



HAL
open science

From Understanding Genetic Drift to a Smart-Restart Mechanism for Estimation-of-Distribution Algorithms *

Benjamin Doerr, Weijie Zheng

► **To cite this version:**

Benjamin Doerr, Weijie Zheng. From Understanding Genetic Drift to a Smart-Restart Mechanism for Estimation-of-Distribution Algorithms *. *Journal of Machine Learning Research*, 2023, 24. hal-04483214

HAL Id: hal-04483214

<https://hal.science/hal-04483214v1>

Submitted on 29 Feb 2024

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.

From Understanding Genetic Drift to a Smart-Restart Mechanism for Estimation-of-Distribution Algorithms*

Weijie Zheng

ZHENGWEIJIE@HIT.EDU.CN

*School of Computer Science and Technology
International Research Institute for Artificial Intelligence
Harbin Institute of Technology, Shenzhen, 518055, China*

Benjamin Doerr[†]

DOERR@LIX.POLYTECHNIQUE.FR

*Laboratoire d'Informatique (LIX)
École Polytechnique, CNRS
Institut Polytechnique de Paris, Palaiseau, 91120, France*

Editor: Sathiya Keerthi

Abstract

Estimation-of-distribution algorithms (EDAs) are optimization algorithms that learn a distribution from which good solutions can be sampled easily. A key parameter of most EDAs is the sample size (population size). Too small values lead to the undesired effect of genetic drift, while larger values slow down the process.

Building on a quantitative analysis of how the population size leads to genetic drift, we design a smart-restart mechanism for EDAs. By stopping runs when the risk for genetic drift is high, it automatically runs the EDA in good parameter regimes.

Via a mathematical runtime analysis, we prove a general performance guarantee for this smart-restart scheme. For many situations where the optimal parameter values are known, this shows that the restart scheme automatically finds these optimal values, leading to the asymptotically optimal performance.

We also conduct an extensive experimental analysis. On four classic benchmarks, the smart-restart scheme leads to a performance close to the one obtainable with optimal parameter values. We also conduct experiments with PBIL (cross-entropy algorithm) on the max-cut problem and the bipartition problem. Again, the smart-restart mechanism finds much better values for the population size than those suggested in the literature, leading to a much better performance.

Keywords: Estimation-of-distribution algorithms, genetic drift, restart strategy, randomized search heuristics, theory of computing

1. Introduction

Different from solution-oriented optimization heuristics such as local search, simulated annealing, or genetic algorithms, estimation-of-distribution algorithms (EDAs) (Larrañaga and Lozano, 2002; Pelikan et al., 2015) try to learn a probability distribution on the search space (“probabilistic model of the search space”) that allows to sample good solutions.

*. A preliminary version (Doerr and Zheng, 2020a), prepared while the first author was with Southern University of Science and Technology, was published in the proceedings of GECCO 2020.

†. Corresponding author.

EDAs are iterative in nature, that is, in each iteration they sample a certain number (“population size”) of solutions, evaluate their quality, and update the previous model based on these solutions and their quality.

The population size is crucial for the optimization behavior and the performance of the EDA. Clearly, a large population size increases the cost of a single iteration. However, a small population size means that the model update relies only on a few samples and thus is heavily influenced by the random nature of the samples. The effect that model updates are influenced more by the randomness in the sampling process than by the guidance of the objective function is known as *genetic drift*.

Genetic drift can be detrimental to the performance of an EDA. As an example, let us discuss the performance of the univariate marginal distribution algorithm (UMDA) (Mühlenbein and Paass, 1996) with artificial frequency margins $\{1/n, 1 - 1/n\}$ on the DECEPTIVE-LEADINGBLOCKS problem with problem size n . Lehre and Nguyen (2019, Theorem 4.9) have shown that if the population size is small, more precisely, $\lambda = \Omega(\log n) \cap o(n)$, and the selective pressure is standard ($\mu/\lambda \geq 14/1000$), then the expected runtime is at least exponential in λ . In contrast, if the population size is large enough, that is, $\lambda = \Omega(n \log n)$ and again $\mu = \Theta(\lambda)$, then with high probability the UMDA finds the optimum in $\lambda(n/2 + 2e \ln n)$ function evaluations (Doerr and Krejca, 2021a, Theorem 5). This runtime bound is roughly proportional to the population size λ , indicating (no lower bounds were shown in (Doerr and Krejca, 2021a)) that the optimal population size is just above the regime leading to the detrimental behavior observed in (Lehre and Nguyen, 2019), but that further increases of the population size are again costly.

The essential reason for this performance pattern, quantified precisely by Doerr and Zheng (2020b), but known already since the ground-breaking works of Shapiro (2002a, 2005, 2006), is that small population sizes can lead to strong genetic drift, that is, the random fluctuations of the sampling frequencies caused by the randomness in the sampling of search points eventually move some sampling frequencies towards a boundary of the frequency range that is not justified by the fitness.

We refer to the recent survey of Krejca and Witt (2020a) for a detailed discussion of the known runtime results for EDAs and only give a high-level summary here. For most of the results presented there, a minimum population size is necessary and then the runtime is roughly proportional to the population size. This suggests again that for small population sizes, no good performance guarantees could be proven (because of the genetic drift effect), whereas from a certain population size on this effect disappears and the runtime becomes roughly proportional to the population size (stemming from the fact that the cost of one iteration is proportional to the population size). In the presence of such a runtime behavior, naturally, choosing the appropriate population size is a key challenge in the effective usage of EDAs.

We note that genetic drift does not in absolutely all cases lead to a bad performance. For example, independently in (Dang et al., 2019; Witt, 2019) it was shown that the UMDA optimizes ONEMAX in time $O(n \log n)$ also for logarithmic population sizes, which are clearly in the regime with strong genetic drift. In (Dang et al., 2019), it was also shown that the runtime of the UMDA on LEADINGONES is at most quadratic when $\lambda = \Omega(\log n) \cap O(n/\log n)$, which is again in the regime with strong genetic drift. So these results show that a good performance is also possible in the presence of strong genetic drift. We note,

however, that the same runtimes of $O(n \log n)$ and $O(n^2)$ can also be obtained in the regime with low genetic drift, see again (Dang et al., 2019; Witt, 2019), and we note further that no example is known where an EDA has an asymptotically better performance in the strong genetic drift regime than outside of it. We finally note that the very careful experimental analysis in (Lengler et al., 2021) shows a good runtime of the cGA on ONEMAX both in a range with strong genetic drift and in a range with low genetic drift (but not in between), but the runtimes observed in the former are still higher than in the latter. The experiments in (Dang et al., 2019) consider only two values for the population size μ , namely \sqrt{n} with strong genetic drift and $\sqrt{n} \log n$ with low genetic drift, of the UMDA optimizing ONEMAX and BINVAL. For both values good runtimes are observed, slightly better ones for the value in the genetic drift regime. However, since only two population sizes are implemented, it is not clear if really the strong genetic drift regime leads to better runtimes or if the best runtimes are observed for a value in the regime with low genetic drift, but different from the one regarded in (Dang et al., 2019). In the light of all these results, and also our experimental results in Section 5, it appears very justified to generally prefer running EDAs in the regime with low genetic drift.

Given the observation that genetic drift often leads to unfavorable results, there have been attempts to define EDAs that are not prone to genetic drift (Shapiro, 2002b; Branke et al., 2007; Friedrich et al., 2016; Doerr and Krejca, 2020). While these led to some promising results, due to their restricted evaluation (only on ONEMAX, BINVAL, LEADING-ONES, NEEDLE, and NK-landscape) and in the light of the negative result (Doerr and Krejca, 2020, Theorem 4), in this work we prefer to discuss how to set the parameters for established EDAs in a way that they do not suffer from genetic drift.

Setting the parameters of an optimization heuristic right is a known challenge. The most direct way is to try to understand how the parameter influences the performance of the algorithm on a given problem and then set the parameter accordingly. When done via experimental means, this approach can be time-consuming, and usually only gives information for a particular problem of a particular size. Mathematical approaches can determine optimal parameter values over larger classes of instances and problems sizes (see, e.g. (Witt, 2013) for such a result), but they require a deep expertise and usually can only be applied to simple benchmark problems.

An easier way to approach the parameter tuning problem is to find ways to automatically set the (or some) parameters. This can be done on-the-fly, that is, the algorithm tries to learn what are good parameter values and adjusts the parameters accordingly while running, or via separate runs of the algorithm with different parameter values. From our understanding of genetic drift, we do not see how an on-the-fly parameter choice of the population size can be successful for an EDA. On the one hand, it is difficult to see during the run of the algorithm whether a model update is justified by the fitness or rather caused by unlucky samplings of the individuals. On the other hand, once the EDA has suffered from genetic drift, it is not clear how to repair the probabilistic model.

For this reason, we shall concentrate on approaches that use several runs of the EDA with different parameter values, and this in a way that the algorithm user does not need to

take care of this parameter in any way. For EDAs having a single parameter such as the compact genetic algorithm, this will result in a parameter-less EDA.¹

Harik and Lobo (1999) proposed two strategies to remove the population size of crossover-based genetic algorithms. One basic strategy is doubling the population size and restarting when all individuals’ genotypes have become identical. The drawback of this strategy is the long time it takes to fulfill this termination criterion after genetic drift has become detrimental. Harik and Lobo proposed a second strategy in which multiple populations with different sizes run simultaneously, smaller population sizes may use more function evaluations, but are removed once their fitness value falls behind the one of larger populations. Their experimental results showed that their genetic algorithm with this second strategy only had a small performance loss over the same genetic algorithm with optimal parameter settings. Many extensions of this strategy and applications to other optimization algorithms (including EDAs) have followed. These gave rise to the extended compact genetic algorithm (Lima and Lobo, 2004), the hierarchical Bayesian optimization algorithm (Pelikan and Lin, 2004), and many other algorithms.

In the IPOP-CMA-ES, Auger and Hansen (2005) use another strategy to remove the population size as a parameter to be set by the algorithm user. They restart the kernel algorithm, the (μ_W, λ) -CMA-ES, with twice the population size (and the other parameters unchanged) once one of five predefined criteria is reached. Four of these criteria build the covariance matrix or evolution paths and thus are specific to the CMA-ES. The other criterion only depends on the objective function and thus can be used also with other algorithms (this is what we shall do in Section 6.2). This criterion triggers a restart if among the last $10 + \lceil 30n/\lambda \rceil$ (where n is the problem size and λ is the population size) generations, the range of the best objective function values is zero, or the range of these values together with all function values of the current generation is below a predefined threshold.

Goldman and Punch (2014) proposed the parameter-less population pyramid, called P3, to iteratively construct a collection of populations. In P3, the population in the pyramid expands iteratively by first adding a currently not existing solution obtained by some local search strategy into the lowest population, and then utilizing some model-building methods to expand the population in all hierarchies of the pyramid. Since initially no population exists in the pyramid, this algorithm frees the practitioner from specifying a population size.

In (Doerr, 2021, Section 2.4) a strategy was proposed that builds on parallel runs of EDAs with exponentially growing population sizes. With a suitable strategy to assign computational resources, this strategy needs no criterion when to abort a run and still leads to a runtime which is only by a logarithmic factor above the runtime stemming from the optimal (problem- and algorithm-specific) population size of the EDA.

