, ABD20b, AD20]. From a broader perspective, this line of work is an example showing that theoretical work not only can help understanding evolutionary algorithms, but it can also propose new operators and algorithms.

3.2 Dynamic Parameter Settings (Parameter Control)

Instead of trying to find a good parameter setting before starting the EA and sticking to this choice throughout the run of the EA, one could also think of optimizing the parameters during the run of the algorithm. This sophisticated-looking idea is called *parameter control* and turns out to be less frightening than it appears at first.

Indeed, the decision space (and thus also the opportunity to take an unsuitable decision) is much larger now – in principle, we could choose different parameter values in each iteration – but there are several powerful ways to overcome this difficulty. The advantage of parameter control is that we can react on the performance observed so far. This has two particularly positive consequences: (i) The need for finding good parameter values before the start of the algorithm, based on a maybe only vague understanding of the problem to be solved, is reduced since a suboptimal initial choice can be corrected. (ii) In the common situation that different parameter settings are optimal during different stages of the optimization process, we have the chance to use the optimal parameters for each stage (whereas a static choice would need to find a suitable trade-off).

It is clear that the large space of different parameter settings for each iteration renders it unlikely to find the absolutely best dynamic choice of the parameters. However, it turns out that often very simple success-based or learning-based approaches lead to a very good performance, and often one that is better than the best static parameter setting. This is confirmed in many practical applications, see, e.g., [KHE15], but also in now a decent number of theoretical works.

The theoretical superiority of dynamic parameter settings over static ones was already demonstrated in [DJW00] (see also [JW06] for an extension of this work), albeit for a simple algorithm with a simple *time-dependent parameter choice* optimizing an artificial problem. Nevertheless, this result has rig-

ously proven that, in principle, dynamic parameter choices can efficiently solve problems where classic static choices would badly fail. Interestingly, the idea of time-dependent mutation rates was recently used again [RW20] to help EAs leaving local optima.

It took ten years until dynamic parameter choices could be shown superior also for classic benchmark problems. The first such work [BDN10] (see also [Doe19a, Section 2.3] for an extension) showed that a constant-factor runtime gain can be obtained from a *fitness-dependent* choice of the mutation rate when optimizing the classic LEADINGONES benchmark via the $(1 + 1)$ EA. Again it took some time until in [BLS14], a super-constant runtime gain (of order $O(\log \log \lambda)$) from a dynamic parameter setting was shown for the $(1 + \lambda)$ EA optimizing ONEMAX. Other fitness-dependent parameter choices were discussed in [DDE15, DDY20]. A main problem with fitness-dependent parameter settings (or more generally speaking, parameter choices that depend on the current state of the algorithm) is that it needs a very good understanding of the problem to define a suitable functional dependence of the parameter value on the algorithm state. For the two examples from [BLS14, DDE15], it appears unlikely that without a mathematical analysis someone would have found the optimal functional dependence. Finding sub-optimal state-dependent parameter values that beat the best static values appears more realistic, but this remains a challenging task requiring a lot of expert knowledge.

Fortunately, there are dynamic parameter settings that need much less expertise. Generally speaking, these observe how the algorithm performs with the current parameter values (and sometimes also the values used in a longer history) and based on this try to adjust the parameter values to more profitable values. The easiest of these on-the-fly parameter choices are *success-based multiplicative parameter updates*. Assume we suspect that increasing a specific parameter assignment also increases the chance to find an improvement, but at the price of increasing the computational cost of one iteration. Then increasing the current parameter value after each iteration without improvement and decreasing it after each iteration with improvement is a simple way to try to move the parameter value into a profitable region. Exactly this was suggested for the offspring population size λ of the $(1 + \lambda)$ EA in [JJW05] and was rigorously analyzed in [LS11], where an asymptotically optimal speed-up of the parallel runtime (number of iterations, ignoring the different costs of the iterations) was shown. The same basic idea was shown to give a (small) asymptotic improvement of the total runtime (number of fitness evaluations) for the $(1 + (\lambda, \lambda))$ GA optimizing ONEMAX [DD18] and certain random SAT instances [BD17].

The usual way to change the parameter value is multiplying or dividing by suitable constant factors. In [LS11], the factor 2 was used, and it is clear that any other constant factor would have given the same asymptotic runtime. In general, as observed in [DD18], smaller update factors can be the safer choice, and also the relation of the factors used in case of success and no success can be important. In [DDL19], a detailed analysis how the choice of these hyperparameters influences the runtime of the $(1 + 1)$ EA with dynamic mutation rate on the LEADINGONES function was conducted. Other theoretical works on multiplicative parameter updates include [DDK18] for multi-valued decision variables, [MS15] for migration intervals of island models, and [DLOW18] for the learning period of a hyperheuristic. We note that the results just described are the first examples of success-based parameter updates in discrete evolutionary optimization. In continuous optimization, a multiplicative update of the step size known as one-fifth rule was already proposed in [Rec73].

Multiplicative update rules work best if there is a simple monotonic influence of the parameter on the success, e.g., as seen for the offspring population size of the $(1 + \lambda)$ EA. Since such a simple relation is harder to find for the mutation rate in the $(1 + \lambda)$ EA, a different success-based scheme was developed in [DGWY19]. Here half of the offspring are generated with twice the current rate, the other half with half the current rate. The mutation rate is then updated to the rate the best offspring was generated with (however, only with probability a half, with the other one-half probability the new rate is chosen randomly from the two alternatives). This mechanism was shown to let the $(1 + \lambda)$ EA optimize ONEMAX in asymptotically the same time as with the optimal fitness-dependent mutation rate developed in [BLS14].

A second way to go beyond multiplicative updates, and to additionally take more stable decisions, was proposed in [DDY16]. Here for a small number of possible values of a parameter, a time-discounted estimate of the effectiveness of this parameter value was computed. In each iteration, with large probability, the best-performing value was used (exploitation) and with small probability a random one of the other values was used. With the right choice of the hyperparameters, this mechanism was shown to arbitrarily well approach the optimal mutation strengths of the $(1 + 1)$ EA optimizing ONEMAX that were computed in [DDY20].

The most generic way to let an EA optimize its parameters itself is *self-adaptation*, which means that the parameters are made part of the encoding of the solution candidates and thus become subject to variation and selection. The first to use this idea from evolution strategies in discrete evolutionary optimization was Bäck [Bäc92]. Taking the mutation rate as example, one appends an encoding of the mutation rate to the representation of the solution

candidates. When mutating such an extended individual, one first mutates the mutation rate encoded in the extended individual and then, with the new rate, the remainder of the individual. The hope is that the suitability of a rate is visible from a higher fitness of the resulting individuals, and that the selection mechanisms of the EA bring these individuals (and thus the good mutation rate) forward in the population. While this way of adjusting parameters is clearly more natural for an EA than parameter adjustment mechanisms outside the evolutionary process, only three rigorous results supporting the usefulness of self-adaptation in discrete evolutionary computation have been published. In a first proof-of-concept work [DL16], an example is constructed that shows that self-adaptation can be useful. In this example, only two different mutation rates are available and it is assumed that the whole initial population starts in a particular search point. In [DWY21], the $(1, \lambda)$ EA with self-adapting mutation rate is analyzed. With the hyperparameters suitably chosen, it can evolve sufficiently good mutation rates to obtain asymptotically the same performance on ONEMAX that was previously obtained with the optimal fitness-dependent setting [BLS14] and the two-population self-adjustment [DGWY19]. A self-adaptive choice of the mutation rate for the (μ, λ) EA optimizing LEADINGONES was proposed and analyzed in [CL20]. Here, the mutation rate is multiplied by a constant $A > 1$ with some probability p , otherwise it is multiplied by a factor $b \in (0, 1)$. Such a multiplicative update was used in a success-based manner in [DDL19]. When the algorithm parameters A, b, p, μ, λ satisfy certain natural range constraints, then this self-adjusting algorithm optimizes the capped LeadingOnes function $\text{LEADINGONES}_k = \min\{\text{LEADINGONES}, k\}$ (defined on bit strings of length n) in time $O(k^2)$ for all $k \in [\Omega(\log^2 n)..n]$.

3.3 The Future of Parameter Research

The existing results show that we are now able to analyze a variety of static and dynamic parameter choices with a precision high enough to clearly distinguish good from bad choices. Some of these works not only analyzed existing algorithms or parameter adjusting mechanisms, but also suggested new approaches. Clearly, as true for all theoretical works, the algorithms and problems that were regarded are much simpler than those occurring in a practical application of EAs. To what extent the recommendations obtained from these simple settings generalize to more realistic ones is a crucial question which can only be answered in a collaboration between theoretical and applied researchers.

From the theory perspective, the following questions appear timely and interesting.

- Interaction of parameters: So far, the vast majority of runtime analyses varies at most one parameter of the algorithm. Experience from practice shows that the interaction of several parameters is even harder to understand. So more runtime analyses discussing several parameters at once are clearly needed. Also, to the best of our knowledge, there is currently no theoretical work regarding two or more independent heavy-tailed parameters or self-adjusting or self-adaptive settings of two or more parameters.
- Self-adaptation: The most natural way to let an algorithm optimize its parameters is self-adaptation, where the parameters are integrated into the evolutionary cycle. So far, only very little theoretical advice exists how to successfully control parameters via self-adaptation. Here clearly more work is required.
- Connections with machine learning: The area of machine learning has made tremendous progress in the last decades. Given that EAs are iterative algorithms in which often the state of the system changes only little in each iteration, one could envisage that dynamic parameter choices can profit from ideas and concepts borrowed from machine learning. While some ideas used in EAs can be related to similar ideas in machine learning, it seems to us that the full power of this connection has not yet been exploited.

4 Analysis of Evolutionary Algorithms in Dynamic and Stochastic Environments

Dynamic and stochastic environments play a key role in real-world applications as information is often uncertain and circumstances change over time. Evolutionary algorithms have the ability to deal with changing circumstances and perform well in noisy environments which makes them well suited for dealing with dynamic and stochastic problems. A key contribution to the design of evolutionary for dynamic problems is the PhD Thesis of Branke [Bra00] which provides a wide range of techniques and investigations of evolutionary algorithms in dynamic environments. Many dynamic problems have been tackled by evolutionary algorithms and other search heuristics. This includes important problems with frequent dynamic components such as vehicle routing [MGRD05, HO07]. Comprehensive presentations on the different techniques and problems investigated can be found in [NYB12, MLY17].

The area of runtime analysis has initially focused on simple toy problem in dynamic and stochastic settings. Again the function OneMax has played a crucial role to get initial insights. An important aspect in the context of dynamic optimization is how often and how drastic a function changes over time. We will describe important results for settings where the function or the constraints of a given problem change dynamically. Furthermore, we will summarize results where the fitness evaluation is impacted by noise and point out different results according to different noise models studied in the literature. Additional investigations regarding dynamic and stochastic constraints in the context of submodular optimization are summarized in Section 5.

4.1 Dynamic Benchmark Functions

The runtime analysis for dynamically changing functions in discrete search spaces has been started by Droste [Dro02, Dro03]. He investigated a dynamic variant of the classical OneMax problem on binary strings. In the first dynamic setting, one randomly chosen bit is flipped in each iteration with probably p . Droste [Dro02] showed that the expected optimization time of the (1+1) EA is polynomial iff $p = O(\log(n)/n)$. In the case where each bit is flipped in each iteration with a given probability p investigated in [Dro03], the runtime becomes super-polynomial if $p = \omega(\log(n)/n^2)$ and is polynomial if $p = O(\log(n)/n^2)$. These investigations were revisited ten years later using drift analysis and generalized to the case where each element is not binary but can take on r different values [KLW15a]. Investigations on the magnitude and frequency of change for artificial benchmark functions have been carried out in [RLY09]. This study provides examples where the expected runtime changes drastically, i.e. from polynomial to exponential, if the frequency of change changes from low to high (or vice versa).

A comparison on the ability of simple evolutionary algorithms and ant colony optimization approaches for dealing with dynamic fitness functions has been carried out in [KM12a, LW16]. These studies show that ant colony optimization can beat evolutionary algorithms due to their ability of adjusting slowly to changes in the fitness functions. Investigations of parallel evolutionary algorithms using island models carried out in [LW18] for the MAZE function, introduced in [KM12a], show that infrequent migration of individuals is necessary for dense models where as infrequent migration becomes less necessary when working with sparse topologies in the island model. Furthermore, the usefulness of using non-elitist populations to track optima of dynamically changing problems has been pointed out in [DJI17].

4.2 Dynamic Combinatorial Optimization Problems

There are also some results on classical combinatorial optimization problems in the dynamic setting. Lissovoi and Witt [LW15] have investigated ACO algorithms and shown how the number of ants can impact different types of changes that can be tracked over time. They also give an example of dynamic oscillations that can not be tracked with a polynomial number of ants. Dynamic makespan scheduling for two machines has been investigated by Neumann and Witt [NW15]. They have studied dynamic settings where solutions of small discrepancy of the two machines have to be recomputed. The results show that a worst case discrepancy of U , where U is an upper bound on the maximal job length, can be maintained. Furthermore, better upper bounds on the runtime and lower discrepancies are shown for the case where the processing times of the jobs change randomly.

Dynamic variants of the minimum vertex cover problem have been considered in [PGN15, PRN20]. Following the edge-based encoding for the minimum vertex cover problem introduced in [JOZ13], the problem formulation makes use of the dual formulation of the problem in order to represent solutions. In [PGN15], the expected time to recompute 2-approximation when edges are added or removed has been studied and improved results have been presented in [PRN20].

Dynamic settings of the classical graph coloring problem have been investigated in [BNPS21]. Here, in particular, bipartite graphs have been studied and the necessity of complex mutation operators has been revealed even if there are only slight dynamic changes to the graph structure. These investigations have recently been extended in [BNPS20] and it has been shown that a dynamic setting where edges are presented to the algorithm in an iterative way can provably lead to better optimization times than presenting the algorithm with the whole input graph at once.

4.3 Noisy Problems

Studies in the area of runtime analysis of evolutionary computing techniques for discrete search spaces involving noisy objective functions have again been started by Droste [Dro04] who analyzed the (1+1) EA on a noisy version of OneMax. He studied a prior noise model. In this case, some bits of a solution x are flipped prior to the fitness evaluation. The studies considered flipping each bit with probability p prior to fitness evaluation and Droste showed that the (1+1) EA can still obtain the optimal solution for OneMax in expected polynomial time if $p = O(n/\log n)$ whereas the expected optimization time becomes super-polynomial if $p = \omega(n/\log n)$. In general, investigations can

be separated into ones investigating prior noise as described above and posterior noise. In the case of posterior noise, the fitness of a solution x is evaluated using the given fitness function f but noise is added afterwards to the fitness value $f(x)$. Gießen and Kötzing [GK16] build on this initial study by Droste and extended the study to population-based evolutionary algorithms and considered prior and posterior noise. Results for prior bit-wise noise for the classical benchmark functions OneMax and LeadingOnes have been obtained in [BQT18, QBJT19]. Additional and improved results including an example where noise helps have been provided by Sudholt [Sud18] and estimation of distribution algorithms (see Section 6) have been studied for OneMax in [FKKS17]. A method that can be used for the analysis of dynamic and noisy fitness functions has been developed in [DNDD⁺18]. Furthermore, the trade-off between reducing noise through resampling and the computational cost of resampling has been investigated by Friedrich et al. [FKQS17]. The authors investigated classical evolutionary algorithms, estimation of distribution algorithms and ant colony optimization for OneMax and the case of additive posterior Gaussian noise. The use of non-elitist population based evolutionary algorithms for noisy versions of OneMax and LeadingOnes has been analyzed in [DL15a].

In terms of classical combinatorial optimization problems, Sudholt and Tyssen [ST12] investigated the stochastic single-destination shortest path problem and have pointed out settings where ant colony optimization is able to solve this problem in expected polynomial time or obtain a good approximation in expected polynomial time. These investigations have been extended in [DHK12] where a slight modification of the previously considered ant colony optimization approach has been analyzed and shown that this modified approach finds shortest path lengths efficiently although it does not necessarily converge. Furthermore, stochastic constraints in form of chance constraints which require that stochastic constraints can only be violated with a small probability have been investigated recently (see Sections 4.4 and 5.3).

4.4 Combinatorial Optimization Problems with Dynamic and Stochastic Constraints

Dynamic constraints reflect the change in resources to solve a given problem. This is often a crucial aspect in many planning problems where resources such as trucks and trains might become unavailable due to failures or become available (again) after maintenance. Considering dynamic constraints, the objective function to be optimized is often assumed to be fixed and only

changes to the constraints are considered. The simplest example is the maximization of a linear function subject to a uniform constraint which limits the number of elements to be at most B . The first runtime analysis in this area considered the case where the bound B changes to B^* and the question is how long an evolutionary algorithm needs to recompute from an optimal solution for a given bound B an optimal solution for the updated bound B^* . The (1+1) EA and simple evolutionary multi-objective algorithms have been studied in [SSF⁺19]. The multi-objective formulations enable to efficiently deal with constraint bound changes and this leads to significantly better experimental results on a wide range of knapsack instances as shown in [RNN18, RNN20] if the magnitude and frequency of change is not too large.