Our contribution: The above parameter-less strategies can be used to automatically find good population sizes, but they are all not specific to the problem of preventing genetic drift. In this work, we aim at profiting from our understanding of genetic drift, in particular, from the recent mathematical analysis (Doerr and Zheng, 2020b) which quantifies when

1. Not surprisingly, many mechanisms to remove parameters have themselves some parameters. The name *parameter-less* might still be justified when these hyperparameters have a less critical influence on the performance of the algorithm.

genetic drift can arise. In very simple words, taking the compact genetic algorithm (cGA), the most simple univariate EDA, as an example, this result shows that genetic drift has a significant influence on the sampling frequency of a bit when the number of iterations exceeds $4\mu^2$ (Doerr and Zheng, 2020b, Proof of Theorem 6), where μ is the hypothetical population size of the cGA, the only parameter of this algorithm, which plays the same role as the population size in other EDAs. We can use this insight to design the following smart-restart version of the cGA. Our *smart-restart mechanism*² starts with running the cGA with hypothetical population size $\mu = 2$, the smallest possible value for this parameter. After $4\mu^2$ iterations, this run is aborted and a new run is started with a randomly initialized probabilistic model, however with twice the hypothetical population size. Each new run is aborted when the time limit of $4\mu^2$ iterations is reached and a new run with twice the parameter value is started. With this procedure, the cGA always runs in a regime in which the risk for genetic drift is considered low. The doubling scheme of the parameter value not only ensures that the possibly more effective smaller values are used first, but also ensures that their influence on the total runtime is small in the case where only a large value is successful. Like any EDA, this is an anytime algorithm, so it can be stopped at any time and then the best solution seen so far is returned.

In more detail and generality, the quantitative analysis in (Doerr and Zheng, 2020b) showed that for each of the three main univariate EDAs cGA, UMDA, and PBIL, both with frequency boundaries and without, there is a number C such that the probability that a particular sampling frequency of the EDA running with (hypothetical) population size μ within the first t function evaluations is strongly affected by genetic drift, is at most $2 \exp(-C\mu^2/t)$. This number C depends on the EDA and on the value of its other parameters (in a manner made precise in (Doerr and Zheng, 2020b)). Hence in all cases, genetic drift affects a sampling frequency after $\Theta(\mu^2)$ function evaluations, however, the implicit constant may depend on the precise setting. For this reason, we shall formulate our general smart-restart scheme, applicable to all these EDAs, with a hyperparameter b , called *budget factor*, such that a run with parameter value μ is aborted after $b\mu^2$ function evaluations. In our asymptotic analyses, we allow that b takes sub-constant values. The main motivation for this is that when taking $b = \Theta(1/\log n)$, we can apply a union-bound argument to show that none of the n sampling frequencies is strongly affected by genetic drift. For the increase of μ in the next stage of the algorithm, for reasons of generality, we not only consider doubling μ , but multiplying it by some number $U > 1$, called update factor. We note that we have just introduced two hyperparameters for a mechanism that controls one algorithm parameter. However, as we shall see in our analyses, both theoretical and experimental, both hyperparameters are not critical for the performance of the smart-restart scheme. Taking $b = 1/\ln(n)$ and $U = 2$ gave the best asymptotic runtimes in our theoretical result and uniformly gave good results in our experiments.

As said before, our smart-restart mechanism can be combined with any of the three main univariate EDAs, and in fact, with any heuristic \mathcal{A} having a parameter μ such one can speculate that the runtime from a certain (unknown) value $\tilde{\mu}$ on is roughly linear in μ . For this general setting, we prove the following *mathematical runtime guarantee*. We assume that there are numbers $\tilde{\mu}$ and T such that \mathcal{A} with all parameter value $\mu \geq \tilde{\mu}$ solves

2. The authors are thankful to an anonymous reviewer of (Doerr and Zheng, 2020a) for suggesting this name.

the given problem in time μT with probability $p > 1 - \frac{1}{U^2}$. Such a runtime behavior is very often observed in EDAs, see, e.g., (Krejca and Witt, 2020a). We prove that under this assumption, our smart-restart mechanism with update factor U and budget factor b solves the problem in expected time

$$\left(\frac{U^2}{U^2 - 1} + \frac{(1-p)U^2}{1 - (1-p)U^2} \right) \max \left\{ b\tilde{\mu}^2, \frac{T^2}{b} \right\} + \frac{pU}{1 - (1-p)U} \tilde{\mu}T,$$

which is $O(\max\{b\tilde{\mu}^2, \frac{T^2}{b}, \tilde{\mu}T\})$ when treating U and p as constants.

When combining this result with several known runtime guarantees for classic EDAs (we refer to Section 4.2 for the details), we easily derive that the smart-restart scheme with $b = \Theta(1/\log n)$ in these situations has the same asymptotic runtime as the original EDA with optimal (problem-specific and often non-trivial to find) value for μ . This in particular holds for the analysis of the cGA on noisy ONEMAX functions, where the optimal value for μ depends also on the intensity of the noise, hence the parameter tuning problem is further complicated by the fact that usually the noise intensity is not known.

We then conduct an extensive *experimental analysis*. We mostly concentrate on the cGA. This algorithm has a single parameter only and this is directly controlling the strength of the model update, so it appears best to study in isolation how the model update strength and automated searches for its best value influence the performance of an EDA.

Since a good experimental understanding of how the model update strength influences the runtime does not yet exist, and since it aids in interpreting our results for the parameter-less versions of the cGA, we also conduct experiments for the original cGA with different static parameter values. We thus ran the original cGA as well as the (parameter-less) parallel-run cGA from (Doerr, 2021) and our (parameter-less) smart-restart cGA on the benchmarks ONEMAX, LEADINGONES, JUMP, and DECEPTIVELEADINGBLOCKS, both in the absence of noise and with additive centered Gaussian posterior noise as considered in (Friedrich et al., 2017). For the original cGA and the noiseless scenario, this analysis confirms, for the first time experimentally for most of these benchmarks, that small population sizes can be detrimental and that from a certain population size on, a roughly linear increase of the runtime can be observed. It also confirms experimentally the insight of the (asymptotic) theory result (Doerr, 2021) that, with the right population size, the cGA can be very efficient on JUMP functions. For example, we measure a median runtime of $4 \cdot 10^6$ on the JUMP function with $n = 50$ and $k = 10$, parameters for which, e.g., the classic $(1+1)$ EA would take more than 10^{17} iterations (the $(1+1)$ EA is the most basic evolutionary algorithm, using a population size of one, creating one offspring from the single parent via standard bit-wise mutation, and replacing the parent by the offspring if the offspring is at least as good as the parent). In the noisy settings, we observe that the population sizes suggested (for ONEMAX) by the theoretical analysis (Friedrich et al., 2017) are much higher (roughly by a factor of 1,000) than what is really necessary, leading to runtime increases of similar orders of magnitudes.

The parameter-less versions of the cGA, namely the parallel-run version and the smart-restart version with budget factors $b = 16$ and $b = 1/\ln(n)$, generally perform very well. Their runtimes are, naturally, larger than the runtimes observed for the optimal static parameter value (which depends heavily on the problem and the noise level), but they clearly

avoid the often catastrophic performances in the strong genetic-drift regime. Overall, the smart-restart cGA with cautious budget factor $b = 1/\ln(n)$ appears best, in particular, for the two more difficult benchmarks JUMP and DECEPTIVELEADINGBLOCKS.

We also extend our smart-restart mechanism to a more complex EDA, population-based incremental learning (PBIL), which can be seen as a variant of the cross-entropy algorithm (Costa et al., 2007; De Boer et al., 2005). Since this algorithm has three parameters, it is not immediately obvious how to use our smart-restart approach. We solve this problem by keeping the selection pressure $\eta = \mu/\lambda$ and the learning rate ρ as parameters and by setting the sample size λ via the smart-restart mechanism (again guided by (Doerr and Zheng, 2020b, Theorem 3)). We apply this algorithm to two combinatorial optimization problems from the literature (Rubinstein and Kroese, 2004), the max-cut problem and the bipartition problem. Our empirical results show that the smart-restart mechanism uses much better values for the population size than the hand-crafted parameters of the previous work (Rubinstein and Kroese, 2004), resulting in significant speed-ups. We also implemented the two restart strategies from (Harik and Lobo, 1999; Auger and Hansen, 2005) discussed earlier. In our experiments, our smart-restart mechanism shows better performance on both combinatorial optimization problems.

We note that this version extends our preliminary version (Doerr and Zheng, 2020a), and differs largely in the following ways. Both in the mathematical runtime analysis and in the experiments, we also consider the presence of additive posterior noise for the cGA. To demonstrate that our general approach is feasible also for other EDAs ((Doerr and Zheng, 2020a) only considers the cGA), we add a runtime analysis for the smart-restart UMDA and evaluate a smart-restart version of PBIL (cross-entropy algorithm) with two other restart strategies for comparison on two combinatorial problems. Finally, this version contains all mathematical proofs that had to be omitted in (Doerr and Zheng, 2020a) for reasons of space.

The remainder of this paper is structured as follows. Section 2 introduces the preliminaries including a detailed description of the related algorithms and benchmark functions. The newly-proposed smart-restart mechanism will be stated in Section 3. Section 4 shows our theoretical results. Section 5 contains our experimental analyses on classic benchmark functions. The extension of our mechanism to the PBIL and an experimental analysis of this algorithm on two combinatorial problems are presented in Section 6. Section 7 concludes our paper.

2. Preliminaries

2.1 Algorithms

In this paper, we consider algorithms maximizing pseudo-Boolean functions $f : \{0, 1\}^n \rightarrow \mathbb{R}$. We regard so-called anytime algorithms, that is, algorithms that can be stopped at any time, and then return the best solution seen so far. In practice, such algorithms are run with some user-specified termination criterion. In our mathematical analyses, we regard the time it takes until an optimum is generated if the algorithm is not stopped prematurely. For that reason, we do not specify a termination criterion here.

Since our smart-restart mechanism builds on an original heuristic, such as the cGA of Harik et al. (1999) or the UMDA of Mühlenbein and Paass (1996), and since we will

compare our smart-restart cGA with the parallel-run cGA (Doerr, 2021), this subsection will give a brief introduction to these algorithms. We shall, in Section 6, also discuss the general performance of our smart-restart PBIL (cross-entropy algorithm), and for the convenience of reading, we will introduce this algorithm in Section 6.

In the following, we shall use $X_i^g = (X_{i,1}^g, \dots, X_{i,n}^g) \in \{0, 1\}^n$ to denote the i -th bitstring in the g -th iteration of the algorithm and $X_{i,j}^g$ for the j -th bit of X_i^g . We use $p^g = (p_1^g, \dots, p_n^g)$ to denote the (univariate) probabilistic model learned in the g -th iteration and p_j^g for the j -th entry of p^g .

2.1.1 THE COMPACT GENETIC ALGORITHM

The compact genetic algorithm (cGA) with hypothetical population size μ samples two individuals in each generation and moves the sampling frequencies by an absolute value of $1/\mu$ towards the bit values of the better individual. Usually, and so do we, in order to avoid frequencies reaching the absorbing boundaries 0 or 1, the artificial margins $1/n$ and $1 - 1/n$ are utilized, that is, we restrict the frequency values to be in the interval $[1/n, 1 - 1/n]$. The following Algorithm 1 shows the details. As is common in runtime analysis, we do not specify a termination criterion. When talking about the runtime of an algorithm, we mean the first time (measured by the number of fitness evaluations) an optimum was sampled.

Algorithm 1 The cGA to maximize a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ with hypothetical population size μ

```

1:  $p^0 = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2}) \in [0, 1]^n$ 
2: for  $g = 1, 2, \dots$  do
   %%Sample two individuals  $X_1^g, X_2^g$ 
3:   for  $i = 1, 2$  do
4:     for  $j = 1, 2, \dots, n$  do
5:        $X_{i,j}^g \leftarrow 1$  with probability  $p_j^{g-1}$  and  $X_{i,j}^g \leftarrow 0$  with probability  $1 - p_j^{g-1}$ 
6:     end for
7:   end for
   %%Update of the frequency vector
8:   if  $f(X_1^g) \geq f(X_2^g)$  then
9:      $p' = p^{g-1} + \frac{1}{\mu}(X_1^g - X_2^g)$ 
10:  else
11:     $p' = p^{g-1} + \frac{1}{\mu}(X_2^g - X_1^g)$ 
12:  end if
13:   $p^g = \min\{\max\{\frac{1}{n}, p'\}, 1 - \frac{1}{n}\}$ 
14: end for

```

2.1.2 THE UNIVARIATE MARGINAL DISTRIBUTION ALGORITHM

The univariate marginal distribution algorithm (UMDA) samples λ individuals in each generation and selects the best μ individuals to learn its probabilistic model. More precisely, the new sampling frequency p_i of the i -th bit is taken as the ratio of ones in the i -th bit of the μ selected individuals. Similar to the cGA, the artificial margins $1/n$ and $1 - 1/n$ are utilized to avoid the frequencies reaching the boundaries 0 or 1. See Algorithm 2 for details.

Algorithm 2 The UMDA (with sample size λ , and selection size μ) to maximize a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$

```

1:  $p^0 = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2}) \in [0, 1]^n$ 
2: for  $g = 1, 2, \dots$  do
    %%Sample  $\lambda$  individuals  $X_1^g, \dots, X_\lambda^g$ 
3:   for  $i = 1, 2, \dots, \lambda$  do
4:     for  $j = 1, 2, \dots, n$  do
5:        $X_{i,j}^g \leftarrow 1$  with probability  $p_j^{g-1}$  and  $X_{i,j}^g \leftarrow 0$  with probability  $1 - p_j^{g-1}$ 
6:     end for
7:   end for
    %%Update of the frequency vector
8:   Let  $\tilde{X}_1^g, \dots, \tilde{X}_\mu^g$  be the best  $\mu$  individuals (tie broken randomly)
9:    $p' = \frac{1}{\mu} \sum_{i=1}^{\mu} \tilde{X}_i^g$ 
10:   $p^g = \min\{\max\{\frac{1}{n}, p'\}, 1 - \frac{1}{n}\}$ 
11: end for

```

2.1.3 THE PARALLEL-RUN cGA

The parallel EDA mechanism was proposed by Doerr (2021) as a side result when discussing the connection between runtime bounds that hold with high probability and the expected runtime. For the cGA, this mechanism yields the following *parallel-run cGA*. In the initial round $\ell = 1$, we start process $\ell = 1$ to run the cGA with population size $\mu = 2^{\ell-1}$ for 1 generation. In round $\ell = 2, 3, \dots$, all running processes $j = 1, \dots, \ell - 1$ run $2^{\ell-1}$ generations and then we start process ℓ to run the cGA with population size $\mu = 2^{\ell-1}$ for $\sum_{i=0}^{\ell-1} 2^i$ generations. The algorithm terminates once any process has found the optimum. Algorithm 3 shows the details of the parallel-run cGA.

Based on the following assumption, Doerr (2021) proved that the expected runtime for this parallel-run cGA is at most $6\tilde{\mu}T(\log_2(\tilde{\mu}T) + 3)$.

Assumption (Doerr, 2021): Consider using the cGA with population size μ to maximize a given function f . Assume that there are unknown $\tilde{\mu}$ and T such that the cGA for all population sizes $\mu \geq \tilde{\mu}$ optimizes this function f in μT fitness evaluations with probability at least $\frac{3}{4}$.

Algorithm 3 The parallel-run cGA to maximize a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$

```

1: Process 1 runs cGA (Algorithm 1) with population size  $\mu = 1$  for 1 generation
2: for round  $\ell = 2, \dots$  do
3:   Processes  $1, \dots, \ell - 1$  continue to run for another  $2^{\ell-1}$  generations, one process after the other one
4:   Start process  $\ell$  to run cGA (Algorithm 1) with population size  $\mu = 2^{\ell-1}$  to maximize  $f$  and run it for  $\sum_{i=0}^{\ell-1} 2^i$  generations
5: end for

```

2.2 Benchmark Functions

To understand the performance of the mechanism proposed in this work, we regard four basic benchmark functions (detailed in Section 2.2.1 below). They are all popular benchmarks in the analysis of evolutionary algorithms, so they are well-understood both from a theoretical and an experimental point of view, which helps interpreting our results. To further test the performance of the proposed approach in a more complicated setting, we also consider a noisy environment (detailed in Section 2.2.2). We shall, in Section 6, also conduct an experimental investigation of the PBIL (cross-entropy algorithm) on particular instances of two combinatorial optimization problems (namely those suggested in (Rubinstejn and Kroese, 2004) to show the power of this algorithm), but since these results cannot be easily compared to the other ones, among others, because no theory exists for these instances, we describe these problems not here, but in Section 6.

2.2.1 BASIC BENCHMARK FUNCTIONS

We selected the four benchmark functions ONEMAX, LEADINGONES, JUMP, and DECEPTIVELEADINGBLOCKS as optimization problems. All four problems are defined on binary representations (bit strings) and we use n to denote their length, that is, all are functions $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

The **OneMax** problem is one of the easiest benchmark problems. The ONEMAX fitness of a bit string is simply the number of ones in the bit string. Formally, the ONEMAX function value of any $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ is defined by

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

Having the perfect fitness-distance correlation, most evolutionary algorithms find it easy to optimize ONEMAX. A common and often easy to prove runtime is $\Theta(n \log n)$ (Mühlenbein, 1992; Garnier et al., 1999; Jansen et al., 2005; Witt, 2006; Rowe and Sudholt, 2014; Doerr and Künnemann, 2015; Antipov and Doerr, 2021; Oliveto et al., 2022). For EDAs, apparently, the runtime of ONEMAX is more complicated. The known results for EDAs are the following. The first mathematical runtime analysis for EDAs by Droste (2006) together with the recent work (Sudholt and Witt, 2019) shows that the cGA can efficiently optimize ONEMAX in time $\Theta(\mu\sqrt{n})$ when $\mu \geq K\sqrt{n} \ln(n)$ for some sufficiently large constant K . As the proofs of this result show (and the same could be concluded from the general result (Doerr and Zheng, 2020b)), in this parameter regime there is little genetic drift. Throughout the runtime, with high probability, all bit frequencies stay above $\frac{1}{4}$. For hypothetical population sizes below the $\sqrt{n} \log n$ threshold, the situation is less understood. However, the lower bound of $\Omega(\mu^{1/3}n)$ valid for all $\mu = O\left(\frac{\sqrt{n}}{\ln(n) \ln \ln(n)}\right)$ proven in (Lengler et al., 2021) together with its proof shows that in parts of this regime the cGA suffers from genetic drift, leading to (mildly) higher runtimes.

For another EDA, the UMDA with sampling population size λ and selected population size μ , a runtime of $\Omega(\lambda\sqrt{n} + n \log n)$ for $\mu = \Theta(\lambda)$ and $\lambda \in \text{poly}(n)$ is shown in (Krejca and Witt, 2020b). Upper bounds were independently proven in (Witt, 2019; Dang et al., 2019). For $\mu = \Theta(\lambda)$, (Witt, 2019) proves a runtime of $O(\lambda n)$ for $\lambda = \Omega(\log n) \cap o(n)$ and a runtime of $O(\lambda\sqrt{n})$ for $\lambda = \Omega(\sqrt{n} \log n) \cap n^{O(1)}$. Replacing the restriction $\mu = \Theta(\lambda)$ with $\lambda = \Omega(\mu)$,

in (Dang et al., 2019), an upper bound of $O(\lambda n)$ is obtained for $\lambda = \Omega(\log n) \cap O(\sqrt{n})$ and one of $O(\lambda\sqrt{n})$ for $\lambda = \Omega(\sqrt{n} \log n)$.

The **LeadingOnes** benchmark is still an easy unimodal problem, however, typically harder than ONEMAX. The LEADINGONES value of a bit string is the number of ones in it, counted from left to right, until the first zero. Formally, the LEADINGONES function value of any $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ is defined by

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

How simple randomized search heuristics optimize LEADINGONES is extremely well understood (Droste et al., 2002; Jansen et al., 2005; Witt, 2006; Böttcher et al., 2010; Sudholt, 2013; Doerr, 2019; Lissovoi et al., 2020; Sudholt, 2021; Doerr et al., 2021), many EAs optimize this benchmark in time $\Theta(n^2)$. Surprisingly, no theoretical results are known on how the cGA optimizes LEADINGONES. However, the runtime of the UMDA with population sizes $\mu = \Theta(\lambda)$ with suitable implicit constants and $\lambda = \Omega(\log n)$ was shown to be $O(n\lambda \log(\lambda) + n^2)$ (Dang et al., 2019) and, recently, $\Theta(n\lambda)$ for $\lambda = \Omega(n \log n)$ (Doerr and Krejca, 2021b). We remark that (Doerr and Zheng, 2020b) for this situation shows that genetic drift occurs when $\lambda = O(n)$ (with suitable implicit constants). Consequently, these results show a roughly linear influence of λ on the runtime when λ is (roughly) at least linear in n , but below this value, there is apparently no big penalty for running the EDA in the genetic drift regime. For the cGA, we will observe a different behavior, which also indicates that translating general behaviors from one EDA to another, even within the class of univariate EDAs, has to be done with caution.

The **Jump** benchmark is a class of multimodal fitness landscapes of scalable difficulty. For a difficulty parameter k , the fitness landscape is isomorphic to the one of ONEMAX except that there is a valley of low fitness of width k around the optimum. More precisely, all search points in distance 1 to $k-1$ from the optimum have a fitness lower than all other search points. Formally, the JUMP $_k$ function (with $k \in [1..n]$) is defined by

$$f(x) = \begin{cases} k + \sum_{i=1}^n x_i, & \text{if } \sum_{i=1}^n x_i \leq n - k \text{ or } x = 1^n, \\ n - \sum_{i=1}^n x_i, & \text{else,} \end{cases}$$

for any $x = (x_1, \dots, x_n) \in \{0, 1\}^n$.