Chance constraints model constraints that are impacted by some noise of the given components and the goal is to optimize a given function f under the condition that the constraint is violated with probability at most α , where α is usually a small value, e.g. $\alpha = 0.001$. Evolutionary algorithms for the chance-constrained knapsack problem have been studied from an experimental perspective by Yue et al. [XHA⁺19, XNN20, XNN21]. Furthermore, Assimi et al. [AHX⁺20] investigated evolutionary multi-objective evolutionary algorithms for the dynamic chance-constrained knapsack problem where the constraint bound for the knapsack dynamically changes over time through experimental studies. A first runtime analysis for problems with chance constraints has been carried out by Neumann and Sutton [NS19] for special instances of the knapsack problem. It shows that even very simple linear functions with a simple linear stochastic constraint can lead to local optima with large inferior neighbourhoods that may make it hard for the (1+1) EA to produce an optimal solution. These investigations have been extended by Yue et al. [XNNS21] to the case of uniform stochastic weights where there are groups of elements that are correlated. For the setting in which every correlated group has the same profit profile, polynomial upper bounds on the expected optimization time of RLS and the (1+1) EA have been given.

Important results on evolutionary algorithms for the optimization of submodular functions under dynamic and stochastic constraints have been obtained recently and are summarized in Section 5.3. Furthermore, a more comprehensive and technical survey on the theory of evolutionary computing in dynamic and stochastic environments can be found in [RPN18].

4.5 Future work

While there has been quite a large progress in recent years in this area of research, there are still a wide range of open problems and questions for future

research. Obviously, the understanding is still limited to basic benchmark functions and some well structured combinatorial problems and to further push the boundary of understanding is an important task. Combining experimental analysis with mathematical investigations seems to be an important way to go as understanding evolutionary algorithms from a mathematical perspective is already challenging in static environments and even more true in the case of dynamic and stochastic problems. The design of evolutionary algorithms for dynamic and stochastic problems can (and has already been) guided by theoretical results. One example are dynamic and stochastic variants of the knapsack problem as investigated in [RNN18, XHA⁺19] where the multi-objective formulations are based on previous theoretical results. Such interactions of experimental and mathematical research seems to be able to bridge the gap and allow to transfer theoretical results into the design of higher performing evolutionary algorithms.

5 Submodular optimization

Submodular functions play a keyrole in the area of optimization. Many real world problems can be stated in terms of a submodular function as they face a diminishing return when adding additional components to a solution (see the survey of Krause and Golovin [KG14]). A wide range of greedy algorithms and local search approaches are available to solve different types of submodular problems with different types of constraints [NWF78, KMN99, LMNS09]. In the area of classical algorithms, these algorithms are generally designed and analyzed in terms of approximation guarantees that can be obtained in polynomial time [CC84, Von10]. The wider area of artificial intelligence also follows this approach but tests the performance of the developed approaches on important real-world problems matching the studied characteristics [ZV16, MKF⁺19]. The previously mentioned studies can roughly be divided by the type of objective functions that are considered and the set of constraints. Furthermore, recent more general investigations consider the case where a function is not necessarily submodular and obtain results dependent on how close it is to being submodular measured in terms of the so-called "submodular ratio". Other important characteristics of the problems considered are noisy or dynamic problems [Mon20, HS17] and problems where additional information is revealed over time [GK11, EKM19].

Evolutionary algorithms have recently been analyzed and applied to different submodular optimization problems. We will summarize some of the key results in this section. For more details on some of the results we refer to the original articles or the recent book by Zhou et al. [ZYQ19]. This book

gives a very comprehensive presentation on submodular optimization by evolutionary algorithms and discusses a wide range of submodular problems in the areas of optimization and machine learning.

We consider the following setting. Given a set $V = \{v_1, \dots, v_n\}$ of elements, the goal is to maximize a function $f: 2^V \rightarrow \mathbb{R}^+$ that maximizes f subject to a given set of constraints. Submodular functions are usually considered in terms of marginal value when adding a new element. We denote by $F_i(A) = f(A \cup \{i\}) - f(A)$ the marginal value of i with respect to A . A function f is submodular iff $F_i(A) \geq F_i(B)$ for all $A \subseteq B \subseteq X$ and $i \in X \setminus B$. Furthermore, a function f is called monotone iff $f(A) \leq f(B)$ for $A \subseteq B$.

The first investigations in terms of runtime behaviour of evolutionary algorithms for submodular functions, we are aware of, have been carried out by Rudolph [Rud97] in the 1990s. More than 15 years later this research area has been re-started by Friedrich and Neumann [FN14] and has since then gained significant attention. The research in the context of static optimization can be grouped with respect to the type of objective functions and the type of constraints that are considered. In terms of objective functions, it is usually differentiated between monotone and non-monotone submodular functions. Furthermore, the submodularity ratio plays a crucial role when broadening the class of functions to functions that are not submodular. This ratio measures how close a function is to being submodular. The submodularity ratio α_f of a given function f is defined as

$$\alpha_f = \min_{X \subseteq Y, v \notin Y} \frac{f(X \cup v) - f(X)}{f(Y \cup v) - f(Y)}.$$

Note that if f is submodular then $\alpha_f = 1$ holds. The other important component in these investigations are the type of constraints that are considered. Constraints are usually of the type $c(X) \leq B$, where $c: 2^V \rightarrow \mathbb{R}^+$ assigns a non negative cost to each set of elements $X \subseteq V$ and B is a given constraint bound. This type of constraints includes the case of a simple uniform constraints where $c(X) = |X|$ holds and limits the number of elements that can be included in a feasible solution by B . The maximization of a monotone submodular function under a uniform constraint is already NP-hard and can be approximated within a factor of $(1 - 1/e)$ by a simple greedy algorithm [NWF78]. More complex constraints involve partition or matroid constraints which are given in form of linear functions. Complex constraints that have been considered include cost values that can only be approximated, i.e. involving NP-hard routing problems. Then the approximation obtained for the submodular function depends on the type of objective function as well as the ability to calculate the cost of the considered constraint. In the

following, we summarize some of the main results in this currently very active research area.

5.1 Monotone Submodular Functions

Optimal solutions for monotone submodular functions $f(X)$ with a cost constraint $c(X) \leq B$ can often be approximated well by simple greedy algorithms (see [KG14] for a comprehensive survey).

Such greedy algorithms start with the empty set and add in each iterations an element with the largest marginal gain

$$(f(X \cup \{x\}) - f(X)) / (c(X \cup \{x\}) - c(X))$$

that does not violate the constraint. The algorithm stops if no element can be added without violating the constraint bound.

Variants of GSEMO (see Algorithm 2) have been widely studied in the context of optimizing submodular functions. The initial analysis carried out in [FN14] considered the maximization of monotone submodular functions with different types of constraints. After this GSEMO has been widely studied in the context of submodular optimization under the umbrella of Pareto Optimization which formulates a given constraint optimization problem as a multi-objective problem by establishing an additional objective based on the considered constraint. Such approaches have been widely used already before this in the context of runtime analysis of GSEMO. Solving single-objective problems by multi-objective formulations is a well-known concept in the evolutionary computation literature and has been studied from a practical and theoretical perspective since mid of the 2000s [Jen04, NW06, BFH⁺09].

For the simplest case of a uniform constraint where $c(X) = |X|$, GSEMO can select in each step an element with the largest marginal again with respect to f . Friedrich and Neumann [FN14] have shown that GSEMO produces a $(1 - 1/e)$ -approximation for monotone submodular functions with a uniform constraint in expected time $O(n^2(\log n + B))$ where $B \leq n$. For monotone submodular functions with k matroid constraints, local search and simple single-objective evolutionary algorithms such as the classical (1+1) EA are able to obtain good approximation results. It has been shown by Lee et al. [LMNS09] that local search introducing at most $2p$ new elements and removing at most $2kp$ elements is able to obtain a $(1/(k + 1/p + \epsilon))$ -approximation in polynomial time if $k \geq 1$ and $p \geq 2$ are constants. This result has later on been used to show that the (1+1) EA is able to obtain the same approximation guarantee in expected time $O(\frac{1}{\epsilon} \cdot n^{2p(k+1)+1} \cdot k \cdot \log n)$. The crucial part of the proof is a result from [LMNS09] which shows that

every solution x for which there is no y in the defined neighborhood with $f(y) \geq (1 + \frac{\epsilon}{n(k+1)}) \cdot f(x)$ is already a $(1/(k + 1/p + \epsilon))$ -approximation.

Further investigations lead to a wide range of results for GSEMO on various submodular problems with cost constraints. The algorithm GSEMO is often called POMC (Pareto Optimization algorithm for maximizing a monotone function with a monotone Cost constraint) (or similar) in such articles. The approaches are referred to as Pareto optimization and specify the class of problems to be tackled. However, usually the difference only lies in the formulation of the objective functions to formulate the constrained submodular problems as a multi-objective optimization problem.

An important result covering a wide range of monotone functions for a broad class of cost constraints has been obtained by Qian et al. [QSYT17]. They investigated monotone functions in terms of submodularity ratio and general cost functions including ones for which it is hard to obtain an optimal solution exactly. Their theoretical results make use of proof ideas used for an adaptive greedy algorithm and show that POMC is able to obtain the same approximation guarantee in expected pseudo-polynomial time. The expected runtime may be exponential with respect to the given input here if both the submodular function and the cost function can take on exponentially many values. In this case, the population size of GSEMO may become exponential during the run. More precisely, they have shown that POMC obtains a $(\alpha/2) \cdot (1 - e^{-\alpha})$ -approximation, α is the submodularity ratio, for a tightened cost constraint bound \hat{B} (instead of B) where \hat{B} depends on how well the given cost constraint can be approximated. Note that this setting includes problems where the cost of a solution may be hard to compute, i.e. for a selection of items it could be an approximation of a minimum Traveling Salesperson tour. The experimental results show that POMC clearly outperforms the adaptive greedy approach if the evolutionary algorithm is given a sufficient large number of fitness evaluations. Recently, an evolutionary multi-objective algorithm called EAMC has been introduced in [BFQY20] which obtains the same worst-case approximation ratio as POMC in expected polynomial time if the submodularity ratio of the given problem is known and used as a parameter in the algorithm. However, EAMC usually performs worse than POMC on important benchmark problems.

Subset selection has also been investigated in the context of sparse regression. Here the submodular ratio α_f of the underlying function to be optimized plays a crucial role for the approximation quality obtained. Again, a variant of GSEMO called POSS [QYZ15] achieves in expected polynomial time the same approximation quality as a greedy approach called forward regression [DK11], namely a solution X with $f(x) \geq (1 - e^{-\alpha}) \cdot OPT$. Furthermore, POSS outperforms forward regression and other simple heuristics in experi-

mental investigations in terms of solution quality when giving it a sufficient amount of time to improve solutions during the evolutionary optimization process.

5.2 Non-monotone Submodular Functions

For symmetric functions which are not necessarily monotone and have k matroid constraints, evolutionary algorithms and local search approaches can increase the function value by local operations to obtain a good approximation. Lee et al [LMNS09] have shown that if the value of a solution x can not be increased by a factor of at least $(1 + \epsilon/n^4)$ by changing at most $k + 1$ elements, then x is a $\frac{1}{(k+2)(1+\epsilon)}$ -approximation. The series of such local improvements requires that the algorithm obtains a solution x of value at least $f(x) \geq OPT/n$. Such a solution can be obtained from the empty set by adding the single element with the largest function value. Consequently local search algorithms building on such a solution and exchanging at most $k + 1$ elements obtain a solution with the stated approximation quality in polynomial time [LMNS09]. It has been shown that GSEMO obtains a $\frac{1}{(k+2)(1+\epsilon)}$ -approximation in expected time $O((1/\epsilon)n^{k+6} \log n)$. The proof analyzes the process until a solution x with $f(x) \geq OPT/n$ is obtained and the required number of local improvements until a solution of the stated approximation quality is obtained.

In their recent work, Qian et al. [QYT⁺19] give other major results which are broadening the setting of previous investigations. They considered an evolutionary multi-objective algorithm called GSEMO-C which differs from GSEMO by producing from the offspring x' a second offspring x'' which is the complement of the first offspring. The selection step of GSEMO is then applied to both x' and x'' . The authors first showed that for the case of non-monotone submodular functions without any constraint, GSEMO-C is able to obtain a $(1/3 - \epsilon/n)$ -approximation in expected time $O(\frac{n^4}{\epsilon} \log n)$. For ϵ -monotone submodular functions, $\epsilon \geq 0$, where $f(X \cup \{x\}) \geq f(X) - \epsilon$ holds for any $X \subseteq V$ and $x \notin X$, and a uniform constraint with bound B , they showed that GSEMO-C achieves a solution x with $f(x) \geq (1 - 1/e) \cdot (OPT - k\epsilon)$ in expected time $O(n^2(B + \log n))$ which generalizes the result given in [FN14] to a wider range of functions by taking their closeness to monotonicity into account. Similar approximation results also hold for GSEMO-C when considering ϵ -approximately submodular functions, i.e. for functions f for which a submodular function g exists such that for all $X \subseteq V$, $(1 - \epsilon)g(X) \leq f(X) \leq (1 + \epsilon)g(X)$ holds. The authors showed that suitable approximation can also be obtained for a wider range of functions

with a cardinality constraint in expected time $O(n^2(B + \log n))$. Specifically, they obtained results that depend on the submodularity ratio of the problem and investigated functions that are ϵ -approximately submodular.

Functions with bounded curvature under partition matroid constraints have been investigated in [FGN⁺19]. The results include the case of non-monotone submodular functions and approximation guarantees have been shown for the generalized greedy algorithm. These investigations have recently been extended by Do and Neumann [DN20a] to the evolutionary multi-objective algorithm GSEMO which is able to guarantee the same approximation quality as greedy but usually outperforms greedy in practice.

5.3 Submodular Functions with Dynamic and Stochastic Constraints

Recent studies extended the investigations for monotone objective and costs functions to problems with dynamic constraints as well as constraints involving stochastic components. Roostapour et al. [RNNF19] investigated the setting of general cost constraints where the constraint bound B changes over time. Generalizing the results of Qian et al. [QSYT17] which are summarized in Section 5.1, they have shown that the evolutionary multi-objective approach POMC computes an approximation for every budget b , $0 \leq b \leq B$. Furthermore, they have shown that if B is increased to B^* , then an approximation for every b , $0 \leq b \leq B^*$ is obtained in pseudo-polynomial time. In contrast to this, it has been pointed out in [RNNF19] that simple adaptations of the generalized greedy algorithm are not able to maintain good approximations when dynamic changes are carried out. Furthermore, POMC is able to learn the dynamic problems over time which gives it significant advantages over the greedy approaches as shown in comprehensive experimental investigations [RNNF18].

Recently, the investigations in the area of submodular optimization have also been extended to stochastic constraints. Chance constraints play an important role in stochastic settings. These model situations where components of a constraint are stochastic and the goal is to optimize a given submodular objective function such that the probability of violating a given constraint bound is at most α . Doerr et al. [DDN⁺20] investigated greedy algorithms for the optimization of monotone submodular functions for two settings. In the first setting, the stochastic weights are identically and independently uniformly distributed within a given interval $[a - \delta, a + \delta]$, $\delta \leq a$, where δ models the uncertainty of the items. In the second setting each element s has its own expected weight and is chosen independently of the others and uniformly at

random in $[a(s) - \delta, a(s) + \delta]$, $\delta \leq \min_{s \in V} a(s)$. The investigations have recently been extended by Neumann and Neumann [NN20] to GSEMO and it has been shown that this algorithm is able to obtain the same approximation guarantee as the greedy approach in expected polynomial time in the case of identically and independently uniformly distributed weights. For the second setting, the same approximation guarantee as the one obtained for the greedy approach is obtained in expected pseudo-polynomial time. Furthermore, experimental investigations carried out for the influence maximization problem in social networks and the maximum coverage problem show that GSEMO significantly outperforms the greedy approach. A comparison of GSEMO to a standard setup of NSGA-II reveals that GSEMO is also often outperforming NSGA-II for the investigated settings which suggests that the ability of GSEMO to construct solutions in a greedy fashion is also crucial for the success of the algorithm in practice.

5.4 Future work

The studies show that evolutionary algorithms often provide the same worst case performance guarantees as classical algorithms but perform much better in practice. Many studies focus on variants of GSEMO which may encounter an exponential population size. Recent investigations show that special type of archived based algorithms such as EAMC only need a polynomial population size and achieve the same approximation guarantees. However, these algorithms usually perform worse in practice. Overall, it would be important to get a better understanding of classical evolutionary multi-objective algorithms for the use of Pareto optimization approaches in the context of submodular functions. Furthermore, the multi-objective model is usually used to enable a greedy process and the question arises whether the whole set of possible trade-offs in the multi-objective formulation is required to achieve good results from a mathematical and experimental perspective. In terms of problem characteristics, it seems to be important to expand previous investigations, in particular, in the areas of dynamic, stochastic problems and further investigate the adaptation capabilities of evolutionary approaches for submodular problems in changing environments.