Recent results (Hasenöhl and Sutton, 2018; Doerr, 2021) show that when μ is large enough, then the cGA can optimize JUMP functions quite efficiently and significantly more efficient than many classic evolutionary algorithms. We omit some details and only mention that for k not too small, a runtime exponential in k results from a population size μ that is also exponential in k . This is much better than the $\Omega(n^k)$ runtime of typical mutation-based evolutionary algorithms (Droste et al., 2002; Doerr et al., 2017; Doerr, 2022; Rajabi and Witt, 2022) or the $n^{O(k)}$ runtime bounds shown for several crossover-based algorithms (Dang et al., 2016; Antipov et al., 2022). We note that $O(n)$ and $O(n \log n)$ runtimes have been shown in (Whitley et al., 2018; Rowe and Aishwaryaprajna, 2019), however, these algorithms appear quite problem-specific (Witt, 2023) and have not been regarded in other contexts so far. It was not known whether the runtime of the cGA becomes worse in the regime with

genetic drift, but our experimental results now show an enormously weak performance in this regime.

The **DeceptiveLeadingBlocks** benchmark was introduced in (Lehre and Nguyen, 2019). It can be seen as a deceptive version of the LEADINGONES benchmark. In DECEPTIVELEADINGBLOCKS, the bits are partitioned into blocks of length two in a left-to-right fashion. The fitness is computed as follows. Counting from left to right, each block that consists of two ones contributes two to the fitness, until the first block is reached that does not consist of two ones. This block contributes one to the fitness if it consists of two zeros, otherwise it contributes zero. All further blocks do not contribute to the fitness. Formally, the DECEPTIVELEADINGBLOCKS (requiring n is even) function value of any $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ is defined by

$$f(x) = \begin{cases} 2m + 1, & \text{if } x_{[1..2m]} = 1^{2m} \text{ and } x_{2m+1} = x_{2m+2} = 0, \\ 2m, & \text{if } x_{[1..2m]} = 1^{2m} \text{ and } x_{2m+1} + x_{2m+2} = 1, \\ n, & \text{if } x = 1^n. \end{cases}$$

The main result in (Lehre and Nguyen, 2019) is that when $\mu = \Theta(\lambda)$ and $\lambda = o(n)$, the expected runtime of the UMDA on DECEPTIVELEADINGBLOCKS is $\exp(\Omega(\lambda))$. We note that when $\lambda = o(n)$, already after a quadratic runtime strong genetic drift is encountered according to (Doerr and Zheng, 2020b). When $\lambda = \Omega(n \log n)$, a runtime guarantee of at most $(1 + o(1))\frac{1}{2}\lambda n$ holds with high probability (Doerr and Krejca, 2021a). Hence for this function and the UMDA as an optimizer, the choice of the population size is again very important. This was the reason for including this function into our set of test problems and the results indicate that indeed the cGA shows a behavior similar to what the mathematical results showed for the UMDA.

We note that (Lehre and Nguyen, 2019) also shows an expected runtime of $O(n\lambda \log \lambda + n^3)$ when $\mu = \Omega(\log n)$ and $\lambda = \Omega(\mu^2)$, which is an unusually high selection pressure. Other runtime results on the DECEPTIVELEADINGBLOCKS function include several $O(n^3)$ runtime guarantees for classic EAs (Lehre and Nguyen, 2019) as well as a $\Theta(n^2)$ runtime for the Metropolis algorithm and an $O(n \log n)$ runtime guarantee for the significance-based cGA (Wang et al., 2021). Till now, there is no theoretical runtime analysis for the cGA.

2.2.2 ADDITIVE CENTERED GAUSSIAN POSTERIOR NOISE

In practical applications, one often encounters various forms of uncertainty. One of these is a noisy access to the objective function. Friedrich et al. (2017) analyzed how the cGA optimizes the ONEMAX problem under additive centered Gaussian posterior noise. They proved that for all noise intensities (variances σ^2 of the Gaussian distribution), there is a population size $\mu = \mu(\sigma^2)$ which depends only polynomially on σ^2 (that is, $\mu(\sigma^2)$ is a polynomial in σ^2) so that the cGA with this population size efficiently solves the ONEMAX problem. This was called *graceful scaling*. They also provided a restart scheme that obtains this performance without knowledge of the noise intensity (however, it requires to know the polynomial $\mu(\sigma^2)$). Hence these results show that the cGA can deal well with the type of noise regarded, and much better than many classic evolutionary algorithms (see the lower bounds in (Gießen and Kötzing, 2016; Friedrich et al., 2017)), but this still needs an action by the algorithm user, namely an appropriate choice of the population size μ .

As we shall show in this work, our restart scheme is also able to optimize noisy versions of ONEMAX and many other problems, but without knowing the polynomial $\mu(\sigma^2)$ and using significantly more efficient values for the population size. For ONEMAX, we prove rigorously that we obtain essentially the performance of the original cGA with the best choice of the population size (Theorem 9), where we profit from the fact that the runtime analysis of (Friedrich et al., 2017) shows that the cGA also for noisy ONEMAX functions essentially satisfies our main assumption that from a certain population size on, the runtime of the cGA is at most proportional to the population size.

We conduct experiments for various benchmark functions in this noise model. They indicate that also for problems different from ONEMAX, the graceful scaling property holds. However, they also show that much smaller population sizes suffice to cope with the noise. Consequently, our smart-restart cGA (as well as the parallel-run cGA from (Doerr, 2021)) optimizes ONEMAX much faster than the algorithms proposed in (Friedrich et al., 2017). This is natural since the parameter-less approaches also try smaller (more efficient in case of success) population sizes, whereas the approaches in (Friedrich et al., 2017) use a population size large enough that one can prove via mathematical means that they will be successful with high probability.

We now make precise the additive centered Gaussian noise model. We take the common assumption that whenever the noisy fitness of a search point is regarded in a run of the algorithm, its noisy fitness is computed anew, that is, with newly sampled noise. This avoids that a single exceptional noise event misguides the algorithm for the remaining run. A comparison of the results in (Sudholt and Thyssen, 2012) (without independent reevaluations) and (Doerr et al., 2012) (with reevaluations) shows how detrimental sticking to previous evaluations can be. We regard posterior noise, that is, the noisy fitness value is obtained from a perturbation of the original fitness value (independent of the argument) as opposed to prior noise, where the algorithm works with the fitness of a perturbed search point. We regard additive perturbations, hence the perceived fitness of a search point x is $f(x) + D$, where f is the original fitness function and D is an independent sample from a distribution describing the noise. Since we consider centered Gaussian noise, we always have $D \sim \mathcal{N}(0, \sigma^2)$, where $\mathcal{N}(0, \sigma^2)$ denotes the Gaussian distribution with expectation zero and variance $\sigma^2 \geq 0$. Obviously, the classic noise-free optimization scenario is subsumed by the special case $\sigma^2 = 0$.

3. The Smart-Restart Mechanism

In this section, we introduce our *smart-restart mechanism*. It can be applied to any randomized search heuristic \mathcal{A} having an integral parameter μ and it makes sense when we can assume that the algorithm has a good performance from a certain value for μ on, a situation often encountered in EDAs. In contrast to the parallel-run mechanism proposed in (Doerr, 2021), which applies to the same scenario, the smart-restart mechanism does not run processes in parallel, which is an advantage from the implementation point of view. The main advantage we aim for is that by predicting when runs with a certain parameter value become hopeless, we can abort these runs and save runtime.

As in (Doerr, 2021), in this exposition we let ourselves be guided by the cGA as base algorithm \mathcal{A} as it is maybe the simplest algorithm in which a parameter behavior as sketched

above is encountered. To decide when to abort a run, we use the first tight quantification of the genetic drift effect of the EDAs by Doerr and Zheng (2020b). In detail, they proved that in a run of the cGA with hypothetical population size μ a frequency of a neutral bit will reach the boundaries of the frequency range in an expected number of at most $4\mu^2$ generations (equivalent to $8\mu^2$ fitness evaluations), which is asymptotically tight. By Markov's inequality the probability that a boundary is reached in $b\mu^2, b > 8$, fitness evaluations, is at least $1 - 8/b$.

This finding inspires the following restart scheme (for any randomized search heuristic \mathcal{A} with a parameter μ), also described in Algorithm 4. We repeat running algorithm \mathcal{A} with increasing values for μ , each time until we decide to abort such a run. For the ℓ -th run, $\ell = 1, 2, \dots$, we use the parameter value $\mu_\ell = 2U^{\ell-1}$, that is, we start with a small value $\mu = 2$ and increase μ by a factor of $U > 1$, called *update factor*, from one run to the next. We abort the ℓ -th run after $B_\ell = b\mu_\ell^2$ fitness evaluations, where b is the second parameter of the restart scheme, called *budget factor*. As before, we do not specify a termination criterion since for our analysis we just count the number of fitness evaluations until a desired solution is found.

We shall discuss the settings of the parameters U and b later in more detail. As a motivating example, we note that for the cGA by taking $b = 16$, the above Markov bound argument shows that in each such run, each bit has a probability of at most $\frac{1}{2}$ to be subject to strong genetic drift. Hence this restart scheme manages to run the cGA in a way that genetic drift does not affect too many bits. We shall later see that a budget factor $b = O(1/\log n)$ can even ensure that none of the bits is subject to strong genetic drift.

Algorithm 4 The smart-restart mechanism with update factor U and budget factor b applied to an algorithm \mathcal{A} with parameter μ for the maximization of a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

- 1: **for** $\ell = 1, 2, \dots$ **do**
 - 2: Run \mathcal{A} with parameter value $\mu_\ell = 2U^{\ell-1}$ for $B_\ell = b\mu_\ell^2$ fitness evaluations on the maximization problem f
 - 3: **end for**
-

4. Theoretical Analyses

In this section, we prove mathematical runtime guarantees for our smart-restart mechanism.

4.1 A General Performance Guarantee

We follow the general approach of (Doerr, 2021, Section 2.4) of assuming that the runtime increases linearly with the population size from a given minimum size $\tilde{\mu}$ on. We do not restrict this property for the cGA with population size μ , but for the general EDA (or any randomized search heuristic) with a parameter μ .

Assumption (L): Let $p \in (0, 1]$. Consider using an EDA (or a randomized search heuristic) with parameter μ to maximize a given function f . Assume that there are unknown $\tilde{\mu}$ and T such that the EDA (or randomized search heuristic) for all parameter values $\mu \geq \tilde{\mu}$ optimizes f within μT fitness evaluations with probability at least p .

This Assumption (L) is identical to the assumption taken in (Doerr, 2021) except that there p was required to be at least $3/4$, whereas we allow a general positive p (but note that we will require $p > 1 - \frac{1}{U^2}$, hence a small p limits the choice of U). Since most existing runtime analyses give bounds with success probability $1 - o(1)$, this difference is, of course, not very important. We note that the proof of the result in (Doerr, 2021) requires p to be at least $3/4$, but we also note that an elementary probability amplification argument (via independent restarts) allows to increase the success probability of a given algorithm, so that the result of (Doerr, 2021) becomes applicable to this modified algorithm.

Under this Assumption (L), we obtain the following result. We note that it is non-asymptotic, which later allows to easily obtain asymptotic results also for non-constant parameters. We note that when assuming p and U to be constants (which is very natural), then the bound becomes $O(\max\{b\tilde{\mu}^2, T^2/b, \tilde{\mu}T\})$.

Theorem 1 *Let $U > 1$ and $b > 0$. Consider using the smart-restart mechanism on an EDA (or a randomized search heuristic) with update factor U and budget $B_\ell = b\mu_\ell^2, \ell = 1, 2, \dots$, optimizing a function f satisfying Assumption (L) with $p \in (1 - \frac{1}{U^2}, 1]$. Then the expected time until the optimum of f is generated is at most*

$$\left(\frac{U^2}{U^2 - 1} + \frac{(1-p)U^2}{1 - (1-p)U^2} \right) \max \left\{ b\tilde{\mu}^2, \frac{T^2}{b} \right\} + \frac{pU}{1 - (1-p)U} \tilde{\mu}T$$

fitness evaluations.