6 Theory of Estimation-of-Distribution Algorithms

Estimation-of-distribution algorithms (EDAs) are a more recent class of evolutionary algorithms (EAs). As a main difference to classic EAs, they do

not evolve a population (that is, a finite set of solution candidates), but a probabilistic model of a solution candidate (that is, a probability distribution over the search space). Whereas a traditional EA selects individuals from a parent population, creates from them offspring via mutation and crossover, evaluates the offspring, and based on this evaluation selects from parents and offspring the next parent population, the EDA samples individuals from the current probabilistic model, evaluates them, and based on this evaluation defines the next probabilistic model. When viewing a parent population of a classic EA as probabilistic model (uniformly distributed on the individuals of the population), one can interpret population-based EAs as particular EDAs, but it is clear that the probabilistic models of EDAs are much more expressive than models building on finite populations. The obvious hope is that this richer class of algorithms contains better optimizers. However, there is also the additional hope that the probabilistic model evolved by an EDA can give insights beyond the good solutions that can be sampled from it.

Most EDAs were defined in the 1990s, first in 1993 in an unpublished work [JBS93] by Ari Juels, Shumeet Baluja, and Alistair Sinclair (see [Lob07]) proposing the *equilibrium genetic algorithm* (similar ideas can already be found in [Ack87, Sys93]). While clearly containing the right ideas, this paper was never published and this algorithm is little known. Acknowledging the joint work with Ari Juels, Shumeet Baluja [Bal94] proposed a very similar algorithm called *population-based incremental learning (PBIL)*. As an important special case of it, Mühlenbein and Paass [MP96] two years later suggested the *univariate marginal distribution algorithm (UMDA)*. In 1999, Harik, Lobo, and Goldberg [HLG99] proposed the *compact genetic algorithm (cGA)*. These and many other EDAs found numerous successful applications in the following years, see, e.g., the surveys [HP11, LL02, PHL15].

First attempts to understand EDAs via theoretical means soon followed, starting – as often – with convergence results such as [HR97]. We note, however, that many of these very early results work with simplifying assumptions such as infinite population models and thus are not fully rigorous in the strict mathematical sense. In a series of works, Shapiro [Sha02, Sha05, Sha06] analyzed how the parameters of EDAs influence the effect of genetic drift. We discuss this central topic in more detail in Section 6.2.3.

The first rigorous runtime analysis for an EDA was presented by Droste at GECCO 2005 (journal version [Dro06]). Chen, Lehre, Tang, and Yao [CLTY09] exhibited an artificial example problem which is easily solved by the UMDA, but for which the $(1 + 1)$ EA with any $\Theta(\frac{1}{n})$ mutation rate needs exponential time to find the optimum. In [CTCY10], Chen, Tang, Chen, and Yao discussed the use of frequency boundaries to prevent prema-

ture convergence. After these early works, it took another five years without theoretical works on EDAs until this area gained significant momentum in 2015–2016 with works like [DLN19] (conference version at GECCO 2015), conducting a runtime analysis of the UMDA on ONEMAX and LEADING-ONES, [FKKS17] (conference version at ISAAC 2015) on the robustness of EDAs to noise, [SW19] (conference version at GECCO 2016) on how the update strength influences the runtime of the cGA, and [FKK16] pointing out that the main known EDAs are balanced, but not stable (that is, subject to genetic drift). These works generated a broad interest in theoretical analyses of EDAs, resulting in a large number of strong papers by a number of different authors. We refer to the recent survey [KW20a] for more details.

There is a natural connection between EDAs and ant colony optimization (ACO), since the pheromone system of an ACO algorithm leads to a probabilistic model of the search space as well. Two notable differences between classic EDAs and ACO algorithms are (i) that some ACO algorithms perform a pheromone update with the best-so-far solution, hence the state of the algorithm is not only described by the probabilistic model, and (ii) that many ACO algorithms, in particular those for graph-based problems, use a very particular way of constructing new solutions. For this reasons, we do not discuss ACO algorithms in this section and refer the reader to the literature, e.g., the early convergence analyses [Gut00, SD02, Gut02, Gut03, ST05], the first steps in runtime analysis for the 1-ANT algorithm and the MMAS with best-so-far improvement [NW09, DJ07, Gut08, GS08, NSW09, DNSW11], runtime analyses on graph-based problems [AF08, Zho09, NW10a, KNRW10, KLNO10a], runtime analyses for ACO algorithms with iteration-best improvement (which are very close to EDAs) [NSW10, SW19], works demonstrating the robustness of ACO in stochastic environments [KM12b, ST12, DHK12, FK13, LW16, FKKS16], or the survey [Gut11], which still is a very valuable resource.

6.1 The Compact Genetic Algorithm

We now describe the compact genetic algorithm (cGA) [HLG99], which will serve as a central example in this section. Other EDAs such as the UMDA or PBIL are substantially different, but appear to have similar strengths and challenges, so expecting similar results for these is a reasonable rule of thumb. However, we only concentrate on EDAs for discrete optimization problems here and we expect very different results in the continuous world. For reasons of simplicity, we only regard *pseudo-Boolean* problems, that is, the optimization of functions $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

The *compact genetic algorithm* (cGA) is a *univariate EDA*, that is, it treats the decision variables independently. For pseudo-Boolean problems, its probabilistic model is described by a *frequency vector* $p \in [0, 1]^n$. This frequency vector determines the following probability distribution on the search space $\{0, 1\}^n$. If $X = (X_1, \dots, X_n) \in \{0, 1\}^n$ is a search point sampled according to this distribution – we write $X \sim \text{Sample}(p)$ to indicate this – then we have $\Pr[X_i = 1] = p_i$ independently for all $i \in [1..n] := \{1, \dots, n\}$. In other words, the probability that X equals some fixed search point y is $\Pr[X = y] = \prod_{i:y_i=1} p_i \prod_{i:y_i=0} (1 - p_i)$.

In each iteration, the cGA samples two search points $x^1, x^2 \sim \text{Sample}(p)$, computes their fitness, sorts them by fitness, that is defines $(y^1, y^2) := (x^1, x^2)$ if x^1 is at least as fit as x^2 and $(y^1, y^2) := (x^2, x^1)$ otherwise, and updates the frequency vector to $p := p + \frac{1}{K}(y^1 - y^2)$, capped into the interval $[0, 1]$, that is, with entries below zero replaced by zero and entries above one replaced by one. This definition ensures that when y^1 and y^2 differ in some bit position i , the i -th frequency moves by a step of $\frac{1}{K}$ into the direction of y_i^1 (but not below zero and above one). The *hypothetical population size* K , often also denoted by μ , is an algorithm parameter that controls how strong this update is. To avoid a premature convergence, one often works with the *frequency boundaries* $\frac{1}{n}$ and $1 - \frac{1}{n}$, that is, one caps the new frequency vector into the interval $[\frac{1}{n}, 1 - \frac{1}{n}]$ instead of $[0, 1]$.

This iterative frequency evolution is pursued until some termination criterion is met. Since we aim at analyzing the time (number of iterations) it takes to sample the optimal solution (this is what we call the *runtime* of the cGA), we do not specify a termination criterion and pretend that the algorithm runs forever.

Algorithm 4: The compact genetic algorithm (cGA) to maximize a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

```

1  $p = (\frac{1}{2}, \dots, \frac{1}{2}) \in [0, 1]^n$ ;
2 repeat
3    $x^1 \sim \text{Sample}(p)$ ;
4    $x^2 \sim \text{Sample}(p)$ ;
5   if  $f(x^1) \geq f(x^2)$  then  $(y^1, y^2) \leftarrow (x^1, x^2)$  else  $(y^1, y^2) \leftarrow (x^2, x^1)$ ;
6    $p \leftarrow p + \frac{1}{K}(y^1 - y^2)$  capped into  $[0, 1]$  or  $[\frac{1}{n}, 1 - \frac{1}{n}]$ ;
7 until forever;
```

6.2 Central Results

In this section, we discuss three main insights which the theoretical analysis of EDAs has produced. For reasons of space, we point out two of them only briefly, namely that EDAs can perform well in noisy optimization and that they can cope well with local optima, and then discuss in detail how to set the parameters of EDA as this might be the biggest obstacle in successfully using EDAs.

6.2.1 EDAs Can Cope Well with Noise

In their remarkable work [FKKS17], Friedrich, Kötzing, Krejca, and Sutton exhibit that the cGA is extremely robust to noise. More precisely, they show that the cGA with a suitable parameter choice can optimize a ONEMAX function subject to additive normally distributed noise in a runtime that only polynomially depends on the variance σ^2 of the noise. As they also show, such a performance cannot be obtained with many classic EAs. The reason for this robustness is the cautious update of the probabilistic model in each iteration (as opposed to the “drastic” alternatives of a classic EA, rejecting an offspring or keeping it and discarding some other individual). This caution of the EDA implies that a single wrong evaluation of a search point only has a small influence on the future run of the algorithm. In the only other study on how EDAs cope with noise, Lehre and Nguyen [LN19b] show that the UMDA with suitable parameter choices can optimize the LEADINGONES problem in time $O(n^2)$ also in the presence of constant-probability one-bit prior noise.

6.2.2 EDAs Can Cope Well with Local Optima

Another difficulty for many EAs are local optima. Once the population of the EA is concentrated on the local optimum, it is difficult to leave this local optimum. As Hasenöhrl and Sutton [HS18] (see also [Doe19c]) show, the larger sampling variance of the cGA (in the regime without genetic drift) enables the algorithm to leave local optima much faster than many classic EAs. More specifically, Hasenöhrl and Sutton show that the cGA can optimize a jump function with jump size k in time $\exp(O(k + \log n))$, whereas many mutation-based EAs need time $\Omega(n^k)$.

6.2.3 Genetic Drift and Optimal Parameter Choices

Both the result on noisy optimization and the one on local optima indicate that EDAs can have significant advantages over classic EAs. For reasons of

brevity, we nevertheless omit further details and now turn to an important topic where a large sequence of works together have greatly increased our understanding, namely how to choose the parameters of EDAs and what is the role of genetic drift in EDAs.

While choosing optimal parameters for EAs is never easy, for many classic EAs a number of easy rules of thumb have been developed. For example, for mutation-based EAs the general recommendation to use standard-bit mutation with mutation rate $p = \frac{1}{n}$ often gives reasonable results (though [DLMN17] suggests that this impression is caused by an overfitting to unimodal problems). For EDAs, such general rules that are true over different classes of problems appear to be harder to find. From a large number of theoretical works, we now understand quite well why and we also have a number of different solutions to this problem.

The main challenge is choosing an appropriate speed of adapting the probabilistic model. If this speed of adaptation is low, then it simply takes a long time to change the initial, usually uniform, model into a model that samples good solutions with reasonable probability. However, if the speed of adaptation is high, then the small random signals stemming from the random choices in the sampling of solutions are over-interpreted and the model is quickly adjusted to an incorrect model. When an EDA without frequency boundaries is used, this means that the model has (at least partially) converged to an incorrect model without the possibility to ever return. With frequency boundaries, there is still the chance to revert to a good model, but practical experience and theory shows (i) that this can take a long time and (ii) that usually the EDA continues to work with degenerate models and thus, to some extent, imitates classic EAs (and consequently does not profit from the more general model-building ability). The effect that frequencies without a justification from the fitness function move to boundary values is known as *genetic drift*.

Since genetic drift can lead to significant performance problems and since the risk of encountering genetic drift via unfortunate parameter choices is high, the question how to avoid genetic drift is, explicitly or implicitly, a common theme of almost all theoretical works on EDAs. Shapiro’s very early works [Sha02, Sha05, Sha06] discussed this question explicitly, Droste’s first rigorous runtime analysis regarded how the cGA optimizes ONEMAX only when the update strength $\frac{1}{K}$ is $O(n^{-0.5-\varepsilon})$ for some constant $\varepsilon > 0$, a parameter regime in which the cGA with high probability finds the optimum of ONEMAX in a way that never a frequency goes below $\frac{1}{3}$, that is, without encountering genetic drift. For reasons of space, we shall not describe in detail the whole history of understanding genetic drift of EDAs, but present immediately the final result only mentioning that both explicit investigations

of genetic drift like [Sha02, Sha05, Sha06, FKK16, DZ20b] and the insights gained from many runtime analyses like [Dro06, DLN19, LN17, SW19, Wit19, KW20b, LSW21, HS18, Doe19c, Doe19b] paved the way towards this result.

Before discussing how to avoid genetic drift, let us quickly describe what is known about the danger of genetic drift. A first indication that genetic drift could be dangerous can be derived from the positive results – the majority of the proven upper bounds for runtimes of EDAs only apply to regimes in which there is provably no genetic drift, and in fact, most proofs heavily exploit this. Rigorous proofs that genetic drift can lead to performance losses are much more rare and appeared only very recently, owing to the fact that lower bound proofs for EDAs are often very difficult. In their deep analysis [LSW21], Lengler, Sudholt, and Witt showed that the cGA with $K = \Theta(n^{0.5}/(\log n \cdot \log \log n))$ needs time $\Omega(n^{7/6}/(\log n \cdot \log \log n))$ to optimize ONEMAX and the proof of this result shows that genetic drift is present. For $K = cn^{0.5} \ln n$, c a sufficiently large constant, the cGA only needs time $O(n \log n)$ and here no genetic drift occurs [SW19]. A more drastic loss from genetic drift, albeit on an artificial example problem, was observed in [LN19a, DK20c]. Lehre and Nguyen [LN19a] define the deceiving-leading-blocks (DLB)¹ problem and show that the UMDA with $\Omega(\log n) \leq \mu = o(n)$ needs time exponential in μ to find the optimum. By [DZ20b], in this parameter regime genetic drift is encountered when the runtime is $\omega(n^2)$. In [DK20c], it is shown that the UMDA with $\mu = \Theta(n \log n)$ can optimize the DLB problem in time $O(n^2 \log n)$ by profiting from the fact that now there is no genetic drift. A few experimental results also discuss the influence of genetic drift on the performance of an EDA, e.g., Figure 3 in [KW20a] shows the runtime of the UMDA on ONEMAX and Figure 1 in [DZ20a] shows the runtimes of the cGA on ONEMAX, LEADINGONES, jump functions and the DLB problem. These results show a mild negative impact of genetic drift in the two ONEMAX experiments, a stronger impact for LEADINGONES, and a drastic impact for jump functions and DLB.

We now discuss how to predict and avoid genetic drift. A good way to measure genetic drift is by regarding a fitness function with a neutral bit,

¹While not essential for the understanding of this section, for the sake of completeness we give a definition of the DLB problem. Let n be even. Then the DLB problem with problem size n is the following function f defined on bit strings of length n . The unique global optimum of f is $x^* = (1, \dots, 1)$, it has fitness $f(x^*) = n$. For all $x \in \{0, 1\}^n \setminus \{x^*\}$, let $\text{LB}(x) = \max\{i \in [0..n/2] \mid \forall j \in [1..2i] : x_j = 1\}$ denote the number of disjoint blocks of length two that are equal to $(1, 1)$, counted from left to right until a different block is encountered. We call this block $(x_{2\text{LB}(x)+1}, x_{2\text{LB}(x)+2})$ the critical block. The fitness $f(x)$ of x is $2\text{LB}(x) + 1$, if the critical block equals $(0, 0)$, and is $f(x) = 2\text{LB}(x)$ if the critical block is $(1, 0)$ or $(0, 1)$.

that is, a bit position that has no influence on the fitness. This might be overly pessimistic, since for such a bit the risk that the frequency approaches an unwanted boundary value might be higher than for a bit with strong influence on the fitness, but (i) a pessimistic view cannot be wrong here as a slightly too weak model update strength only slightly increases the runtime, whereas genetic drift as just seen can be detrimental, and (ii) the results just described show that the estimates from regarding neutral bits, for these examples, cannot be far from the truth.

The up to now most complete answer to the question of genetic drift was given in [DZ20b], as said, a work that would not exist without the long sequence of previous works named above. We discuss this result in detail for the cGA and note that similar results are true for the UMDA and PBIL.