Proof Let $\ell' = \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, B_\ell \geq 2U^{\ell-1}T\}$. With $B_\ell = b\mu_\ell^2$ and $\mu_\ell = 2U^{\ell-1}$, we have $\ell' = \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, b(2U^{\ell-1})^2 \geq 2U^{\ell-1}T\} = \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, 2U^{\ell-1} \geq T/b\}$, that is, $\ell' = \min\{\ell \mid 2U^{\ell-1} \geq \max\{\tilde{\mu}, T/b\}\}$. Then it is not difficult to see that $2U^{\ell'-1} \leq U \max\{\tilde{\mu}, T/b\}$ and that for any $\ell \geq \ell'$, the population size $\mu_\ell := 2U^{\ell-1}$ satisfies $\mu_\ell \geq \tilde{\mu}$ and $B_\ell \geq \mu_\ell T$. Hence, according to the assumption, we know the EDA (randomized search heuristic) with such a μ_ℓ optimizes f with probability at least p in time $\mu_\ell T$. We pessimistically assume that the optimum is not reached before the parameter value increases to $\mu_{\ell'}$. Now the expected time when the smart-restart EDA (randomized search heuristic) finds the optimum of f is at most

$$\begin{aligned} & \sum_{i=1}^{\ell'-1} B_i + p \cdot 2U^{\ell'-1}T + \sum_{i=1}^{\infty} (1-p)^i p \left(\sum_{j=0}^{i-1} B_{\ell'+j} + 2U^{\ell'+i-1}T \right) \\ &= \sum_{i=1}^{\ell'-1} B_i + p \cdot 2U^{\ell'-1}T + \sum_{i=1}^{\infty} (1-p)^i p \sum_{j=0}^{i-1} B_{\ell'+j} + 2pU^{\ell'-1}T \sum_{i=1}^{\infty} (1-p)^i U^i \\ &= \sum_{i=1}^{\ell'-1} B_i + 2pU^{\ell'-1}T + \sum_{j=0}^{\infty} B_{\ell'+j} p \sum_{i=j+1}^{\infty} (1-p)^i + 2pU^{\ell'-1}T \frac{(1-p)U}{1 - (1-p)U} \\ &= \sum_{i=1}^{\ell'-1} B_i + \sum_{j=0}^{\infty} B_{\ell'+j} p \frac{(1-p)^{j+1}}{1 - (1-p)} + \frac{2pU^{\ell'-1}T}{1 - (1-p)U} \\ &= \sum_{i=1}^{\ell'-1} B_i + \sum_{j=0}^{\infty} (1-p)^{j+1} B_{\ell'+j} + \frac{2pU^{\ell'-1}T}{1 - (1-p)U}, \end{aligned}$$

where the second equality uses $(1-p)U \in [0, 1)$ from $p \in (1 - \frac{1}{U^2}, 1]$. With $B_\ell = b\mu_\ell^2 = b(2U^{\ell-1})^2 = 4bU^{2\ell-2}$, we further compute

$$\begin{aligned}
 & \sum_{i=1}^{\ell-1} B_i + \sum_{j=0}^{\infty} (1-p)^{j+1} B_{\ell+j} + \frac{2pU^{\ell-1}T}{1 - (1-p)U} \\
 &= \sum_{i=1}^{\ell-1} 4bU^{2i-2} + \sum_{j=0}^{\infty} (1-p)^{j+1} 4bU^{2\ell+2j-2} + \frac{2pU^{\ell-1}T}{1 - (1-p)U} \\
 &= \frac{4b(U^{2\ell-2} - 1)}{U^2 - 1} + \frac{4b(1-p)U^{2\ell-2}}{1 - (1-p)U^2} + \frac{2pU^{\ell-1}T}{1 - (1-p)U} \\
 &\leq \frac{bU^2 \max\{\tilde{\mu}^2, T^2/b^2\}}{U^2 - 1} + \frac{b(1-p)U^2 \max\{\tilde{\mu}^2, T^2/b^2\}}{1 - (1-p)U^2} + \frac{pU\tilde{\mu}T}{1 - (1-p)U} \\
 &= \left(\frac{U^2}{U^2 - 1} + \frac{(1-p)U^2}{1 - (1-p)U^2} \right) \max\left\{ b\tilde{\mu}^2, \frac{T^2}{b} \right\} + \frac{pU}{1 - (1-p)U} \tilde{\mu}T,
 \end{aligned}$$

where the second equality uses $(1-p)U^2 \in [0, 1)$ from $p \in (1 - \frac{1}{U^2}, 1]$ and the first inequality uses $2U^{\ell-1} \leq U \max\{\tilde{\mu}^2, T^2/b^2\}$. \blacksquare

For comparison, we recall that the complexity of the parallel-run cGA from (Doerr, 2021).

Theorem 2 (Doerr (2021, Theorem 2)) *The expected number of fitness evaluations for the parallel-run cGA optimizing a function f satisfying Assumption (L) with $p \geq 3/4$ is $O(\tilde{\mu}T \log(\tilde{\mu}T))$.*

Since the choice $b = \Theta(T/\tilde{\mu})$ gives an asymptotic runtime of $O(\tilde{\mu}T)$ for the smart-restart cGA, we see that with the right choice of the parameters the smart-restart cGA can outperform the parallel-run cGA slightly. This shows that it indeed gains from its ability to abort unprofitable runs.

Our main motivation for regarding Assumption (L) was that this runtime behavior is often observed both in theoretical results (see, e.g., the survey (Krejca and Witt, 2020a)) and in experiments (see Section 5). Unfortunately, some theoretical results were only proven under the additional assumption that μ is polynomially bounded in n , that is, that $\mu = O(n^C)$ for some, possibly large, constant C . For most of these results, we are convinced that the restriction on μ is not necessary, but was only taken for convenience and in the light that super-polynomial values for μ would imply not very interesting super-polynomial runtimes. To extend such results to our smart-restart cGA in a formally correct manner, we now prove a version of Theorem 1 applying to such settings. More precisely, we regard the following assumption. Similar to Assumption (L), we do not restrict this property to the cGA with the population size as parameter.

Assumption (L'): Let $p \in (0, 1]$. Consider using an EDA (or a randomized search heuristic) with parameter μ to maximize a given function f . Assume that there are unknown $\tilde{\mu}$, μ^+ , and T such that the EDA (or randomized search heuristic) for all parameter values $\tilde{\mu} \leq \mu \leq \mu^+$ optimizes f within μT fitness evaluations with probability at least p .

We prove the following result.

Theorem 3 *Let $U > 1$ and $b > 0$. Consider using the smart-restart mechanism on an EDA (or a randomized search heuristic) with update factor U and budget $B_\ell = b\mu_\ell^2$, $\ell = 1, 2, \dots$, optimizing a function f satisfying Assumption (L') with $p \in (1 - \frac{1}{U^2}, 1)$. Let $\ell' := \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, B_\ell \geq 2U^{\ell-1}T\}$ and $\mathcal{L} := \{\ell \in \mathbb{Z} \mid \ell \geq \ell', 2U^{\ell-1} \leq \mu^+\}$. Then, apart from when an exceptional event of probability at most $(1-p)^{|\mathcal{L}|}$ holds, the expected time until the optimum of f is generated is at most*

$$\left(\frac{U^2}{U^2 - 1} + \frac{(1-p)U^2}{1 - (1-p)U^2} \right) \max \left\{ b\tilde{\mu}^2, \frac{T^2}{b} \right\} + \frac{pU}{1 - (1-p)U} \tilde{\mu}T$$

fitness evaluations.

Proof Let A be the event that none of the runs of the EDA (or randomized search heuristic) with parameter $\mu_\ell = 2U^{\ell-1}$ at most μ^+ finds the optimum of f . As in the proof of Theorem 1, each of the runs using parameter μ_ℓ , $\ell \in \mathcal{L}$, with probability at least p finds the optimum. Hence the event A occurs with probability at most $(1-p)^{|\mathcal{L}|}$.

Let us condition on the event $\neg A$. Under this event, the smart-restart EDA (or randomized search heuristic) surely finds the optimum before the parameter value is increased beyond μ^+ . Note that when running the EDA (or randomized search heuristic) with a parameter value of at most μ^+ , the Assumptions (L) and (L') are identical. For that reason, analogous to the first paragraph of the proof of Theorem 1, we see that any $\ell \in \mathcal{L}$ a run of the EDA (or randomized search heuristic) with parameter $\mu_\ell = 2U^{\ell-1}$ finds the optimum of f in time $\mu_\ell T$ with probability at least $p / \Pr[\neg A] \geq p$. Consequently, analogous to that proof, the expected runtime conditional on $\neg A$ is at most

$$\sum_{i=1}^{\ell'-1} B_i + p \cdot 2U^{\ell'-1}T + \sum_{i=1}^{\max\{\mathcal{L}\}-\ell'} (1-p)^i p \left(\sum_{j=0}^{i-1} B_{\ell'+j} + 2U^{\ell'+i-1}T \right). \quad (1)$$

This expression is identical to the corresponding one in the proof of Theorem 1 except that the second sum is not taken over the range $i \in \mathbb{Z}_{\geq 1}$, but only the range $i \in [1, \max\{\mathcal{L}\} - \ell']$. Since these sums involve positive terms only, we can bound (1) from above in exactly the same way as in the proof of Theorem 1. This shows our claim. \blacksquare

4.2 Specific Runtime Results

The following examples show how to combine our general runtime analysis with known runtime results to obtain performance guarantees for smart-restart EDAs on several specific problems.

4.2.1 ONEMAX AND JUMP

We recall the following runtime results for the cGA on ONEMAX (Sudholt and Witt, 2019) and JUMP (Doerr, 2021) as well as for the UMDA on ONEMAX (Witt, 2019; Dang et al., 2019). As common, by *runtime* we mean the number of fitness evaluations until the optimum is sampled. This is, essentially, two times the number of generations until the optimum is sampled for the cGA, and λ times this generation number for the UMDA.

Theorem 4 (Sudholt and Witt (2019); Doerr (2021); Witt (2019); Dang et al. (2019)) *Let $K > 0$ be a sufficiently large constant and let $C > 0$ be any constant. Consider the cGA with $K\sqrt{n}\ln n \leq \mu \leq n^C$ and the UMDA with $K\sqrt{n}\ln n \leq \mu \leq n^C$ and $\lambda = \Theta(\mu)$ (Witt, 2019) (or the UMDA with $\mu \geq K\sqrt{n}\ln n$ and $\lambda \geq a\mu$ for sufficiently large constant $a > 1$ (Dang et al., 2019)).*

- *The expected runtimes on the ONEMAX function are $O(\mu\sqrt{n})$ for the cGA (Sudholt and Witt, 2019, Theorem 2) and $O(\lambda\sqrt{n})$ for the UMDA, see (Witt, 2019, Theorem 10) and (Dang et al., 2019, Theorem 9).*
- *With probability $1 - o(1)$, the optimum of the JUMP function with jump size $k < \frac{1}{20} \ln n$ is found by the cGA in time $O(\mu\sqrt{n})$ (Doerr, 2021, Theorem 9).*

With the optimal parameters, that is, with the smallest applicable population sizes, these runtimes are all $O(n \log n)$.