Theorem 1. *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$. Assume that the i -th bit of f is neutral, that is, $f(x) = f(y)$ for all $x, y \in \{0, 1\}^n$ with $x_j = y_j$ for all $j \in [1..n] \setminus \{i\}$. Consider optimizing f via the cGA with hypothetical population size K using the frequency range $[\varepsilon, 1 - \varepsilon]$ for some $\varepsilon \in [0, \frac{1}{4}]$. Denote by $p^{(t)}$ the frequency vector resulting from the t -th iteration.*

1. *Let $T^* = \min\{t \mid p_i^{(t)} \in \{\varepsilon, 1 - \varepsilon\}\}$. Then $E[T^*] = O(K^2)$.*
2. *Let $T_{1/4} = \min\{t \mid p_i^{(t)} \in [0, \frac{1}{4}] \cup [\frac{3}{4}, 1]\} \leq T^*$ be the first time the i -th frequency leaves the interval $(\frac{1}{4}, \frac{3}{4})$ of the frequency range. Then $E[T_{1/4}] = \Omega(K^2)$.*
3. *For all $\gamma > 0$ and $T \in \mathbb{N}$, we have $\Pr[\forall t \in [0..T] : |p_i^{(t)} - \frac{1}{2}| < \gamma] \geq 1 - 2 \exp\left(-\frac{\gamma^2 K^2}{2T}\right)$.*

In very simple words the above result states that if we run the cGA for less than roughly K^2 iterations, then we do not encounter genetic drift, whereas after more than roughly K^2 iterations, genetic drift is likely to occur.

The tail bound (3) together with a simple union bound admits more precise guarantees, e.g., the following two formulations.

- If our aim is to run the cGA for T iterations on some pseudo-Boolean function of dimension n , then by taking $K \geq \sqrt{32T \ln(2n^2)}$ we can ensure that with probability at least $1 - \frac{1}{n}$ no neutral bit has its frequency leave the interval $(\frac{1}{4}, \frac{3}{4})$ within these T iterations.
- When K is given, the probability that within $T \leq \frac{K^2}{32 \ln(2n^2)}$ iterations a neutral frequency leaves the interval $(\frac{1}{4}, \frac{3}{4})$ is at most $\frac{1}{n}$.

We note without further details that similar statements hold for bits which are not neutral, but which have a preference for a particular value b , that is, where changing the bit-value to b can never decrease the fitness. Here the above statements hold for the undesired events that the frequency of this bit approaches the wrong boundary $1 - b$. We refer to [DZ20b] for a precise statement of this result. This extension allows one to determine good values for the hypothetical population size for simple test functions like ONEMAX or LEADINGONES: If we run the cGA on one of these functions for T iterations, then taking $K \geq \sqrt{32T \ln(2n^2)}$ ensures that with probability at least $1 - \frac{1}{n}$ no frequency will go below $\frac{1}{4}$.

For bit-values that have no uniform preference for a particular value (which is, naturally, the typical case for difficult optimization problems), we would still recommend to stick to the above-derived recommendations for setting K since this at least avoids that frequencies reach the wrong value due to genetic drift. If a fitness landscape is strongly deceptive (the fitness drags the typical heuristics away from the global optimum), clearly, such arguments cannot avoid that frequencies approach the wrong end of the frequency range due to the misleading fitness signal. We note though that the heuristic argument for setting K along the lines from above gives a good value and a good optimization behavior for the non-unimodal jump function class [HS18, Doe19c].

We finally note that there are three “automated” ways to approach the difficulty of finding the right parameter value. Inspired by the above insight, Doerr and Zheng [DZ20a] proposed to start with a small value of K , run the cGA until either a satisfying solution is found or the time exceeds a limit up to which we are sure to not observe genetic drift and then restart with twice the K -value. In [Doe19c], a strategy is proposed that in parallel works with different K -values. Both approaches were proven to optimize simple test functions in a time that is by at most a logarithmic factor larger than the runtime that can be obtained from using the optimal value of K . An experimental comparison [DZ20a] gives no clear picture which of the two approaches is superior. Clearly, both perform better than what results from a static, but inappropriate choice of K .

Instead of solving the genetic drift problem via a suitable choice of K , the significance-based cGA proposed in [DK20b] tries to avoid genetic drift outright. We recall that genetic drift is caused by random fluctuations of the frequencies, which again are caused by sampling search points from the probabilistic model and updating the model based on these. Therefore, the significance-based EDA avoids updating the model based on such short-sighted insights. Instead, this algorithm does not update the model until the history of the process gives sufficient evidence that some bit should bet-

ter have a particular value. In this case, a drastic model update is performed by setting the corresponding frequency to $\frac{1}{n}$ or $1 - \frac{1}{n}$. This algorithm was shown to optimize both `ONEMAX` and `LEADINGONES` in time $O(n \log n)$, a performance not observed with any other classic EA or EDA so far. With no research on non-unimodal objective functions and no practical experience so far, of course, this is still a very preliminary line of research.

6.3 Open Problems

Being a very recent research topic, it is clear that the theory of EDAs contains more open problems than solved ones, and many open problems are fundamental for our understanding and the future use of EDAs. We first mention briefly research topics where we feel that more results would greatly help and then give more details for two particular research questions.

- **Robust optimization:** The only two results [FKKS17, LN19b] here show that the cGA can efficiently optimize `ONEMAX` in the presence of normally distributed additive posterior noise and that the UMDA can efficiently optimize `LEADINGONES` in the presence of one-bit prior noise. Having such results for other EDAs, other optimization problems, and other noise models (and other stochastic disturbances such as dynamically changing problem instances) would be highly desirable.
- **Combinatorial optimization:** While for classic EAs a large number of runtime analyses for combinatorial optimization problems exist [NW10b], no such results have been shown for EDAs.
- **Representations different from bit strings:** For classic EAs, a number of results exist for problem representations different from bit strings, e.g., [STW04, DJ10, DDK18], and these results show that the choice of the representation and the choice of the variation operators for these can make a crucial difference. For EDAs, all results so far only discuss bit-string representations.

We now discuss in more detail two possible directions for future research.

6.3.1 Runtime Results in the Regime with Genetic Drift

In the regime without genetic drift, EDAs often show a regular optimization behavior which often allows one to prove matching upper and lower bounds for runtimes. In the presence of genetic drift, the runtime is strongly influenced by how some frequencies approach the boundaries of the frequency

range. It is thus rare events that determine the runtime and this makes it much harder to prove tight bounds. One could argue that runtime analyses in this regime are less interesting since we rather expect larger runtimes and rather an undesired behavior (e.g., imitating EAs), but this is not the full truth. For example, the UMDA optimizes ONEMAX in time $\Theta(n \log n)$ both for $\mu = \Theta(\log n)$ in the regime with (strong) genetic drift and for $\mu = \Theta(\sqrt{n} \log n)$ in the regime without genetic drift.

Apart from sporadic results, which most likely are not tight in most of the cases, not much is known about the runtimes of EDAs in the genetic drift regime. In particular, the following questions are not understood.

- Runtimes of EDAs for very high update strengths: The few runtime analyses in the genetic drift regime all assume that the update strength is at least so small that a (sufficiently large) logarithmic number of frequency updates is necessary to bring a frequency to a boundary value, e.g., that $K \geq C \ln n$ for some sufficiently large constant C when considering the cGA. Nothing nontrivial is known for even larger update strengths, but it is conjectured that one will typically encounter a super-polynomial runtime here.
- Runtimes of EDAs on ONEMAX for moderate update strengths. For the case that the update strength is smaller than in the previous paragraph, but still high enough to lead to genetic drift, a general $\Omega(n \log n)$ lower bound for the cGA and UMDA [KW20b, SW19] and an $O(n\lambda)$ upper bound for the UMDA [DLN19, Wit19] are known. For the cGA, a slightly stronger lower bound of $\Omega(K^{1/3}n)$ was shown in the regime $K = \Omega(\log^3 n) \cap O(\sqrt{n}/\log^2 n)$ [LSW21]. With this being all that is known, a true understanding of this regime is far from established.
- Runtimes of EDAs on jump functions: In the regime without genetic drift, a reasonable understanding of the runtime of the cGA on jump functions has been obtained in [HS18, Doe19c, Doe19b]. The lower bound [Doe19b], exponential in the jump size k , also applies to the regime with genetic drift, but is by far not sufficient to explain the huge runtimes observed experimentally [DZ20a] in this regime. Hence a proof that the cGA optimizing jump functions suffers significantly from genetic drift is still missing.

6.3.2 Multivariate EDAs

While multivariate EDAs are used a lot and very successfully in practice, essentially all theoretical research so far regarded univariate EDAs only, that

is, EDAs which evolve a univariate probabilistic model in which the bits are sampled independently. This is not surprising given how difficult it already was to obtain our limited understanding of univariate EDAs.

So far, only two theoretical results on multivariate EDAs exist, and both rely on theory-driven experiments and not on proven results. In [LN19a], the authors claim that the bivariate EDA *mutual information maximization for input clustering (MIMIC)* can cope better with fitness landscape in which the decision variables are interdependent. They define an artificial fitness landscape with strong inter-variable dependencies, the DLB-problem, prove that the UMDA with $\mu = o(n)$ needs time exponential in μ and show experimentally that the MIMIC can optimize this landscape in time polynomial in n . Based on this finding, they suggest “that one should consider EDAs with more complex probabilistic models when optimizing problems with some degree of epistasis and deception.” As discussed in Section 6.2.3, the lower bound on the runtime of the UMDA only applies to the regime with strong genetic drift and from $\mu = \Omega(n \log n)$ on, the runtime of the UMDA on DLB becomes $O(\mu n)$ [DK20c]. For this reason, it is not clear if the MIMIC, and more generally, bivariate EDAs, are also superior to the UMDA with the right choice of the parameters.

While so far no example exists in which a multivariate EDA was proven to have a better optimization behavior than a univariate one (with good parameters), the recent work [DK20a] shows (again only experimentally) that bivariate EDAs can evolve very expressive probabilistic models. For a simple fitness landscape with $2^{n/2}$ global optima it is shown that the MIMIC very quickly evolves a probabilistic model which allows it to sample global optima with constant chance and in a way that very rarely an optimum is sampled repeatedly. Hence the model evolved indeed represents to some extent the structure of the set of optimal solutions. It is clear that this would not be possible with a univariate EDA or a population-based EA.

In summary, there is a cautious indication from theoretical work that multivariate EDAs could be interesting both from the viewpoint of good optimization times and good representations of the structure of the fitness landscape, but almost all of the work in this direction still needs to be done.

7 Drift Analysis

Drift analysis has become one of the most heavily employed tools in the mathematical analysis of evolutionary algorithms (EAs). Interestingly, it is one of the few tool sets which were not imported from the classic algorithms

field. Rather, the classic algorithms field is now starting to use the drift theorems developed in our field, see, e.g., [BLM⁺20, GKK18, KU18, OE12].

Using drift analysis as a tool in the performance analysis of EAs builds on the insight that it is often easy to estimate the expected progress (with regard to some suitable measure) of an EA in one iteration. Drift analysis therefore tries to translate this information into estimates for the first time that a particular goal is achieved.

As a simple humorous example, inspired by a similar one from [Doe11], consider the following question. You have an initial capital of \$1,000. Each day you go to your favorite pub and drink a random number of beers for an expected total price of \$10. After how many days you are bankrupt?

If there was no randomness involved, that is, if you would spend exactly \$10 each day, then obviously it takes exactly 100 days to spend your money. So does the answer change with randomness? Interestingly, it does not (of course, we can now only talk about the expected number of days to bankruptcy): The expected number of days until you have spend all your money is exactly 100, regardless of the distribution of the amount you spent per day (which could be different for each day, could depend on previous days, and could also take negative values). This is a simple application of the additive drift theorem (Theorem 2 below).

The additive drift theorem is intuitive, but is in fact a deep mathematical result. Also, we have to note that it is not true that “randomness never changes things”. Take for example the opposite process: You start with no money, but each day you earn an expected number of ten dollars. What is the expected time it takes until you have at least \$1,000? Now we can only say that it is at least 100 days (with a slightly less direct application of the additive drift theorem), but it could be much larger. For example, if each day we earn \$10,000 with probability 0.001 and \$0 otherwise, then it takes an expected number of 1000 days until we have at least \$1,000.

Drift analysis was introduced to the field of evolutionary computation in the seminal paper [HY01] of He and Yao, which builds on Hajek’s fundamental work [Haj82]. Since some of the early uses of drift arguments led to quite technical proofs, many researchers first shied away from using this new method and preferred classic arguments like Wegener’s fitness level technique [Weg01]. Over time, however, more elegant applications of the additive drift theorem, e.g., in [Jäg08], and drift theorems capturing better particular scenarios, e.g., the multiplicative drift theorem [DJW12a], paved the way to drift analysis becoming perhaps the most powerful tool in the mathematical analysis of EAs.

7.1 Three True Drift Theorems

To show the beauty, simplicity, and power of drift analysis, we now present three central drift theorems. We call them *true drift theorems* to reflect that all three translate information on the expected one-step progress into a hitting time without further assumptions on the distribution of the one-step progress. We state these theorems in their most basic version and trust that the reader is able to derive seemingly more general, but equivalent versions via scaling, shifting, or mirroring the random process.

7.1.1 Additive Drift.

From a deeper mathematical result of Hajek [Haj82], He and Yao [HY01] derived the *additive drift theorem* and used it to prove several runtime bounds.

Theorem 2 (additive drift theorem). *Let X_0, X_1, \dots be a sequence of random variables taking values in some finite set $S \subseteq \mathbb{R}_{\geq 0}$ with $0 \in S$. Let $T = \inf\{t \mid X_t = 0\}$.*

- *Assume that there is a $\delta > 0$ such that for all $t \geq 0$ and $s \in S \setminus \{0\}$, we have $E[X_t - X_{t+1} \mid X_t = s] \geq \delta$. Then $E[T \mid X_0] \leq \frac{X_0}{\delta}$.*
- *Assume that there is a $\delta > 0$ such that for all $t \geq 0$ and $s \in S \setminus \{0\}$, we have $E[X_t - X_{t+1} \mid X_t = s] \leq \delta$. Then $E[T \mid X_0] \geq \frac{X_0}{\delta}$.*

Without going into details, we note that the assumptions can be weakened slightly, e.g., one can replace the “point-wise drift requirement”, that is, the conditioning on $X_t = s$, by an “average drift condition”, that is, conditioning only on $X_t > 0$ [HY01]. Also, the first part is also true for arbitrary infinite sets $S \subseteq \mathbb{R}_{\geq 0}$ and the second part is true also for bounded infinite sets S ; see [LS18], where also a short and elegant proof of this result is presented.

The additive drift theorem gives good results if there is a roughly uniform progress regardless of time and state. In fact, as the two estimates together show, the additive drift theorem gives an exact estimate for the hitting time T when the expected progress is known to be exactly δ at all times before hitting the target.

7.1.2 Multiplicative Drift.

For many natural optimization processes, the progress towards the optimum slows down when getting closer to the optimum. To use the additive drift theorem in such situations, the natural distance measure has to be transformed in such a way that the resulting expected progress is roughly uniform. Since

the expected transformed progress is usually not just the transformation of the expected progress, such proofs can become technical and unintuitive.

Noting that a common situation is that the expected progress is roughly proportional to the distance to the target, in [DJW12a] a *multiplicative drift theorem* was derived from the additive drift theorem. With a simpler direct proof, the following variant was later shown in [DG13]. According to [Len20], the multiplicative drift theorem is the most often used drift theorem in the theory of evolutionary algorithms.

Theorem 3 (multiplicative drift theorem). *Let X_0, X_1, \dots be a sequence of random variables over a state space $S \subseteq \{0\} \cup \mathbb{R}_{\geq 1}$ with $0 \in S$. Let $T = \min\{t \mid X_t = 0\}$. Assume that there is a $\delta > 0$ such that for all $t \geq 0$ and $s \in S \setminus \{0\}$, we have $E[X_{t+1} \mid X_t = s] \leq (1 - \delta)s$. Then the following estimates hold.*

- $E[T \mid X_0] \leq \frac{\ln(X_0)+1}{\delta}$.
- For all $\lambda > 0$, we have $\Pr[T > \lceil \frac{\ln(X_0)+\lambda}{\delta} \rceil] \leq \exp(-\lambda)$.

7.1.3 Variable Drift.

While indeed very many processes occurring in the analysis of evolutionary algorithms display an additive or multiplicative drift behavior, there remain processes in which the drift is decreasing when approaching the target (so that the additive drift theorem is hard to use), but not in a multiplicative fashion (so that the multiplicative drift theorem is hard to use). For these, so-called *variable drift theorems* can be applied. The first variable drift theorem for the analysis of evolutionary algorithms was proposed by Mitavskiy, Rowe, and Cannings [MRC09], however, the independently developed result of Johannsen [Joh10] appears to be used more often. The following is a variant of Johannsen’s result avoiding the use of integrals.