With Theorem 3, we have the following result.

Theorem 5 *Consider the smart-restart cGA with update factor $U > 1$ optimizing the JUMP function with jump size $k < \frac{1}{20} \ln n$ or the ONEMAX function, or the smart-restart UMDA with update factor $U > 1$ optimizing the ONEMAX function. Then, apart from a rare event of probability at most $n^{-\omega(1)}$, we have the following estimates for the expected runtime.*

- *If the budget factor b is $\Theta(1/\log n)$, then the expected runtime is $O(n \log n)$.*
- *If the budget factor b is between $\Omega(1/\log^2 n)$ and $O(1)$, then the expected runtime is $O(n \log^2 n)$.*

Proof For the smart-restart cGA, we note that the JUMP result (more specifically, the “with probability $1 - o(1)$ ” clause) also applies to ONEMAX simply because the JUMP function with jump size $k = 1$ has a fitness landscape that can in a monotonic manner be transformed into the one of the ONEMAX function. Hence for n sufficiently large, we have Assumption (L’) satisfied with $\tilde{\mu} = K\sqrt{n}\ln n$, $\mu^+ = n^C$, $T = O(\sqrt{n})$, and $p = 1 - o(1)$. Consequently, for any (constant) $U > 1$ we have $p \in (1 - \frac{1}{U^2}, 1)$. Given the above information on $\tilde{\mu}$ and T , we see that any $b \in n^{-o(1)} \cap n^{o(1)}$ gives that $\ell' = \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, B_\ell \geq 2U^{\ell-1}T\} = (1 \pm o(1))\frac{1}{2} \log_U(n)$. Since $\mu^+ = n^C$, we have $|\mathcal{L}| = (1 \pm o(1))(C - \frac{1}{2}) \log_U(n)$.

For the results where Theorem 4 only gives bounds on the expected runtime, we note that such statements can easily be transferred to a statement with a given probability via Markov’s bound $\Pr[\xi \geq cE[\xi]] \leq 1/c$ for any $c \geq 1$. Consequently, for the smart-restart UMDA with the population size λ as parameter of interest, we can set $p \in (1 - \frac{1}{U^2}, 1)$ for any (constant) $U > 1$ to satisfy Assumption (L’) together with $\tilde{\mu} = K\sqrt{n}\ln n$, $\mu^+ = n^C$, and $T = O(\sqrt{n})$. Hence, the same arguments as above also show the other claims. ■

Hence, the smart-restart cGA and smart-restart UMDA with $b = \Theta(1/\log n)$ have essentially the same time complexity as the original cGA and UMDA with optimal population size (see Theorem 4). A constant value for b results in a slightly inferior runtime of $O(n \log^2 n)$, which is also the runtime guarantee for the parallel-run cGA (Theorem 2).

4.2.2 LEADINGONES AND DECEPTIVELEADINGBLOCKS

As discussed in Section 2.2.1, no theoretical runtime guarantees exists for the cGA on the LEADINGONES and DECEPTIVELEADINGBLOCKS functions. For the UMDA, the following results are known.

Theorem 6 (Doerr and Krejca (2021a,b)) *Let $K > 0$ and $C > 1$ be sufficiently large constants. Consider the UMDA selecting $\mu \geq Kn \ln n$ best from the sampling population with size $\lambda \geq C\mu$.*

- *The expected runtime on the LEADINGONES function is $O(\lambda n)$ (Doerr and Krejca, 2021b, Theorem 5).*
- *With probability $1 - o(1)$, the optimum of the DECEPTIVELEADINGBLOCKS is found in time $O(\lambda n)$ (Doerr and Krejca, 2021a, Theorem 3).*

With the optimal parameter choice, that is, the smallest admissible population sizes, these runtimes are $O(n^2 \log n)$.

For simplicity, we could just require $\lambda = C\mu$ and Theorem 6 still holds. For n sufficiently large, we have Assumption (L) with respect to the parameter λ satisfied with $\tilde{\mu} = CKn \ln n, T = O(n)$, and $p = 1 - o(1)$. Consequently, for any (constant) $U > 1$ we have $p \in (1 - \frac{1}{U^2}, 1)$. Given the above information on $\tilde{\mu}$ and T , and with Theorem 1, we have the following result.

Theorem 7 *Consider the smart-restart UMDA with update factor $U > 1$ and constant selection pressure in UMDA optimizing the LEADINGONES or DECEPTIVELEADINGBLOCKS function. Then we have the following estimates for the expected runtime.*

- *If the budget factor b is $\Theta(1/\log n)$, then the expected runtime is $O(n^2 \log n)$.*
- *If the budget factor b is between $\Omega(1/\log^2 n)$ and $O(1)$, then the expected runtime is $O(n^2 \log^2 n)$.*

Hence, our smart-restart UMDA with $b = \Theta(1/\log n)$ has essentially the same time complexity as the original UMDA (Theorem 6) with optimal population sizes λ and μ .

4.2.3 NOISY ONEMAX

For another example, we recall the runtime of the cGA (without artificial margins) on the ONEMAX function with additive centered Gaussian noise from (Friedrich et al., 2017).

Theorem 8 (Friedrich et al. (2017, Theorem 5)) *Let $h : [0, \infty) \rightarrow [0, \infty)$ and $h \in \omega(1) \cap n^{o(1)}$. Consider the n -dimensional ONEMAX function with additive centered Gaussian noise with variance $\sigma^2 > 0$. Then with probability $1 - o(1)$, the cGA (without margins) with population size $\mu = h(n)\sigma^2\sqrt{n} \log n$ has all frequencies at 1 in $O(\mu\sigma^2\sqrt{n} \log(\mu n)) = O(h(n)\sigma^4 n \log^2 n)$ iterations.*

We note that here the cGA is used without restricting the frequencies to the interval $[1/n, 1 - 1/n]$, whereas more commonly (and in the remainder of this paper) the cGA is equipped with the margins $1/n$ and $1 - 1/n$ to avoid that frequencies reach the absorbing boundaries 0 or 1. Since our general runtime results do not rely on such implementation details but merely lift a result for a particular cGA to its smart-restart version, this poses no problems for us. As a side remark, though, we note that we are very optimistic that the above result from (Friedrich et al., 2017) holds equally for the setting with frequency margins.

More interestingly, the runtime result above is not of the type that for μ sufficiently large, the expected runtime is $O(\mu T)$ for some T (since μ appears also in the $\log(\mu n)$ term). Fortunately, with Theorem 3 at hand, we have an easy solution. By only regarding values of μ that are at most n^C for some constant C (which we may choose), the $\log(\mu n)$ term can be bounded by $O(\log n)$. Since the minimal applicable μ (the $\tilde{\mu}$ in the notation of Theorem 3) depends on σ^2 , this also implies that we can only regard polynomially bounded variances, but it is clear that any larger variances can be only of a purely academic interest. We thus formulate and prove the following result. We note that with more work, we could also have extended Theorem 1 to directly deal with the runtime behavior described in Theorem 8. For example, we could exploit that the geometric series showing up in the analysis do not change significantly when an extra logarithmic term is present. However, this appears to be a lot of work for a logarithmic term for which it is not even clear if it is necessary in the original result. Hence, we will not discuss them and focus on the following result.

Theorem 9 *Let $C \geq 1$ and $U > 1$. Let $h : [0, \infty) \rightarrow [0, \infty)$ and $h \in \omega(1) \cap n^{o(1)}$. Consider the smart-restart cGA with the update factor U and budget factor b optimizing the n -dimensional ONEMAX function with additive centered Gaussian noise with variance $\sigma^2 \leq n^C$. Then outside a rare event holding with probability $n^{-\omega(1)}$, the following runtime estimates are true.*

- *If $b = \Theta(1/h(n))$, then the expected runtime is $O(h(n)\sigma^4 n \log^2 n)$.*
- *If $b = O(1) \cap \Omega(1/\log^2 n)$, then the expected runtime is $O(h(n)\sigma^4 n \log^3 n)$.*

Proof By Theorem 8, we have Assumption (L') satisfied with $\tilde{\mu} = h(\mu)\sigma^2\sqrt{n}\ln(n)$, $\mu^+ = n^{2C}$, $T = O(\sigma^2\sqrt{n}\log n)$, and $p = 1 - o(1)$. Consequently, any $b \in n^{-o(1)} \cap n^{o(1)}$ gives that $\ell' = \min\{\ell \mid 2U^{\ell-1} \geq \tilde{\mu}, B_\ell \geq 2U^{\ell-1}T\} = (1 \pm o(1))(\frac{1}{2}\log_U(n) + \log_U(\sigma^2)) \leq (1 + o(1))C \log_U(n)$. Since $\mu^+ = n^{2C}$, we have $|\mathcal{L}| \geq (1 \pm o(1))C \log_U(n)$. With Theorem 3, we have proven our claim. ■

We remark that the parallel-run cGA has an expected runtime of $O(h(n)\sigma^4 n \log^3 n)$ outside a rare event of probability $n^{-\omega(n)}$.

5. Experimental Results

In this section, we experimentally analyze the smart-restart mechanism proposed in this work. We concentrate on the smart-restart cGA, since the cGA is one of the best-understood

EDAs and since it has only a single parameter which governs exactly how strong the update of the probabilistic model is in each iteration. We regard PBIL as a more complex EDA in Section 6.

Since, apart from the analysis of the cGA on ONEMAX by Lengler et al. (2021), such data is not yet available, we start with an investigation of how the runtime of the original cGA depends on the (hypothetical) population size μ . This will in particular support the assumption, underlying our smart-restart strategy and the parallel-run strategy from (Doerr, 2021), that the runtime can be excessively large when μ is below some threshold, and moderate and linearly increasing with μ when μ is larger than this threshold.

We then analyze the performance of the two existing approaches to automatically find good values for μ . Our focus is on understanding how one can relieve the user of an EDA from the difficult task of setting this parameter, not on finding the most efficient algorithm for the benchmark problems we regard. For this reason, we do not include other algorithms in this investigation. We note, though, that EDAs have shown a superior performance on the JUMP and DECEPTIVELEADINGBLOCKS benchmarks (Hasenöhrl and Sutton, 2018; Doerr, 2021; Wang et al., 2021), so clearly these are interesting algorithms for these two problems.

5.1 Experimental Settings

We ran the original cGA (with varying population sizes), the parallel-run cGA, and our smart-restart cGA (with two budget factors) on four benchmark problems, both without noise and in the presence of Gaussian posterior noise of four different strengths. For each experiment for the parallel-run cGA and our smart-restart cGA, we conducted 20 independent trials. Due to the often extremely large runtimes in the regime with genetic drift, only 10 independent trials were conducted for the original cGA. The detailed settings for our experiments were as follows.