Theorem 4 (variable drift theorem). *Let X_0, X_1, \dots be a sequence of random variables over a finite space S . Assume that $S = \{s_0, \dots, s_M\}$ with $0 = s_0 < s_1 < \dots < s_M$. Let $T = \min\{t \mid X_t = 0\}$. Assume that there is a monotonically non-decreasing function $h : S \setminus \{0\} \rightarrow \mathbb{R}$ such that for all $t \geq 0$ and $s \in S \setminus \{0\}$, we have $E[X_t - X_{t+1} \mid X_t = s] \geq h(s)$. Then $E[T \mid X_0] \leq \sum_{i=1}^{X_0} \frac{s_i - s_{i-1}}{h(s_i)}$.*

The above are, most likely, the three most important drift theorems. We mention that the only other real drift theorem (that is, not requiring additional assumptions on the one-step distribution) we are aware of is the following result proven in [DLO19]: Let a random process as in

the multiplicative drift theorem be given, but with the drift condition $E[X_t - X_{t+1} \mid X_t = s] \geq \delta s$ replaced by the slightly stronger condition $E[X_t - X_{t+1} \mid X_t = s] \geq \delta s(\log_\gamma(s) + 1)$ for some $\gamma > 1$. Then $E[T \mid X_0] \leq \frac{3+4\ln\gamma+\max\{0,2\log_2\log_\gamma X_0\}}{\delta}$. We do not know if this result will find other applications, so we state the result here mainly to demonstrate that an only slightly stronger assumption on the drift – $\Omega(s \log s)$ instead of $\Omega(s)$ – can lead to a drastically smaller hitting time – $O(\log \log X_0)$ instead of $O(\log X_0)$.

7.2 Drift Results With Additional Requirements

The results presented in the previous section derive estimates for hitting times solely from the expected one-step progress; however, with two important restrictions: (i) except for the additive drift theorem, only upper bounds for hitting times can be obtained, and (ii) only processes can be analyzed in which there is a drift towards the target that can be uniformly bounded or that decreases when approaching the target.

Consequently, these drift theorems fail to describe a large number of behaviors of random processes that occur in the analysis of evolutionary algorithms. In this section, we briefly describe such behaviors and what solutions for their analysis exist. Unfortunately, and this is the reason why we shall state no precise result, all these tools not only require information on the expected one-step change, but also on the distribution of the one-step change (typically, that the one-step change is concentrated around its expectation). For all results, this is not a weakness of the result, but an intrinsic necessity.

7.2.1 Lower Bounds for Hitting Times.

From classic algorithms theory we know that it is very valuable to also have lower bounds on runtimes as these quantify how good our performance guarantees (upper bounds) are. If we have derived an upper bound from a certain drift behavior, say additive, multiplicative, or a certain variable drift, then the most natural approach would be to show a matching or near-matching upper bound on the expected one-step progress and derive (via a suitable drift theorem) from it a lower bound on the runtime. This works perfectly for the additive drift theorem as it contains such matching upper and lower bound results.

For multiplicative and variable drift, the theorems presented in the previous section are missing such matching results, and this for good reason, namely because in general they are not true. As a simple example, consider

the process on the state space $S = \{0, n\}$, starting with probability one in $X_0 = n$, which leaves state n to 0 with probability $1/n$ and stays in n otherwise. Apparently, we have $E[X_{t+1} | X_t = n] = n - 1 = (1 - \frac{1}{n})n$, that is, we have perfect multiplicative drift with $\delta = \frac{1}{n}$. The multiplicative drift theorem thus gives an estimate for the expected hitting time of $E[T] = O(n \log n)$. This is best-possible in the sense that there are processes with multiplicative drift with $\delta = \frac{1}{n}$ which indeed need $\Omega(n \log n)$ time, but for this particular process, the truth obviously is $E[T] = n$. This shows that a matching lower bound cannot exist without additional assumptions.

The assumption that usually gives the desired behavior (and the desired lower bounds) is that the one-step progress is concentrated around its expectation, typically with some exponential tails or by forbidding large progresses at all. We spare the details and point the reader to [Wit13, DDK18] for a multiplicative drift theorem for lower bounds and to [DFW11, GW18, DDY20] for variable drift theorems for lower bounds.

7.2.2 Increasing Drift.

All three main drift theorems require that the one-step progress is not increasing when approaching the target. This is a behavior often observed in evolutionary computation: The better the current solutions are, the harder it is to make progress. However, also the opposite behavior can be found, for example, when we consider how a better individual takes over a population. Here we would expect that the number of copies of the good individual increases in a multiplicative fashion (of course, only up to the point that a certain saturation is reached). Processes showing an increasing multiplicative drift have been analyzed in several papers dealing with population-based EAs, most notably in Lehre's [Leh11] fitness-level approach to non-elitist algorithms and in the so-called level-based theorem [CDEL18]. In both works, however, the increasing drift is mostly visible in the mathematical proofs. An explicit formulation of a drift result for increasing-drift processes was given in [DK19]. Again, an expected multiplicative one-step progress is not enough, but some additional concentration assumptions are necessary. Motivated by the application to population processes, the additional assumption was made that the one-step progress stochastically dominates a binomial distribution.

7.2.3 Negative Drift.

A different situation is that a process shows a drift away from the target and that we want to argue that it takes a long time to reach this target. Such a situation naturally arises again in lower-bound proofs. The first such

drift theorem was given by Oliveto and Witt [OW11, OW12]. Like many results proven later, see again the survey [Len20], it shows that if there is a constant negative expected progress in some interval of length ℓ and the one-step changes have both-sided exponential tails, then with probability $1 - \exp(-\Omega(\ell))$, the process takes time exponential in ℓ to reach the target.

A different approach to analyze a negative drift situation was taken in [ADY19]. Instead of the true process X_t , one regards an exponential transformation $Y_t = \exp(c(X_t - d))$ for suitable constants c, d , shows that Y_t has at most a constant additive drift, and then uses the lower bound part of the additive drift theorem to derive the desired result. Depending on how easy it is to compute the drift of the transformed process, this approach might be technically simpler than using the existing negative-drift theorems. Different from all existing negative-drift theorems, it allows one to derive explicit constants in the exponent. As shown in [Doe20a], this approach can also give super-exponential lower bounds.

Very recently, a negative drift theorem without additional constraints was presented in [Doe20b]. At the moment, it is hard to foresee if it will find other applications than those presented in [Doe20b].

7.3 Challenge: Finding the Right Potential Function

The results discussed so far show that we now have a decent number of drift theorems, which cover many different random processes. While surely new drift theorems will come up and existing ones will be polished, we are optimistic that the drift theorems developed in the last twenty years allow us to analyze most random processes occurring in the analysis of EAs.

What is less understood, and often still a challenge, is defining the right random process. To be able to apply a drift theorem, we need to define a random process (X_t) that describes some aspect of the run of our EA on some problem. Formally speaking, we need a function g that maps the full state S_t of the algorithm after iteration t into a real number $X_t = g(S_t)$, and this in a way that the process (X_t) still contains some relevant information of the run of the EA (e.g., that a suitable hitting time of (X_t) corresponds to the time when an optimum was first found) and in a way that a drift theorem can be applied. While there are some generic solutions to this technical problem, many questions are still open here and this might be the biggest challenge in the future of drift analysis.

A natural way to define the potential function g is to take the fitness distance of the current-best solution to the optimum. This works well when there is a good correlation between the remaining optimization time and the fitness distance as observed, e.g., for the simple benchmarks ONEMAX and

LEADINGONES (note that the classic analyses [DJW02] stem from the time before drift analysis was introduced and hence use Wegener’s [Weg01] fitness level method) or combinatorial problems such as the minimum spanning tree problem (again, the classic proof [NW07] does not use drift, but the expected multiplicative weight decrease method) or the maximum satisfiability problem with clauses of length 3 [DNS17].

An equally natural potential is the structural distance to the optimum, e.g., the Hamming distance in the case of pseudo-Boolean optimization. This was used, e.g., to show that the $(1 + 1)$ EA with mutation rate c/n , c a constant strictly between 0 and 1, optimizes any strictly monotonic function in time $O(n \log n)$ [DJS⁺13].

Unfortunately, these two potentials do not suffice. The most famous example for which neither the fitness distance nor the structural distance work well is the linear functions problem, that is, the question how fast the simple $(1 + 1)$ EA optimizes a general linear pseudo-Boolean function $f(x) = \sum_{i=1}^n a_i x_i$. With a sequence of more powerful potential functions, all different from fitness and structural distance, increasingly strong results were obtained [DJW02, DJW12a, DG13, Wit13].² Unfortunately, it remains unclear how to easily derive such potential functions. In fact, the only result regarding this question is a negative one, namely that to prove the results for larger mutation rates such as [DG13, Wit13], it is not possible to use one “universal” potential function for all linear functions, but the potential has to be chosen depending on the problem instance [DJW12b].

In three particular directions, we currently see the greatest lack of understanding how to define potential functions to use drift analysis. These are the following.

7.3.1 Drift Analysis for Representations Other Than Bit Strings.

Once a relatively compact analysis of the runtime of the $(1 + 1)$ EA with standard mutation rate $1/n$ on linear functions was found [DJW12a], the question was raised how far these methods could be extended. One direction are linear functions defined not on bit strings, but on higher-arity representations $\{0, \dots, r\}^n$. While the $O(n \log n)$ runtime estimate could be shown for the search space $\{0, 1, 2\}^n$ [DJS11], it was also shown in this work that there is no universal potential function from $r \geq 43$ on. With instance-specific potential functions, an $O(rn \log n + r^3 n \log \log n)$ upper bound was shown in [DP12]. This extends the $O(rn \log n)$ bound to all $r = O((\frac{\log n}{\log \log n})^{1/2})$, but not beyond. It is an open problem whether larger r indeed lead to an

²We note that Jägersküpfer with a clever averaging argument could also use the structural distance as potential function.

inferior runtime behavior or not. This example and the general shortage of works analyzing EAs with representations different from bit strings via drift analysis (we are only aware of [KLW15b, LW16, DDK18]) suggest that more work is needed in this direction.

7.3.2 Drift Analysis for Population-based EAs.

All works described above, and in general the vast majority of runtime analyses building on drift arguments, only regard very simple EAs such as the $(1 + 1)$ EA or, occasionally, the $(1 + \lambda)$ EA or the $(1 + (\lambda, \lambda))$ GA. For such EAs, a potential function only needs to estimate the quality of the single parent individual. For EAs working with a non-trivial parent population, it is much harder to define a suitable potential function. In fact, the main work on lower bounds for such algorithms by Lehre [Leh10] used drift arguments only in the ancestral lines of single individuals and captured the effect of the whole population via family trees (see [Doe20b] for an alternative approach). Again for lower bound proofs, Neumann, Oliveto, and Witt [NOW09] and later [OW15, ADY19] used $\sum_{x \in P} c^{\text{ONEMAX}(x)}$ as potential (to be maximized) of a population P in an algorithm maximizing the ONEMAX benchmark, where $c > 1$ is a suitable constant. The main line of works on upper bounds [Leh11, DL15b, CDEL18] defines an extremely complicated potential function on the populations that is, most likely, not easy to transfer to problems not covered by these results. While these results show that drift analysis can be employed in the analysis of population-based algorithms, more work seems necessary, in particular, on upper bounds analyses, for drift analysis to replace the classic fitness level arguments more often employed here, see, e.g., [Wit06, CHS⁺09, ADFH18].

7.3.3 Drift Analysis for Dynamic Parameters.

With the popularity of dynamic parameter choices both in theory (see also Section 3.2) and practice, there is a strong need for mathematical methods to analyze such algorithms. From the perspective of drift analysis, again the challenge is to define a suitable potential function on the cross product of populations (which in the simplest case are just single individuals) and parameter values (or more generally, the full inner state of the algorithm). So far, we are only aware of the five works [DDK18, AAG18, Row18, DWY21, CL20] providing solutions to this problem. In the interest of brevity, we refer to [DWY21, Section 1.3] for a more detailed discussion, and state here only that our impression is that more work on this problem is necessary (and desirable) to ease future analyses of dynamic parameter settings.

8 Final Words

We provided an overview on areas of research in the field of theory of evolutionary computation in discrete search spaces that have gained significant attention during the last 10 years. The survey tried to capture the most important aspects from the perspective of the authors. We refer to the recent edited book [DN20b] for a more comprehensive overview, which also includes other evolutionary computing techniques such as genetic programming and artificial immune systems. For the true technical details, naturally, we invite the reader to consult the original articles.

There are many areas where we see a lot of room for progress. Analyses for constrained problems static, dynamic, or stochastic have just recently been started and understanding the behavior of evolutionary algorithms for linear functions even very special simple constraints is still a challenging task [NPW19]. A first theoretical analysis of differential evolution in discrete search spaces has been carried out in [ZYD18], however, indicating that our current methods cannot cope well with the complicated stochastic dependencies arising in this optimization process. The entropy compression method has found a first application in evolutionary computation [LMS19], but other applications of this powerful methods are not in sight. From a broader perspective, our understanding of the impact of populations, crossover operators, and diversity mechanisms still lags behind their practical success and proving the usefulness of such modules of an evolutionary algorithm for complex optimization problems is a challenging task.

We hope that the survey helps the readers to pursue their own research in this area. Although tremendous progress has been made during the last 10 years, there are still a lot of open questions and problems, some of which have been outlined in this article. We encourage the reader to make their own contribution to this field of research and help to transfer theoretical knowledge into the design of high performing evolutionary computing techniques.

Acknowledgements

Frank Neumann has been support by the Alexander von Humboldt Foundation through a Humboldt Fellowship for Experienced Researchers and by the Australian Research Council through grant FT200100536.

References

- [AAG18] Youhei Akimoto, Anne Auger, and Tobias Glasmachers. Drift theory in continuous search spaces: expected hitting time of the $(1 + 1)$ -ES with $1/5$ success rule. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 801–808. ACM, 2018.
- [ABD20a] Denis Antipov, Maxim Buzdalov, and Benjamin Doerr. Fast mutation in crossover-based algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2020*, pages 1268–1276. ACM, 2020.
- [ABD20b] Denis Antipov, Maxim Buzdalov, and Benjamin Doerr. First steps towards a runtime analysis when starting with a good solution. In *Parallel Problem Solving From Nature, PPSN 2020, Part II*, pages 560–573. Springer, 2020.
- [Ack87] David H. Ackley. *A Connectionist Machine for Genetic Hill Climbing*. Kluwer, 1987.
- [AD20] Denis Antipov and Benjamin Doerr. Runtime analysis of a heavy-tailed $(1 + (\lambda, \lambda))$ genetic algorithm on jump functions. In *Parallel Problem Solving From Nature, PPSN 2020, Part II*, pages 545–559. Springer, 2020.
- [ADFH18] Denis Antipov, Benjamin Doerr, Jiefeng Fang, and Tangi Hetet. Runtime analysis for the $(\mu + \lambda)$ EA optimizing OneMax. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1459–1466. ACM, 2018.
- [ADY19] Denis Antipov, Benjamin Doerr, and Quentin Yang. The efficiency threshold for the offspring population size of the (μ, λ) EA. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1461–1469. ACM, 2019.
- [AF08] Nattapat Attiratanasunthron and Jittat Fakcharoenphol. A running time analysis of an ant colony optimization algorithm for shortest paths in directed acyclic graphs. *Information Processing Letters*, 105:88–92, 2008.
- [AHX⁺20] Hiran Assimi, Oscar Harper, Yue Xie, Aneta Neumann, and Frank Neumann. Evolutionary bi-objective optimization for

- the dynamic chance-constrained knapsack problem based on tail bound objectives. In *ECAI*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 307–314. IOS Press, 2020.
- [Aro98] Sanjeev Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.
- [Bäc92] Thomas Bäck. Self-adaptation in genetic algorithms. In *European Conference on Artificial Life, ECAL 1992*, pages 263–271. MIT Press, 1992.
- [Bäc93] Thomas Bäck. Optimal mutation rates in genetic search. In *International Conference on Genetic Algorithms, ICGA 1993*, pages 2–8. Morgan Kaufmann, 1993.
- [Bäc96] Thomas Bäck. *Evolutionary Algorithms in Theory and Practice – Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, 1996.
- [Bal94] Shumeet Baluja. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Technical report, Carnegie Mellon University, 1994.
- [BD17] Maxim Buzdalov and Benjamin Doerr. Runtime analysis of the $(1 + (\lambda, \lambda))$ genetic algorithm on random satisfiable 3-CNF formulas. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 1343–1350. ACM, 2017.
- [BDDV20] Maxim Buzdalov, Benjamin Doerr, Carola Doerr, and Dmitry Vinokurov. Fixed-target runtime analysis. In *GECCO*, pages 1295–1303. ACM, 2020.
- [BDM15] Weiwei Bi, Graeme C. Dandy, and Holger R. Maier. Improved genetic algorithm optimization of water distribution system design by incorporating domain knowledge. *Environ. Model. Softw.*, 69:370–381, 2015.
- [BDN10] Süntje Böttcher, Benjamin Doerr, and Frank Neumann. Optimal fixed and adaptive mutation rates for the LeadingOnes problem. In *Parallel Problem Solving from Nature, PPSN 2010*, pages 1–10. Springer, 2010.

- [BFH⁺09] Dimo Brockhoff, Tobias Friedrich, Nils Hebbinghaus, Christian Klein, Frank Neumann, and Eckart Zitzler. On the effects of adding objectives to plateau functions. *IEEE Trans. Evol. Comput.*, 13(3):591–603, 2009.
- [BFM97] Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz. *Handbook of Evolutionary Computation*. IOP Publishing Ltd., 1997.
- [BFQY20] Chao Bian, Chao Feng, Chao Qian, and Yang Yu. An efficient evolutionary algorithm for subset selection with general cost constraints. In *AAAI*, pages 3267–3274. AAAI Press, 2020.
- [BLM⁺20] Daniel Bertschinger, Johannes Lengler, Anders Martinsson, Robert Meier, Angelika Steger, Milos Trujic, and Emo Welzl. An optimal decentralized $(\delta + 1)$ -coloring algorithm. In *European Symposium on Algorithms, ESA 2020*, volume 173, pages 17:1–17:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [BLS14] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. Unbiased black-box complexity of parallel search. In *Parallel Problem Solving from Nature, PPSN 2014*, pages 892–901. Springer, 2014.
- [BM16] Mohammad Reza Bonyadi and Zbigniew Michalewicz. Evolutionary computation for real-world problems. In *Challenges in Computational Statistics and Data Mining*, volume 605 of *Studies in Computational Intelligence*, pages 1–24. Springer, 2016.
- [BNPS20] Jakob Bossek, Frank Neumann, Pan Peng, and Dirk Sudholt. More effective randomized search heuristics for graph coloring through dynamic optimization. In *GECCO*, pages 1277–1285. ACM, 2020.
- [BNPS21] Jakob Bossek, Frank Neumann, Pan Peng, and Dirk Sudholt. Time complexity analysis of randomized search heuristics for the dynamic graph coloring problem. *Algorithmica*, 2021. doi:10.1007/s00453-021-00838-3.
- [BP14] Thomas Bartz-Beielstein and Mike Preuss. Experimental analysis of optimization algorithms: Tuning and beyond. In *Theory and Principled Methods for the Design of Metaheuristics*, Natural Computing Series, pages 205–245. Springer, 2014.

- [BQT18] Chao Bian, Chao Qian, and Ke Tang. Towards a running time analysis of the (1+1)-EA for OneMax and LeadingOnes under general bit-wise noise. In *Parallel Problem Solving from Nature, PPSN 2018, Part II*, pages 165–177. Springer, 2018.
- [Bra00] Jürgen Branke. *Evolutionary optimization in dynamic environments*. PhD thesis, Universität Karlsruhe, 2000.
- [CC84] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete Applied Mathematics*, 7(3):251–274, 1984.
- [CDEL18] Dogan Corus, Duc-Cuong Dang, Anton V. Eremeev, and Per Kristian Lehre. Level-based analysis of genetic algorithms and other search processes. *IEEE Transactions on Evolutionary Computation*, 22:707–719, 2018.
- [CHS⁺09] Tianshi Chen, Jun He, Guangzhong Sun, Guoliang Chen, and Xin Yao. A new approach for analyzing average time complexity of population-based evolutionary algorithms on unimodal problems. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 39:1092–1106, 2009.
- [CL20] Brendan Case and Per Kristian Lehre. Self-adaptation in nonelitist evolutionary algorithms on discrete problems with unknown structure. *IEEE Transactions on Evolutionary Computation*, 24:650–663, 2020.
- [CLTY09] Tianshi Chen, Per Kristian Lehre, Ke Tang, and Xin Yao. When is an estimation of distribution algorithm better than an evolutionary algorithm? In *Congress on Evolutionary Computation, CEC 2009*, pages 1470–1477. IEEE, 2009.
- [CTCY10] Tianshi Chen, Ke Tang, Guoliang Chen, and Xin Yao. Analysis of computational time of simple estimation of distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 14:1–22, 2010.
- [DD18] Benjamin Doerr and Carola Doerr. Optimal static and self-adjusting parameter choices for the $(1 + (\lambda, \lambda))$ genetic algorithm. *Algorithmica*, 80:1658–1709, 2018.

- [DDE15] Benjamin Doerr, Carola Doerr, and Franziska Ebel. From black-box complexity to designing new genetic algorithms. *Theoretical Computer Science*, 567:87–104, 2015.
- [DDK18] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Static and self-adjusting mutation strengths for multi-valued decision variables. *Algorithmica*, 80:1732–1768, 2018.
- [DDK19] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Solving problems with unknown solution length at almost no extra cost. *Algorithmica*, 81:703–748, 2019.
- [DDL19] Benjamin Doerr, Carola Doerr, and Johannes Lengler. Self-adjusting mutation rates with provably optimal success rules. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1479–1487. ACM, 2019.
- [DDN⁺20] Benjamin Doerr, Carola Doerr, Aneta Neumann, Frank Neumann, and Andrew M. Sutton. Optimization of chance-constrained submodular functions. In *AAAI*, pages 1460–1467. AAAI Press, 2020.
- [DDY16] Benjamin Doerr, Carola Doerr, and Jing Yang. k -bit mutation with self-adjusting k outperforms standard bit mutation. In *Parallel Problem Solving from Nature, PPSN 2016*, pages 824–834. Springer, 2016.
- [DDY20] Benjamin Doerr, Carola Doerr, and Jing Yang. Optimal parameter choices via precise black-box analysis. *Theoretical Computer Science*, 801:1–34, 2020.
- [DF99] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer, 1999.
- [DFW11] Benjamin Doerr, Mahmoud Fouz, and Carsten Witt. Sharp bounds by probability-generating functions and variable drift. In *Genetic and Evolutionary Computation Conference, GECCO 2011*, pages 2083–2090. ACM, 2011.
- [DG13] Benjamin Doerr and Leslie A. Goldberg. Adaptive drift analysis. *Algorithmica*, 65:224–250, 2013.
- [DGWY19] Benjamin Doerr, Christian Gießen, Carsten Witt, and Jing Yang. The $(1 + \lambda)$ evolutionary algorithm with self-adjusting mutation rate. *Algorithmica*, 81:593–631, 2019.

- [DHK12] Benjamin Doerr, Ashish Ranjan Hota, and Timo Kötzing. Ants easily solve stochastic shortest path problems. In *Genetic and Evolutionary Computation Conference, GECCO 2012*, pages 17–24. ACM, 2012.
- [DHOW06] Vladimir G. Deineko, Michael Hoffmann, Yoshio Okamoto, and Gerhard J. Woeginger. The traveling salesman problem with few inner points. *Oper. Res. Lett.*, 34(1):106–110, 2006.
- [DJ07] Benjamin Doerr and Daniel Johannsen. Refined runtime analysis of a basic ant colony optimization algorithm. In *Congress on Evolutionary Computation, CEC 2007*, pages 501–507. IEEE, 2007.
- [DJ10] Benjamin Doerr and Daniel Johannsen. Edge-based representation beats vertex-based representation in shortest path problems. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 759–766. ACM, 2010.
- [DJL17] Duc-Cuong Dang, Thomas Jansen, and Per Kristian Lehre. Populations can be essential in tracking dynamic optima. *Algorithmica*, 78(2):660–680, 2017.
- [DJS11] Benjamin Doerr, Daniel Johannsen, and Martin Schmidt. Runtime analysis of the (1+1) evolutionary algorithm on strings over finite alphabets. In *Foundations of Genetic Algorithms, FOGA 2011*, pages 119–126. ACM, 2011.
- [DJS⁺13] Benjamin Doerr, Thomas Jansen, Dirk Sudholt, Carola Winzen, and Christine Zarges. Mutation rate matters even when optimizing monotone functions. *Evolutionary Computation*, 21:1–21, 2013.
- [DJW00] Stefan Droste, Thomas Jansen, and Ingo Wegener. Dynamic parameter control in simple evolutionary algorithms. In *Foundations of Genetic Algorithms, FOGA 2000*, pages 275–294. Morgan Kaufmann, 2000.
- [DJW02] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.
- [DJW12a] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64:673–697, 2012.

- [DJW12b] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Non-existence of linear universal drift functions. *Theoretical Computer Science*, 436:71–86, 2012.
- [DJWZ13] Benjamin Doerr, Thomas Jansen, Carsten Witt, and Christine Zarges. A method to derive fixed budget results from expected optimisation times. In *GECCO*, pages 1581–1588. ACM, 2013.
- [DK11] Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *ICML*, pages 1057–1064. Omnipress, 2011.
- [DK15] Benjamin Doerr and Marvin Künnemann. Optimizing linear functions with the $(1 + \lambda)$ evolutionary algorithm—different asymptotic runtimes for different instances. *Theoretical Computer Science*, 561:3–23, 2015.
- [DK19] Benjamin Doerr and Timo Kötzing. Multiplicative up-drift. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1470–1478. ACM, 2019.
- [DK20a] Benjamin Doerr and Martin S. Krejca. Bivariate estimation-of-distribution algorithms can find an exponential number of optima. In *Genetic and Evolutionary Computation Conference, GECCO 2020*, pages 796–804. ACM, 2020.
- [DK20b] Benjamin Doerr and Martin S. Krejca. Significance-based estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24:1025–1034, 2020.
- [DK20c] Benjamin Doerr and Martin S. Krejca. The univariate marginal distribution algorithm copes well with deception and epistasis. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2020*, pages 51–66. Springer, 2020.
- [DL15a] Duc-Cuong Dang and Per Kristian Lehre. Efficient optimisation of noisy fitness functions with population-based evolutionary algorithms. In *FOGA*, pages 62–68. ACM, 2015.
- [DL15b] Duc-Cuong Dang and Per Kristian Lehre. Simplified runtime analysis of estimation of distribution algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2015*, pages 513–518. ACM, 2015.

- [DL16] Duc-Cuong Dang and Per Kristian Lehre. Self-adaptation of mutation rates in non-elitist populations. In *Parallel Problem Solving from Nature, PPSN 2016*, pages 803–813. Springer, 2016.
- [DLMN17] Benjamin Doerr, Huu Phuoc Le, Régis Makhlara, and Ta Duy Nguyen. Fast genetic algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 777–784. ACM, 2017.
- [DLN19] Duc-Cuong Dang, Per Kristian Lehre, and Phan Trung Hai Nguyen. Level-based analysis of the univariate marginal distribution algorithm. *Algorithmica*, 81:668–702, 2019.
- [DLO19] Benjamin Doerr, Andrei Lissovoi, and Pietro Simone Oliveto. Evolving boolean functions with conjunctions and disjunctions via genetic programming. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1003–1011. ACM, 2019.
- [DLOW18] Benjamin Doerr, Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. On the runtime analysis of selection hyper-heuristics with adaptive learning periods. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1015–1022. ACM, 2018.
- [DN20a] Anh Viet Do and Frank Neumann. Maximizing submodular or monotone functions under partition matroid constraints by multi-objective evolutionary algorithms. In *PPSN (2)*, volume 12270 of *Lecture Notes in Computer Science*, pages 588–603. Springer, 2020.
- [DN20b] Benjamin Doerr and Frank Neumann, editors. *Theory of Evolutionary Computation—Recent Developments in Discrete Optimization*. Springer, 2020. Also available at <https://cs.adelaide.edu.au/~frank/papers/TheoryBook2019-selfarchived.pdf>.
- [DNDD⁺18] Raphaël Dang-Nhu, Thibault Dardinier, Benjamin Doerr, Gautier Izacard, and Dorian Nogneng. A new analysis method for evolutionary optimization of dynamic and noisy objective functions. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1467–1474. ACM, 2018.

- [DNS17] Benjamin Doerr, Frank Neumann, and Andrew M. Sutton. Time complexity analysis of evolutionary algorithms on random satisfiable k -CNF formulas. *Algorithmica*, 78:561–586, 2017.
- [DNSW11] Benjamin Doerr, Frank Neumann, Dirk Sudholt, and Carsten Witt. Runtime analysis of the 1-ANT ant colony optimizer. *Theoretical Computer Science*, 412:1629–1644, 2011.
- [Doe11] Benjamin Doerr. Drift analysis. In *Genetic and Evolutionary Computation Conference, GECCO 2011, Companion Material*, pages 1311–1320. ACM, 2011.
- [Doe19a] Benjamin Doerr. Analyzing randomized search heuristics via stochastic domination. *Theoretical Computer Science*, 773:115–137, 2019.
- [Doe19b] Benjamin Doerr. An exponential lower bound for the runtime of the compact genetic algorithm on jump functions. In *Foundations of Genetic Algorithms, FOGA 2019*, pages 25–33. ACM, 2019.
- [Doe19c] Benjamin Doerr. A tight runtime analysis for the cGA on jump functions: EDAs can cross fitness valleys at no extra cost. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1488–1496. ACM, 2019.
- [Doe20a] Benjamin Doerr. Does comma selection help to cope with local optima? In *Genetic and Evolutionary Computation Conference, GECCO 2020*, pages 1304–1313. ACM, 2020.
- [Doe20b] Benjamin Doerr. Lower bounds for non-elitist evolutionary algorithms via negative multiplicative drift. In *Parallel Problem Solving From Nature, PPSN 2020, Part II*, pages 604–618. Springer, 2020.
- [Dos13] Martin Dostál. Evolutionary music composition. In *Handbook of Optimization*, volume 38 of *Intelligent Systems Reference Library*, pages 935–964. Springer, 2013.
- [DP12] Benjamin Doerr and Sebastian Pohl. Run-time analysis of the (1+1) evolutionary algorithm optimizing linear functions over a finite alphabet. In *Genetic and Evolutionary Computation Conference, GECCO 2012*, pages 1317–1324. ACM, 2012.

- [Dro02] Stefan Droste. Analysis of the (1+1) EA for a dynamically changing OneMax-variant. In *Congress on Evolutionary Computation, CEC 2002*, pages 55–60. IEEE, 2002.
- [Dro03] Stefan Droste. Analysis of the (1+1) EA for a dynamically bitwise changing onemax. In *GECCO*, volume 2723 of *Lecture Notes in Computer Science*, pages 909–921. Springer, 2003.
- [Dro04] Stefan Droste. Analysis of the (1+1) EA for a noisy onemax. In *GECCO (1)*, volume 3102 of *Lecture Notes in Computer Science*, pages 1088–1099. Springer, 2004.
- [Dro06] Stefan Droste. A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing*, 5:257–283, 2006.
- [DWY21] Benjamin Doerr, Carsten Witt, and Jing Yang. Runtime analysis for self-adaptive mutation rates. *Algorithmica*, 83:1012–1053, 2021.
- [DZ20a] Benjamin Doerr and Weijie Zheng. From understanding genetic drift to a smart-restart parameter-less compact genetic algorithm. In *Genetic and Evolutionary Computation Conference, GECCO 2020*, pages 805–813. ACM, 2020.
- [DZ20b] Benjamin Doerr and Weijie Zheng. Sharp bounds for genetic drift in estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24:1140–1149, 2020.
- [EKM19] Hossein Esfandiari, Amin Karbasi, and Vahab S. Mirrokni. Adaptivity in adaptive submodularity. *CoRR*, abs/1911.03620, 2019.
- [FGN⁺19] Tobias Friedrich, Andreas Göbel, Frank Neumann, Francesco Quinzan, and Ralf Rothenberger. Greedy maximization of functions with bounded curvature under partition matroid constraints. In *AAAI*, pages 2272–2279. AAAI Press, 2019.
- [FGQW18a] Tobias Friedrich, Andreas Göbel, Francesco Quinzan, and Markus Wagner. Evolutionary algorithms and submodular functions: Benefits of heavy-tailed mutations. *CoRR*, abs/1805.10902, 2018.

- [FGQW18b] Tobias Friedrich, Andreas Göbel, Francesco Quinzan, and Markus Wagner. Heavy-tailed mutation operators in single-objective combinatorial optimization. In *Parallel Problem Solving from Nature, PPSN 2018, Part I*, pages 134–145. Springer, 2018.
- [FHH⁺10] Tobias Friedrich, Jun He, Nils Hebbinghaus, Frank Neumann, and Carsten Witt. Approximating covering problems by randomized search heuristics using multi-objective models. *Evolutionary Computation*, 18:617–633, 2010.
- [FK13] Matthias Feldmann and Timo Kötzing. Optimizing expected path lengths with ant colony optimization using fitness proportional update. In *Foundations of Genetic Algorithms, FOGA 2013*, pages 65–74. ACM, 2013.
- [FKK16] Tobias Friedrich, Timo Kötzing, and Martin S. Krejca. EDAs cannot be balanced and stable. In *Genetic and Evolutionary Computation Conference, GECCO 2016*, pages 1139–1146. ACM, 2016.
- [FKKS16] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. Robustness of ant colony optimization to noise. *Evolutionary Computation*, 24:237–254, 2016.
- [FKKS17] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. The compact genetic algorithm is efficient under extreme Gaussian noise. *IEEE Transactions on Evolutionary Computation*, 21:477–490, 2017.
- [FKQS17] Tobias Friedrich, Timo Kötzing, Francesco Quinzan, and Andrew M. Sutton. Resampling vs recombination: a statistical run time estimation. In *FOGA*, pages 25–35. ACM, 2017.
- [FN14] Tobias Friedrich and Frank Neumann. Maximizing submodular functions under matroid constraints by multi-objective evolutionary algorithms. In *13th International Conference on Parallel Problem Solving from Nature (PPSN)*, volume 8672 of *Lecture Notes in Computer Science*, pages 922–931. Springer, 2014.
- [FQW18] Tobias Friedrich, Francesco Quinzan, and Markus Wagner. Escaping large deceptive basins of attraction with heavy-tailed mutation operators. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 293–300. ACM, 2018.