- Benchmark functions and problem sizes: ONEMAX (problem size $n = 100$), LEADINGONES ($n = 50$), JUMP ($n = 50$ and jump size $k = 10$), and DECEPTIVELEADINGBLOCKS ($n = 30$). The population size for ONEMAX was chosen identical to the one used in (Friedrich et al., 2017), namely $n = 100$. For the other three problems, taking into account the longer runtimes, we chose relatively small problem sizes. Since for these our experimental results fit the known theoretical results (see Sections 2.2 and 4.2), we are confident that they are still representative.
- Noise model: additive centered Gaussian posterior noise with variances $\sigma^2 = \{0, n/2, n, 2n, 4n\}$ as described in Section 2.2.2.
- Termination criterion: Since the original cGA with unsuitable population sizes often did not find the optimum in a reasonable time, we imposed the following maximum numbers of generations and aborted the run after this number of generations: $\lceil n^4 \ln n \rceil$ for ONEMAX, n^5 for LEADINGONES, $n^{k/2}$ for JUMP, and $10n^5$ for DECEPTIVELEADINGBLOCKS. We did not define such a termination criterion for the parameter-less versions of the cGA since they always found the optimum in an affordable time.
- Population size of the original cGA: $\mu = 2^{\lceil 1..10 \rceil}$ for ONEMAX and LEADINGONES, $\mu = 2^{\lceil 9..18 \rceil}$ for JUMP, and $\mu = 2^{\lceil 1..14 \rceil}$ for DECEPTIVELEADINGBLOCKS. The reason

for omitting the range $\mu = 2^{[1..8]}$ for JUMP is the large runtime observed on this benchmark for small population sizes.

- Budget factor b for the smart-restart cGA: 16 and $1/\ln n$. As explained in the introduction, the budget factors $b = 16$ and $\Theta(1/\ln n)$ are two proper choices. We chose the constant 1 based on the experimental results on JUMP and DECEPTIVELEADING-BLOCKS without noise (noise variance $\sigma^2 = 0$) in Figures 3 and 4.
- Update factor U for the smart-restart cGA: 2. Doubling the parameter value after each unsuccessful run ($U = 2$) is a natural choice. We note that in (Doerr and Zheng, 2020a), we also did some experiments with $U = \sqrt{2}$, but these mostly gave inferior results.

5.2 Experimental Results and Analysis I: The cGA with Different Population Sizes

The curves in Figures 1–4 (excluding the three right-most points on the x -axis) show the runtime (number of fitness evaluations) of the original cGA with different population sizes when optimizing our four benchmarks under Gaussian noise with different variances (including the noise-free setting $\sigma^2 = 0$). Given are the median runtime together with the first and third quartiles. When a run was stopped because the maximum number of function evaluations was reached, we simply and bluntly counted the runtime up to this point as runtime. Clearly, there are better ways to handle such incomplete runs, but since a fair computation for these inefficient parameter ranges is not too important, we did not start a more elaborate evaluation. To ease the comparison, we also plotted the run budgets $B = b\mu^2$ which the smart-restart algorithm with budget factor b would have with population size μ .

The results displayed in Figures 1–4 typically show that the runtime of the cGA is roughly unimodal in the population size μ . For values of μ smaller than the optimal value, the runtime steeply increases and is accompanied by larger variances. For larger values of μ , we typically observe a moderate, roughly linear increase of the runtime. The variances are relatively small here.

Let us regard these results in some more detail. For ONEMAX, we see a perfect unimodal runtime behavior. The minima of the runtime curves are not strongly pronounced, but for population sizes smaller than the optimum one by a factor of four or more, drastic performance losses are observed. For larger population sizes, a roughly linear increase of the runtime is well visible. We also note that the optimal population size increases with the noise level, which fits the intuition that larger noise levels lead to larger runtimes, which need larger population sizes to prevent genetic drift throughout the runtime.

Our experiments do not show the bimodal runtime behavior observed in (Lengler et al., 2021). This is not surprising given that the bimodal pattern is very weak. In Figure 2 in (Lengler et al., 2021), the runtime pattern for relatively large problem size $n = 1000$ shows a global minimum at $\mu \approx 130$ (where the genetic drift is low as also shown in that figure). There is a second (local) minimum at around $\mu \approx 12$, but its runtime is only around 4% smaller than the runtime at the local maximum between the two minima. Given these small differences, seen for problem size $n = 1000$ in 3000 independent runs, it is not surprising that our experiments, conducted for smaller problem size and larger ranges of

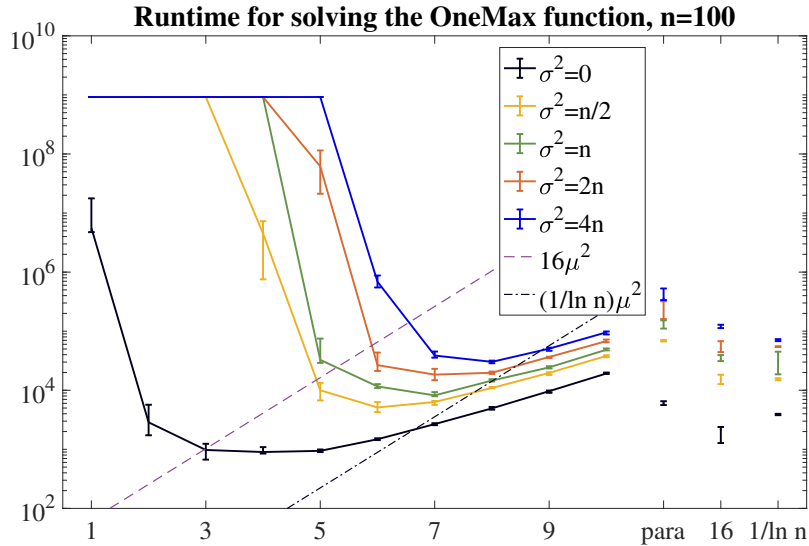


Figure 1: The median number of fitness evaluations (with the first and third quartiles) of the original cGA with different μ ($\log_2 \mu \in \{1, 2, \dots, 10\}$), the parallel-run cGA (“para”), and the smart-restart cGA with two budget factors ($b = 16$ and $b = 1/\ln n$) on the ONEMAX function ($n = 100$) under Gaussian noise with variances $\sigma^2 = 0, n/2, n, 2n, 4n$ in 20 independent runs (10 runs for the original cGA).

μ (resulting in much larger ranges of the runtimes), cannot detect this bimodal runtime behavior.

For LEADINGONES, we observe a more pronounced optimal value for μ . Reducing μ below this level leads to a clear increase of the runtime, typically at least by a factor of 10 for each halving of μ (which is still less drastic than for ONEMAX or JUMP functions). Interestingly, the optimal population size is relatively independent from the noise level (say compared to the ONEMAX results). We have no explanation for this.

For JUMP functions, the optimal μ -value is again less pronounced, however reducing the population size μ below the efficient values immediately gives a catastrophic increase of the runtime, almost always leading to all runs being stopped because the maximum number of generations is reached. In the very narrow transition regime between efficient and catastrophic optimization, we observe large variances of the runtime. This could indicate that there is not a continuous increase of the typical runtime, but rather an increase of the probability that a run enters an unfavorable situation, e.g., caused by genetic drift.

We recall that we ran the parameter-less versions of the algorithm without a termination criterion (since they always found the optimum), so it is for this reason that some of these runtimes appear larger than those for small static values of μ (which were just stopped after $n^{k/2} = 50^5 = 312,500,000$ iterations, i.e., $6.25 \cdot 10^8$ fitness evaluations, when not optimum was found before that).

The runtime behavior on DECEPTIVELEADINGBLOCKS is harder to understand. There is a clear “linear regime” from $\mu = 2^9$ or $\mu = 2^{10}$ on, again with very small variances. There is also a steep increase of the runtimes roughly starting at $\mu 2^4$. In between these

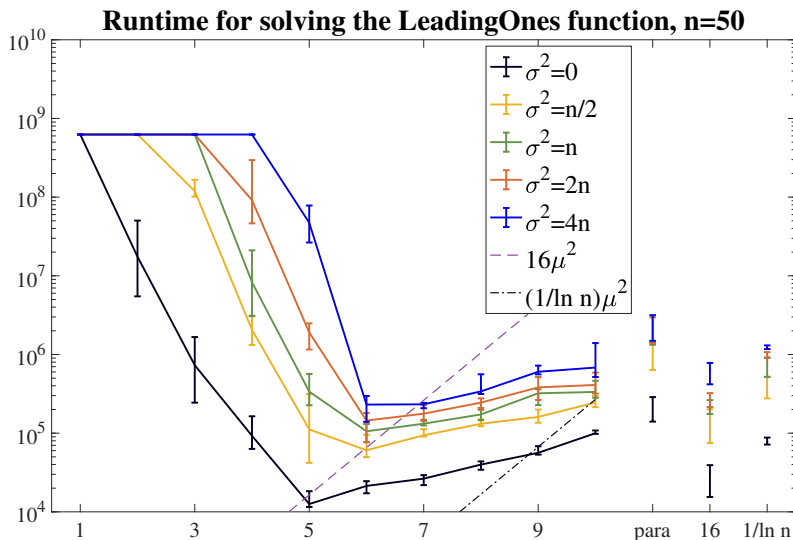


Figure 2: The median number of fitness evaluations (with the first and third quartiles) of the original cGA with different μ ($\log_2 \mu \in \{1, 2, \dots, 10\}$), the parallel-run cGA (“para”), and the smart-restart cGA with two budget factors ($b = 16$ and $b = 1/\ln n$) on the LEADING-ONES function ($n = 50$) under Gaussian noise with variances $\sigma^2 = 0, n/2, n, 2n, 4n$ in 20 independent runs (10 runs for the original cGA).

two regimes, the runtime behavior is hard to understand. The noisy runs show a small increase of the runtime in this middle regime together with slightly increased variances. The noise-free runs, however, are massively slower than the noisy ones, with large variances and a decent number of unsuccessful runs. We have no explanation for this.

Apart from the runtimes on DECEPTIVELEADINGBLOCKS (though to some extent also here, namely in the noisy runs), our results indicate a runtime behavior as described in Assumption (L). We have no proof for the fact that this behavior is caused by the effect of genetic drift and such a proof is most likely not easy to give. For this work, however, such a proof is not indispensable – what counts is that our understanding of genetic drift led to the development of the smart-restart scheme, which both in mathematical runtime analyses and in experiments showed a good performance and this is without the need to tune the hypothetical population size of the cGA (which is, in turn, indispensable when using static parameters as shown by our experiments).