- [GK11] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res.*, 42:427–486, 2011.
- [GK16] Christian Gießen and Timo Kötzing. Robustness of populations in stochastic environments. *Algorithmica*, 75(3):462–489, 2016.
- [GKK18] Andreas Göbel, Timo Kötzing, and Martin S. Krejca. Intuitive analyses via drift theory. *CoRR*, abs/1806.01919, 2018. arXiv:1806.01919.
- [GKS99] Josselin Garnier, Leila Kallel, and Marc Schoenauer. Rigorous hitting times for binary mutations. *Evolutionary Computation*, 7:173–203, 1999.
- [GL10] Oliver Giel and Per Kristian Lehre. On the effect of populations in evolutionary multi-objective optimisation. *Evol. Comput.*, 18(3):335–356, 2010.
- [GS08] Walter J. Gutjahr and Giovanni Sebastiani. Runtime analysis of ant colony optimization with best-so-far reinforcement. *Methodology and Computing in Applied Probability*, 10:409–433, 2008.
- [Gut00] Walter J. Gutjahr. A graph-based ant system and its convergence. *Future Generation Computer Systems*, 16:873–888, 2000.
- [Gut02] Walter J. Gutjahr. ACO algorithms with guaranteed convergence to the optimal solution. *Information Processing Letters*, 82:145–153, 2002.
- [Gut03] Walter J. Gutjahr. A converging ACO algorithm for stochastic combinatorial optimization. In *Stochastic Algorithms: Foundations and Applications, SAGA 2003*, pages 10–25. Springer, 2003.
- [Gut08] Walter J. Gutjahr. First steps to the runtime complexity analysis of ant colony optimization. *Computers & Operations Research*, 35:2711–2727, 2008.
- [Gut11] Walter J. Gutjahr. Ant colony optimization: recent developments in theoretical analysis. In Anne Auger and Benjamin Doerr, editors, *Theory of Randomized Search Heuristics: Foundations and Recent Developments*, volume 1, pages 225–254. World Scientific, 2011.

- [GW17] Christian Gießen and Carsten Witt. The interplay of population size and mutation probability in the $(1 + \lambda)$ EA on OneMax. *Algorithmica*, 78:587–609, 2017.
- [GW18] Christian Gießen and Carsten Witt. Optimal mutation rates for the $(1 + \lambda)$ EA on OneMax through asymptotically tight drift analysis. *Algorithmica*, 80:1710–1731, 2018.
- [Haj82] Bruce Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 13:502–525, 1982.
- [HLG99] Georges R. Harik, Fernando G. Lobo, and David E. Goldberg. The compact genetic algorithm. *IEEE Transactions on Evolutionary Computation*, 3:287–297, 1999.
- [HO07] Franklin Hanshar and Beatrice M. Ombuki-Berman. Dynamic vehicle routing using genetic algorithms. *Appl. Intell.*, 27(1):89–99, 2007.
- [HP11] Mark Hauschild and Martin Pelikan. An introduction and survey of estimation of distribution algorithms. *Swarm and Evolutionary Computation*, 1:111–128, 2011.
- [HR97] Markus Hohfeld and Günter Rudolph. Towards a theory of population-based incremental learning. In *Conference on Evolutionary Computation*, pages 1–5. IEEE Press, 1997.
- [HS17] Avinatan Hassidim and Yaron Singer. Submodular optimization under noise. In *COLT*, volume 65 of *Proceedings of Machine Learning Research*, pages 1069–1122. PMLR, 2017.
- [HS18] Václav Hasenöhrl and Andrew M. Sutton. On the runtime dynamics of the compact genetic algorithm on jump functions. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 967–974. ACM, 2018.
- [HY01] Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127:51–81, 2001.
- [Jäg08] Jens Jägersküpper. A blend of Markov-chain and drift analysis. In *Parallel Problem Solving From Nature, PPSN 2008*, pages 41–51. Springer, 2008.

- [Jan07] Thomas Jansen. On the brittleness of evolutionary algorithms. In *Foundations of Genetic Algorithms, FOGA 2007*, pages 54–69. Springer, 2007.
- [JBS93] Ari Juels, Shumeet Baluja, and Alistair Sinclair. The equilibrium genetic algorithm and the role of crossover. Unpublished, 1993.
- [Jen04] Mikkel T. Jensen. Helper-objectives: Using multi-objective evolutionary algorithms for single-objective optimisation. *J. Math. Model. Algorithms*, 3(4):323–347, 2004.
- [JJW05] Thomas Jansen, Kenneth A. De Jong, and Ingo Wegener. On the choice of the offspring population size in evolutionary algorithms. *Evolutionary Computation*, 13:413–440, 2005.
- [Joh10] Daniel Johannsen. *Random Combinatorial Structures and Randomized Search Heuristics*. PhD thesis, Universität des Saarlandes, 2010.
- [JOZ13] Thomas Jansen, Pietro Simone Oliveto, and Christine Zarges. Approximating vertex cover using edge-based representations. In *FOGA*, pages 87–96. ACM, 2013.
- [JS07] Jens Jägersküpper and Tobias Storch. When the plus strategy outperforms the comma strategy and when not. In *Foundations of Computational Intelligence, FOCI 2007*, pages 25–32. IEEE, 2007.
- [JW00] Thomas Jansen and Ingo Wegener. On the choice of the mutation probability for the (1+1) EA. In *Parallel Problem Solving from Nature, PPSN 2000*, pages 89–98. Springer, 2000.
- [JW06] Thomas Jansen and Ingo Wegener. On the analysis of a dynamic evolutionary algorithm. *Journal of Discrete Algorithms*, 4:181–199, 2006.
- [JZ12] Thomas Jansen and Christine Zarges. Fixed budget computations: a different perspective on run time analysis. In *GECCO*, pages 1325–1332. ACM, 2012.
- [JZ14a] Thomas Jansen and Christine Zarges. Performance analysis of randomised search heuristics operating with a fixed budget. *Theor. Comput. Sci.*, 545:39–58, 2014.

- [JZ14b] Thomas Jansen and Christine Zarges. Reevaluating immune-inspired hypermutations using the fixed budget perspective. *IEEE Trans. Evol. Comput.*, 18(5):674–688, 2014.
- [KG14] Andreas Krause and Daniel Golovin. Submodular function maximization. In *Tractability*, pages 71–104. Cambridge University Press, 2014.
- [KHE15] Giorgos Karafotias, Mark Hoogendoorn, and Ágoston E. Eiben. Parameter control in evolutionary algorithms: trends and challenges. *IEEE Transactions on Evolutionary Computation*, 19:167–187, 2015.
- [KLNO10a] Timo Kötzing, Per Kristian Lehre, Frank Neumann, and Pietro S. Oliveto. Ant colony optimization and the minimum cut problem. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 1393–1400. ACM, 2010.
- [KLNO10b] Stefan Kratsch, Per Kristian Lehre, Frank Neumann, and Pietro Simone Oliveto. Fixed parameter evolutionary algorithms and maximum leaf spanning trees: A matter of mutation. In *PPSN (1)*, volume 6238 of *Lecture Notes in Computer Science*, pages 204–213. Springer, 2010.
- [KLW15a] Timo Kötzing, Andrei Lissovoi, and Carsten Witt. (1+1) EA on generalized dynamic onemax. In *FOGA*, pages 40–51. ACM, 2015.
- [KLW15b] Timo Kötzing, Andrei Lissovoi, and Carsten Witt. (1+1) EA on generalized dynamic OneMax. In *Foundations of Genetic Algorithms, FOGA 2015*, pages 40–51. ACM, 2015. doi:10.1145/2725494.2725502.
- [KM12a] Timo Kötzing and Hendrik Molter. ACO beats EA on a dynamic pseudo-boolean function. In *PPSN (1)*, volume 7491 of *Lecture Notes in Computer Science*, pages 113–122. Springer, 2012.
- [KM12b] Timo Kötzing and Hendrik Molter. ACO beats EA on a dynamic pseudo-boolean function. In *Parallel Problem Solving from Nature, PPSN 2012, Part I*, pages 113–122. Springer, 2012.

- [KMN99] Samir Khuller, Anna Moss, and Joseph Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, 1999.
- [KN13] Stefan Kratsch and Frank Neumann. Fixed-parameter evolutionary algorithms and the vertex cover problem. *Algorithmica*, 65(4):754–771, 2013.
- [KNRW10] Timo Kötzing, Frank Neumann, Heiko Röglin, and Carsten Witt. Theoretical properties of two ACO approaches for the traveling salesman problem. In *Swarm Intelligence, ANTS 2010*, pages 324–335. Springer, 2010.
- [KP19] Pascal Kerschke and Mike Preuss. Exploratory landscape analysis. In *GECCO (Companion)*, pages 1137–1155. ACM, 2019.
- [KU18] Adrian Kosowski and Przemyslaw Uznanski. Brief announcement: Population protocols are fast. In *Symposium on Principles of Distributed Computing, PODC 2018*, pages 475–477. ACM, 2018.
- [KW20a] Martin Krejca and Carsten Witt. Theory of estimation-of-distribution algorithms. In Benjamin Doerr and Frank Neumann, editors, *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 405–442. Springer, 2020. Also available at <https://arxiv.org/abs/1806.05392>.
- [KW20b] Martin S. Krejca and Carsten Witt. Lower bounds on the run time of the Univariate Marginal Distribution Algorithm on OneMax. *Theoretical Computer Science*, 832:143–165, 2020. doi:10.1016/j.tcs.2018.06.004.
- [Leh10] Per Kristian Lehre. Negative drift in populations. In *Parallel Problem Solving from Nature, PPSN 2010*, pages 244–253. Springer, 2010.
- [Leh11] Per Kristian Lehre. Fitness-levels for non-elitist populations. In *Genetic and Evolutionary Computation Conference, GECCO 2011*, pages 2075–2082. ACM, 2011.
- [Len20] Johannes Lengler. Drift analysis. In Benjamin Doerr and Frank Neumann, editors, *Theory of Evolutionary*

Computation: Recent Developments in Discrete Optimization, pages 89–131. Springer, 2020. Also available at <https://arxiv.org/abs/1712.00964>.

- [Lew08] Matthew R. Lewis. Evolutionary visual art and design. In *The Art of Artificial Evolution*, Natural Computing Series, pages 3–37. Springer, 2008.
- [LL02] Pedro Larrañaga and José Antonio Lozano, editors. *Estimation of Distribution Algorithms*. Genetic Algorithms and Evolutionary Computation. Springer, 2002.
- [LLM07] Fernando G. Lobo, Cláudio F. Lima, and Zbigniew Michalewicz, editors. *Parameter Setting in Evolutionary Algorithms*. Springer, 2007.
- [LMNS09] Jon Lee, Vahab S. Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Non-monotone submodular maximization under matroid and knapsack constraints. In *STOC*, pages 323–332. ACM, 2009.
- [LMS19] Johannes Lengler, Anders Martinsson, and Angelika Steger. When does hillclimbing fail on monotone functions: an entropy compression argument. In *Analytic Algorithmics and Combinatorics, ANALCO 2019*, pages 94–102. SIAM, 2019.
- [LN17] Per Kristian Lehre and Phan Trung Hai Nguyen. Improved runtime bounds for the univariate marginal distribution algorithm via anti-concentration. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 1383–1390. ACM, 2017.
- [LN19a] Per Kristian Lehre and Phan Trung Hai Nguyen. On the limitations of the univariate marginal distribution algorithm to deception and where bivariate EDAs might help. In *Foundations of Genetic Algorithms, FOGA 2019*, pages 154–168. ACM, 2019.
- [LN19b] Per Kristian Lehre and Phan Trung Hai Nguyen. Runtime analysis of the univariate marginal distribution algorithm under low selective pressure and prior noise. In *Genetic and Evolutionary Computation Conference, GECCO 2019*, pages 1497–1505. ACM, 2019.

- [Lob07] Fernando G. Lobo. Lost gems of EC: The equilibrium genetic algorithm and the role of crossover. *SIGEVolution*, 2(2):14–15, 2007. URL: <https://doi.org/10.1145/1329465.1329468>.
- [LS11] Jörg Lässig and Dirk Sudholt. Adaptive population models for offspring populations and parallel evolutionary algorithms. In *Foundations of Genetic Algorithms, FOGA 2011*, pages 181–192. ACM, 2011.
- [LS18] Johannes Lengler and Angelika Steger. Drift analysis and evolutionary algorithms revisited. *Combinatorics, Probability & Computing*, 27:643–666, 2018.
- [LSW21] Johannes Lengler, Dirk Sudholt, and Carsten Witt. The complex parameter landscape of the compact genetic algorithm. *Algorithmica*, 83:1096–1137, 2021.
- [LW15] Andrei Lissovoi and Carsten Witt. Runtime analysis of ant colony optimization on dynamic shortest path problems. *Theor. Comput. Sci.*, 561:73–85, 2015.
- [LW16] Andrei Lissovoi and Carsten Witt. MMAS versus population-based EA on a family of dynamic fitness functions. *Algorithmica*, 75:554–576, 2016.
- [LW18] Andrei Lissovoi and Carsten Witt. The impact of a sparse migration topology on the runtime of island models in dynamic optimization. *Algorithmica*, 80(5):1634–1657, 2018.
- [LY12] Per Kristian Lehre and Xin Yao. On the impact of mutation-selection balance on the runtime of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 16:225–241, 2012.
- [MBT⁺13] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Markus Wagner, Jakob Bossek, and Frank Neumann. A novel feature-based approach to characterize algorithm performance for the traveling salesperson problem. *Ann. Math. Artif. Intell.*, 69(2):151–182, 2013.
- [MD10] Christie Myburgh and Kalyanmoy Deb. Evolutionary algorithms in large-scale open pit mine scheduling. In *GECCO*, pages 1155–1162. ACM, 2010.

- [MGRD05] Roberto Montemanni, Luca Maria Gambardella, Andrea Emilio Rizzoli, and Alberto V. Donati. Ant colony system for a dynamic vehicle routing problem. *J. Comb. Optim.*, 10(4):327–343, 2005.
- [MKF⁺19] Marko Mitrovic, Ehsan Kazemi, Moran Feldman, Andreas Krause, and Amin Karbasi. Adaptive sequence submodularity. In *NeurIPS*, pages 5353–5364, 2019.
- [MLY17] Michalis Mavrovouniotis, Changhe Li, and Shengxiang Yang. A survey of swarm intelligence for dynamic optimization: Algorithms and applications. *Swarm Evol. Comput.*, 33:1–17, 2017.
- [Mon20] Morteza Monemizadeh. Dynamic submodular maximization. In *NeurIPS*, 2020.
- [MP96] Heinz Mühlenbein and Gerhard Paass. From recombination of genes to the estimation of distributions I. Binary parameters. In *Parallel Problem Solving from Nature, PPSN 1996*, pages 178–187. Springer, 1996.
- [MRC09] Boris Mitavskiy, Jonathan E. Rowe, and Chris Cannings. Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *International Journal on Intelligent Computing and Cybernetics*, 2:243–284, 2009.
- [MS15] Andrea Mambrini and Dirk Sudholt. Design and analysis of schemes for adapting migration intervals in parallel evolutionary algorithms. *Evolutionary Computation*, 23:559–582, 2015.
- [Müh92] Heinz Mühlenbein. How genetic algorithms really work: mutation and hillclimbing. In *Parallel Problem Solving from Nature, PPSN 1992*, pages 15–26. Elsevier, 1992.
- [NAN20] Aneta Neumann, Bradley Alexander, and Frank Neumann. Evolutionary image transition and painting using random walks. *Evol. Comput.*, 28(4):643–675, 2020.
- [NAW20] Mehdi Neshat, Bradley Alexander, and Markus Wagner. A hybrid cooperative co-evolution algorithm framework for optimising power take off and placements of wave energy converters. *Inf. Sci.*, 534:218–244, 2020.

- [NN20] Aneta Neumann and Frank Neumann. Optimising monotone chance-constrained submodular functions using evolutionary multi-objective algorithms. In *PPSN (1)*, volume 12269 of *Lecture Notes in Computer Science*, pages 404–417. Springer, 2020.
- [NOW09] Frank Neumann, Pietro S. Oliveto, and Carsten Witt. Theoretical analysis of fitness-proportional selection: landscapes and efficiency. In *Genetic and Evolutionary Computation Conference, GECCO 2009*, pages 835–842. ACM, 2009.
- [NPW19] Frank Neumann, Mojgan Pourhassan, and Carsten Witt. Improved runtime results for simple randomised search heuristics on linear functions with a uniform constraint. In *GECCO*, pages 1506–1514. ACM, 2019.
- [NS19] Frank Neumann and Andrew M. Sutton. Runtime analysis of the $(1 + 1)$ evolutionary algorithm for the chance-constrained knapsack problem. In *FOGA*, pages 147–153. ACM, 2019.
- [NSN13a] Samadhi Nallaperuma, Andrew M. Sutton, and Frank Neumann. Fixed-parameter evolutionary algorithms for the euclidean traveling salesperson problem. In *IEEE Congress on Evolutionary Computation*, pages 2037–2044. IEEE, 2013.
- [NSN13b] Samadhi Nallaperuma, Andrew M. Sutton, and Frank Neumann. Parameterized complexity analysis and more effective construction methods for ACO algorithms and the euclidean traveling salesperson problem. In *IEEE Congress on Evolutionary Computation*, pages 2045–2052. IEEE, 2013.
- [NSW09] Frank Neumann, Dirk Sudholt, and Carsten Witt. Analysis of different MMAS ACO algorithms on unimodal functions and plateaus. *Swarm Intelligence*, 3:35–68, 2009.
- [NSW10] Frank Neumann, Dirk Sudholt, and Carsten Witt. A few ants are enough: ACO with iteration-best update. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 63–70. ACM, 2010.
- [NW06] Frank Neumann and Ingo Wegener. Minimum spanning trees made easier via multi-objective optimization. *Nat. Comput.*, 5(3):305–319, 2006.

- [NW07] Frank Neumann and Ingo Wegener. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Theoretical Computer Science*, 378:32–40, 2007.
- [NW09] Frank Neumann and Carsten Witt. Runtime analysis of a simple ant colony optimization algorithm. *Algorithmica*, 54:243–255, 2009.
- [NW10a] Frank Neumann and Carsten Witt. Ant colony optimization and the minimum spanning tree problem. *Theoretical Computer Science*, 411:2406–2413, 2010.
- [NW10b] Frank Neumann and Carsten Witt. *Bioinspired Computation in Combinatorial Optimization – Algorithms and Their Computational Complexity*. Springer, 2010.
- [NW15] Frank Neumann and Carsten Witt. On the runtime of randomized local search and simple evolutionary algorithms for dynamic makespan scheduling. In *IJCAI*, pages 3742–3748. AAAI Press, 2015.
- [NWF78] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.*, 14(1):265–294, 1978.
- [NWN15] Samadhi Nallaperuma, Markus Wagner, and Frank Neumann. Analyzing the effects of instance features and algorithm parameters for max-min ant system and the traveling salesperson problem. *Frontiers Robotics AI*, 2:18, 2015.
- [NY12] T.T. Nguyen and X. Yao. Continuous dynamic constrained optimization: The challenges. *IEEE Transactions on Evolutionary Computation*, 16(6):769–786, 2012. doi:10.1109/TEVC.2011.2180533.
- [NYB12] Trung Thanh Nguyen, Shengxiang Yang, and Jürgen Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm Evol. Comput.*, 6:1–24, 2012.
- [Och02] Gabriela Ochoa. Setting the mutation rate: scope and limitations of the $1/L$ heuristic. In *Genetic and Evolutionary Computation Conference, GECCO 2002*, pages 495–502. Morgan Kaufmann, 2002.

- [OE12] Adrian Ogierman and Robert Elsässer. The impact of the power law exponent on the behavior of a dynamic epidemic type process. In *Symposium on Parallelism in Algorithms and Architectures, SPAA 2012*, pages 131–139. ACM, 2012.
- [OV18] Gabriela Ochoa and Nadarajen Veerapen. Mapping the global structure of TSP fitness landscapes. *J. Heuristics*, 24(3):265–294, 2018.
- [OW11] Pietro S. Oliveto and Carsten Witt. Simplified drift analysis for proving lower bounds in evolutionary computation. *Algorithmica*, 59:369–386, 2011.
- [OW12] Pietro S. Oliveto and Carsten Witt. Erratum: Simplified drift analysis for proving lower bounds in evolutionary computation. *CoRR*, abs/1211.7184, 2012. [arXiv:1211.7184](https://arxiv.org/abs/1211.7184).
- [OW15] Pietro S. Oliveto and Carsten Witt. Improved time complexity analysis of the simple genetic algorithm. *Theoretical Computer Science*, 605:21–41, 2015.
- [OWBM13] Yuki Osada, R. Lyndon While, Luigi Barone, and Zbigniew Michalewicz. Multi-mine planning using a multi-objective evolutionary algorithm. In *IEEE Congress on Evolutionary Computation*, pages 2902–2909. IEEE, 2013.
- [PA12] Erik Pitzer and Michael Affenzeller. A comprehensive survey on fitness landscape analysis. In *Recent Advances in Intelligent Engineering Systems*, volume 378 of *Studies in Computational Intelligence*, pages 161–191. Springer, 2012.
- [PGN15] Mojgan Pourhassan, Wanru Gao, and Frank Neumann. Maintaining 2-approximations for the dynamic vertex cover problem using evolutionary algorithms. In *GECCO*, pages 903–910. ACM, 2015.
- [PHL15] Martin Pelikan, Mark Hauschild, and Fernando G. Lobo. Estimation of distribution algorithms. In Janusz Kacprzyk and Witold Pedrycz, editors, *Springer Handbook of Computational Intelligence*, pages 899–928. Springer, 2015.
- [PRN20] Mojgan Pourhassan, Vahid Roostapour, and Frank Neumann. Runtime analysis of RLS and (1+1) EA for the dynamic

- weighted vertex cover problem. *Theor. Comput. Sci.*, 832:20–41, 2020.
- [Prü04] Adam Prügel-Bennett. When a genetic algorithm outperforms hill-climbing. *Theoretical Computer Science*, 320:135–153, 2004.
- [PSN19] Mojgan Pourhassan, Feng Shi, and Frank Neumann. Parameterized analysis of multiobjective evolutionary algorithms and the weighted vertex cover problem. *Evol. Comput.*, 27(4):559–575, 2019.
- [QBJT19] Chao Qian, Chao Bian, Wu Jiang, and Ke Tang. Running time analysis of the $(1 + 1)$ -EA for OneMax and LeadingOnes under bit-wise noise. *Algorithmica*, 81:749–795, 2019.
- [QSYT17] Chao Qian, Jing-Cheng Shi, Yang Yu, and Ke Tang. On subset selection with general cost constraints. In Carles Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 2613–2619. ijcai.org, 2017. doi:10.24963/ijcai.2017/364.
- [QYT⁺19] Chao Qian, Yang Yu, Ke Tang, Xin Yao, and Zhi-Hua Zhou. Maximizing submodular or monotone approximately submodular functions by multi-objective evolutionary algorithms. *Artif. Intell.*, 275:279–294, 2019.
- [QYZ15] Chao Qian, Yang Yu, and Zhi-Hua Zhou. Subset selection by pareto optimization. In *NIPS*, pages 1774–1782, 2015.
- [Rec73] Ingo Rechenberg. *Evolutionstrategie*. Friedrich Fromman Verlag (Günther Holzboog KG), Stuttgart, 1973.
- [RKD17] Pratyusha Rakshit, Amit Konar, and Swagatam Das. Noisy evolutionary optimization algorithms - A comprehensive survey. *Swarm and Evolutionary Computation*, 33:18–45, 2017. doi:10.1016/j.swevo.2016.09.002.
- [RLY09] Philipp Rohlfshagen, Per Kristian Lehre, and Xin Yao. Dynamic evolutionary optimisation: an analysis of frequency and magnitude of change. In *GECCO*, pages 1713–1720. ACM, 2009.

- [RNN18] Vahid Roostapour, Aneta Neumann, and Frank Neumann. On the performance of baseline evolutionary algorithms on the dynamic knapsack problem. In *PPSN (1)*, volume 11101 of *Lecture Notes in Computer Science*, pages 158–169. Springer, 2018.
- [RNN20] Vahid Roostapour, Aneta Neumann, and Frank Neumann. Evolutionary multi-objective optimization for the dynamic knapsack problem. *CoRR*, abs/2004.12574, 2020. Conference version appeared at PPSN 2018.
- [RNNF18] Vahid Roostapour, Aneta Neumann, Frank Neumann, and Tobias Friedrich. Pareto optimization for subset selection with dynamic cost constraints. *CoRR*, abs/1811.07806, 2018.
- [RNNF19] Vahid Roostapour, Aneta Neumann, Frank Neumann, and Tobias Friedrich. Pareto optimization for subset selection with dynamic cost constraints. In *AAAI*, pages 2354–2361. AAAI Press, 2019.
- [RNRN21] William Reid, Aneta Neumann, Simon Ratcliffe, and Frank Neumann. Advanced ore mine optimisation under uncertainty using evolution. *CoRR*, abs/2102.05235, 2021.
- [Row18] Jonathan E. Rowe. Linear multi-objective drift analysis. *Theoretical Computer Science*, 736:25–40, 2018.
- [RP13] Ulrike Ritzinger and Jakob Puchinger. Hybrid metaheuristics for dynamic and stochastic vehicle routing. In *Hybrid Metaheuristics*, volume 434 of *Studies in Computational Intelligence*, pages 77–95. Springer, 2013.
- [RPN18] Vahid Roostapour, Mojgan Pourhassan, and Frank Neumann. Analysis of evolutionary algorithms in dynamic and stochastic environments. *CoRR*, abs/1806.08547, 2018.
- [RS14] Jonathan E. Rowe and Dirk Sudholt. The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theoretical Computer Science*, 545:20–38, 2014.
- [Rud97] Günter Rudolph. *Convergence properties of evolutionary algorithms*. Kovac, 1997.
- [RW20] Amirhossein Rajabi and Carsten Witt. Self-adjusting evolutionary algorithms for multimodal optimization. In *Genetic and*

Evolutionary Computation Conference, GECCO 2020, pages 1314–1322. ACM, 2020.

- [SD02] Thomas Stützle and Marco Dorigo. A short convergence proof for a class of ant colony optimization algorithms. *IEEE Transactions on Evolutionary Computation*, 6:358–365, 2002.
- [SFAM18] Moshe Sipper, Weixuan Fu, Karuna Ahuja, and Jason H. Moore. Investigating the parameter space of evolutionary algorithms. *BioData Min.*, 11(1):2:1–2:14, 2018.
- [Sha02] Jonathan L. Shapiro. The sensitivity of PBIL to its learning rate, and how detailed balance can remove it. In *Foundations of Genetic Algorithms, FOGA 2002*, pages 115–132. Morgan Kaufmann, 2002.
- [Sha05] Jonathan L. Shapiro. Drift and scaling in estimation of distribution algorithms. *Evolutionary Computing*, 13:99–123, 2005.
- [Sha06] Jonathan L. Shapiro. Diversity loss in general estimation of distribution algorithms. In *Parallel Problem Solving from Nature, PPSN 2006*, pages 92–101. Springer, 2006.
- [SN12] Andrew M. Sutton and Frank Neumann. A parameterized runtime analysis of evolutionary algorithms for the Euclidean traveling salesperson problem. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI’12)*, pages 1105–1111. AAAI Press, 2012.
- [SNN14] Andrew M. Sutton, Frank Neumann, and Samadhi Nallaperuma. Parameterized runtime analyses of evolutionary algorithms for the planar euclidean traveling salesperson problem. *Evol. Comput.*, 22(4):595–628, 2014.
- [SSF⁺19] Feng Shi, Martin Schirneck, Tobias Friedrich, Timo Kötzing, and Frank Neumann. Reoptimization time analysis of evolutionary algorithms on linear functions under dynamic uniform constraints. *Algorithmica*, 81(2):828–857, 2019.
- [ST05] Giovanni Sebastiani and Giovanni Luca Torrisi. An extended ant colony algorithm and its convergence analysis. *Methodology and Computing in Applied Probability*, 7:249–263, 2005.

- [ST12] Dirk Sudholt and Christian Thyssen. A simple ant colony optimizer for stochastic shortest path problems. *Algorithmica*, 64:643–672, 2012.
- [Sto06] Tobias Storch. How randomized search heuristics find maximum cliques in planar graphs. In *GECCO*, pages 567–574. ACM, 2006.
- [Sto07] Tobias Storch. Finding large cliques in sparse semi-random graphs by simple randomized search heuristics. *Theor. Comput. Sci.*, 386(1-2):114–131, 2007. doi:10.1016/j.tcs.2007.06.008.
- [STW04] Jens Scharnow, Karsten Tinnefeld, and Ingo Wegener. The analysis of evolutionary algorithms on sorting and shortest paths problems. *Journal of Mathematical Modelling and Algorithms*, 3:349–366, 2004.
- [Sud13] Dirk Sudholt. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 17:418–435, 2013.
- [Sud18] Dirk Sudholt. On the robustness of evolutionary algorithms to noise: refined results and an example where noise helps. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1523–1530. ACM, 2018.
- [SvH11] Kate Smith-Miles and Jano I. van Hemert. Discovering the suitability of optimisation algorithms by learning from evolved instances. *Ann. Math. Artif. Intell.*, 61(2):87–104, 2011.
- [SW19] Dirk Sudholt and Carsten Witt. On the choice of the update strength in estimation-of-distribution algorithms and ant colony optimization. *Algorithmica*, 81:1450–1489, 2019.
- [Sys93] Gilbert Syswerda. Simulated crossover in genetic algorithms. In *Foundations of Genetic Algorithms, FOGA 1992*, pages 239–255. Morgan Kaufmann, 1993.
- [The09] Madeleine Theile. Exact solutions to the traveling salesperson problem by a population-based evolutionary algorithm. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2009*, pages 145–155. Springer, 2009.

- [TWD⁺13] Raymond Tran, Junhua Wu, Christopher Denison, Thomas Ackling, Markus Wagner, and Frank Neumann. Fast and effective multi-objective optimisation of wind turbine placement. In *GECCO*, pages 1381–1388. ACM, 2013.
- [Von10] Jan Vondrák. Submodularity and curvature: The optimal algorithm. *RIMS Kôkyûroku Bessatsu*, B23:253—266, 2010.
- [Weg01] Ingo Wegener. Theoretical aspects of evolutionary algorithms. In *Automata, Languages and Programming, ICALP 2001*, pages 64–78. Springer, 2001.
- [Wit06] Carsten Witt. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation*, 14:65–86, 2006.
- [Wit13] Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing*, 22:294–318, 2013.
- [Wit19] Carsten Witt. Upper bounds on the running time of the univariate marginal distribution algorithm on OneMax. *Algorithmica*, 81:632–667, 2019.
- [WQT18] Mengxi Wu, Chao Qian, and Ke Tang. Dynamic mutation based Pareto optimization for subset selection. In *Intelligent Computing Methodologies, ICIC 2018, Part III*, pages 25–35. Springer, 2018.
- [WSOC14] L. Darrell Whitley, Andrew M. Sutton, Gabriela Ochoa, and Francisco Chicano. The component model for elementary landscapes and partial neighborhoods. *Theor. Comput. Sci.*, 545:59–75, 2014.
- [XHA⁺19] Yue Xie, Oscar Harper, Hirađ Assimi, Aneta Neumann, and Frank Neumann. Evolutionary algorithms for the chance-constrained knapsack problem. In *GECCO*, pages 338–346. ACM, 2019.
- [XNN20] Yue Xie, Aneta Neumann, and Frank Neumann. Specific single- and multi-objective evolutionary algorithms for the chance-constrained knapsack problem. In *GECCO*, pages 271–279. ACM, 2020.

- [XNN21] Yue Xie, Aneta Neumann, and Frank Neumann. Heuristic strategies for solving complex interacting stockpile blending problem with chance constraints. In *GECCO*, pages 1079–1087. ACM, 2021.
- [XNNS21] Yue Xie, Aneta Neumann, Frank Neumann, and Andrew M. Sutton. Runtime analysis of RLS and the (1+1) EA for the chance-constrained knapsack problem with correlated uniform weights. In *GECCO*, pages 1187–1194. ACM, 2021.
- [Zho09] Yuren Zhou. Runtime analysis of an ant colony optimization algorithm for TSP instances. *IEEE Transactions on Evolutionary Computation*, 13:1083–1092, 2009.
- [ZV16] Haifeng Zhang and Yevgeniy Vorobeychik. Submodular optimization with routing constraints. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI 2016*, pages 819–826. AAAI Press, 2016.
- [ZYD18] Weijie Zheng, Guangwen Yang, and Benjamin Doerr. Working principles of binary differential evolution. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1103–1110. ACM, 2018.
- [ZYQ19] Zhi-Hua Zhou, Yang Yu, and Chao Qian. *Evolutionary Learning: Advances in Theories and Algorithms*. Springer, 2019. doi:10.1007/978-981-13-5956-9.