As a side result, this data confirms that the cGA has a good performance on noise-free JUMP functions, not only in asymptotic terms as proven in (Hasenöhr and Sutton, 2018; Doerr, 2021), but also in terms of actual runtimes for concrete problem sizes. On a JUMP function with parameters $n = 50$ and $k = 10$, a classic mutation-based algorithm would run into the local optimum and from there would need to generate the global optimum via one mutation. For standard bit mutation with mutation rate $\frac{1}{n}$, this last step would take an expected time of $n^k \left(\frac{n}{n-1}\right)^{n-k}$, which for our values of n and k is approximately $2.2 \cdot 10^{17}$. With the asymptotically optimal mutation rate of $\frac{k}{n}$ determined in (Doerr et al., 2017), this

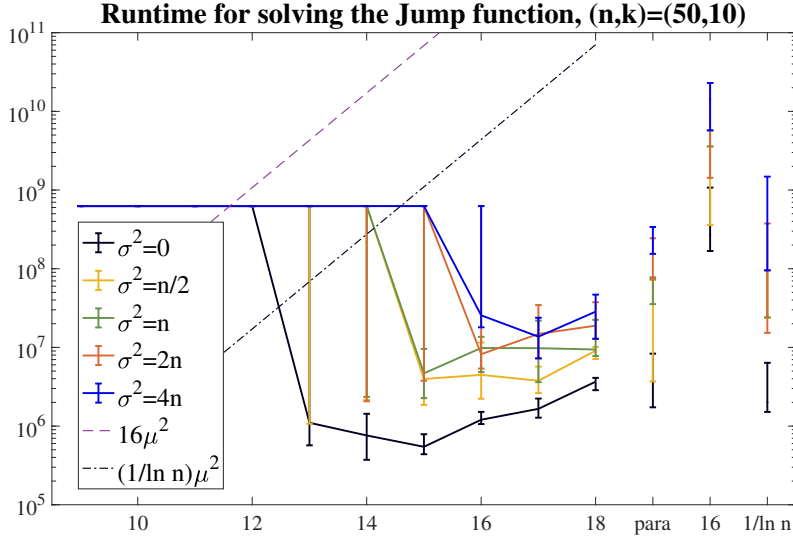


Figure 3: The median number of fitness evaluations (with the first and third quartiles) of the original cGA with different μ ($\log_2 \mu \in \{9, 10, \dots, 18\}$), the parallel-run cGA (“para”), and the smart-restart cGA with two budget factors ($b = 16$ and $b = 1/\ln n$) on the JUMP function with $(n, k) = (50, 10)$ under Gaussian noise with variances $\sigma^2 = 0, n/2, n, 2n, 4n$ in 20 independent runs (10 runs for the original cGA).

time would still be approximately $7.3 \cdot 10^{10}$. In contrast, the median optimization time of the cGA with $\mu \in 2^{[15..18]}$ is always below $4 \cdot 10^6$.

Our data also indicates that a good performance of the cGA can often be obtained with much smaller population sizes (and thus more efficiently) than what previous theoretical works suggest. For example, in (Friedrich et al., 2017) a population size of $\omega(\sigma^2 \sqrt{n} \log n)$ was required for the optimization of a noisy ONEMAX function via the cGA. In their experiments on a noisy ONEMAX function with $n = 100$ and $\sigma^2 = n$, a population size (called K in (Friedrich et al., 2017) to be consistent with some previous works) of $\mu = 7\sigma^2 \sqrt{n} (\ln n)^2 \approx 148,000$ was used, which led to a runtime of approximately 200,000 (data point for $\sigma^2 = 100$ interpolated from the two existing data points for $\sigma^2 = 64$ and $\sigma^2 = 128$ in the left chart of Figure 1 in (Friedrich et al., 2017)). In contrast, our experiments displayed in Figure 1 suggest that population sizes between 64 and 256 are already well sufficient and give runtimes clearly below 20,000.

We have to admit that we do not fully understand this number 200,000 from (Friedrich et al., 2017) and expect that it should be much larger. Our skepticism is based both on theoretical and experimental considerations. On the theoretical side, we note that even in the absence of noise and with the frequency vector having the (for this purpose) ideal value $\tau = (\frac{1}{2}, \dots, \frac{1}{2})$, the sum $\|\tau\|_1$ of the frequency values increases by an expected value of $O(\frac{1}{\mu} \sqrt{n})$ only (with small leading constant; an absolute upper bound of $\frac{1}{2\mu} \sqrt{n}$ follows, e.g., easily from (Berend and Kontorovich, 2013)). Hence after only 200,000 iterations, the frequency sum $\|\tau\|_1$ should still be relatively close to $n/2$. Since the probability to sample the optimum is $\prod_{i=1}^n (1 - \tau_i) \leq \exp(-\|\tau\|_1)$, it appears unlikely that the optimum

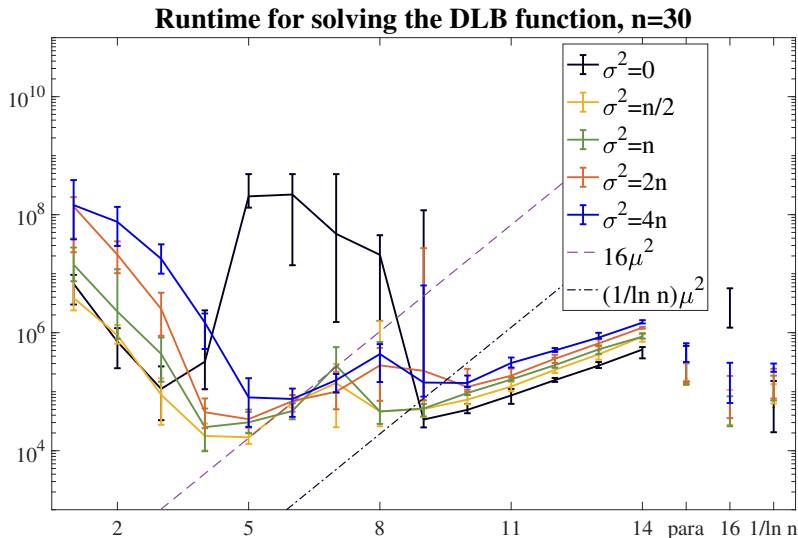


Figure 4: The median number of fitness evaluations (with the first and third quartiles) of the original cGA with different μ ($\log_2 \mu \in \{1, 2, \dots, 14\}$), the parallel-run cGA (“para”), and the smart-restart cGA with two budget factors ($b = 8$ and $b = 0.5/\ln n$) on the DECEPTIVELEADINGBLOCKS function ($n = 30$) under Gaussian noise with variances $\sigma^2 = 0, n/2, n, 2n, 4n$ in 20 independent runs (10 runs for the original cGA).

is sampled within that short time. Our experimental data displayed in Figure 1 suggests an affine-linear dependence of the runtime on μ when μ is at least 2^8 . From the median runtimes for $\mu = 2^9$ and $\mu = 2^{10}$, which are $T_9 = 24,384$ and $T_{10} = 48,562$, we would thus estimate a runtime of $T(\mu) = T_9 + (T_{10} - T_9)(\mu - 2^9)2^{-9}$ for $\mu \geq 2^{10}$, in particular, $T(7\sigma^2\sqrt{n}(\ln n)^2) = 7,010,551$ for the data point $\sigma^2 = 100$ and $n = 100$. To resolve this discrepancy, we conducted 20 runs of the cGA with $\mu = \lceil 7\sigma^2\sqrt{n}(\ln n)^2 + \frac{1}{2} \rceil$, $\sigma^2 = 100$, $n = 100$ and observed a median runtime of 5,728,969 (and a low variance, in fact, all 20 runtimes were in the interval $[5,042,714; 6,131,522]$). So most likely, the number of 200,000 given in (Friedrich et al., 2017) is not correct, and the price for the large value of μ is significantly larger than what the 200,000 suggests.

5.3 Experimental Results and Analysis II: Runtimes of the Parallel-Run cGA and the Smart-Restart cGA

The three right-most items on the x -axis in Figures 1–4 show the runtimes of the parallel-run cGA and the smart-restart cGA (with two budget factors b). Figures 1–4 also plot the two budgets (number of fitness evaluations) $16\mu^2$ and $(1/\ln n)\mu^2$ corresponding to $b = 16$ and $1/\ln n$ respectively.

The intersection point of the runtime curve of the cGA and the budget curve is a good indication for the μ -value with which the smart-restart cGA finds the optimum, and thus it is also a good indication for the runtime of the smart restart cGA. For example, the smallest μ such that the $16\mu^2$ curve is above the noiseless LEADINGONES curve is $\mu = 2^5$. Consequently,

we expect the smart-restart cGA with budget factor 16 to not find the optimum of the noiseless LEADINGONES functions in the runs with $\mu = 2, 4, 8, 16$, but only in the run with $\mu = 32$. For this reason, the runtime of this smart-restart cGA should be equal to the runtime of the original cGA with $\mu = 32$ plus $16 \cdot 2^2 + 16 \cdot 4^2 + 16 \cdot 8^2 + 16 \cdot 16^2 = 5440$, which fits roughly to our experimental data.

We note that for most runtime curves of the classic cGA, the intersection points with the budget curves are in the linear regime, which means that the corresponding smart-restart cGA avoids spending much time in the inefficient genetic drift regime. Some intersection points, e.g., those for the $(1/\ln(n))\mu^2$ budget in the LEADINGONES figure, are far in the linear regime. This indicates that the corresponding smart restart cGA misses the better (smaller) μ -values that are already outside the genetic drift regime. As the curves for the original cGA and the performances of the smart-restart cGA show, the performance loss of this miss is not too large. It is clearly much less than the catastrophic performance loss from running the cGA in the regime with strong genetic drift.

In more detail, we see that for the easy functions ONEMAX and LEADINGONES under all noise assumptions, the smart-restart cGA with both values of b has a smaller runtime than the parallel-run cGA. This can be explained from the runtime data of the original cGA in the corresponding figures: Since the runtimes are similar for several population sizes, the parallel-run cGA with its strategy to assign a similar budget to different population sizes wastes computational power, which the smart-restart cGA saves by aborting some processes early and not starting others. For both functions, the larger budget factor typically is superior. This fits again to the data on the original cGA, where we see the smaller budget factor curve intersecting the runtime curve clearly in the linear regime.

More interesting are the results for JUMP and DECEPTIVELEADINGBLOCKS. We recall that here a wrong choice of the population size can be catastrophic, so these are the two functions where not having to choose the population size is a big advantage for the user. What is clearly visible from the data is that here the smaller budget factor is preferable for the smart-restart cGA. This fits our previously gained intuition that for these two functions, genetic drift is detrimental. Hence there is no gain from continuing a run that is suffering from genetic drift (we note that there is no way to detect genetic drift on the fly – a frequency can be at a (wrong) boundary value due to genetic drift or at a (correct) boundary value because of a sufficiently strong fitness signal).

What is clear as a general rule is that both algorithms, the parallel-run cGA and the smart-restart cGA with the small fitness evaluation budget factor, clearly do a good job in successfully running the cGA with a reasonable population size – recall that for both of the difficult functions, a wrong choice of the population size can easily imply that the cGA does not find the optimum in 10^8 iterations.

6. Smart-Restart PBIL (Cross-Entropy Algorithm)

To see how smart-restart EDAs perform on combinatorial optimization problems and also to discuss a third EDA in this work, we now conduct an experimental analysis of *population-based incremental learning (PBIL)* (cross-entropy algorithm) on two optimization problems it was applied to in the literature. For comparison, two other restart strategies originally designed for evolutionary algorithms are also adapted to PBIL and implemented.

6.1 Smart-Restart Population-Based Incremental Learning (Smart-Restart Cross-Entropy)

Besides the cGA and UMDA, in (Doerr and Zheng, 2020b) also a theoretical analysis of the boundary hitting time caused by genetic drift in the algorithm PBIL (Baluja, 1994; Baluja and Caruana, 1995) was conducted. This algorithm is identical to the basic version of the cross-entropy (CE) algorithm for discrete optimization (Costa et al., 2007; De Boer et al., 2005).³ It further includes the UMDA (Mühlenbein and Paass, 1996) and the λ -max-min ant system (λ -MMAS) (Stützle and Hoos, 2000), a classic ant colony optimization algorithm, as special cases.

The general procedure of PBIL is to sample λ individuals and select μ best individuals to learn the current probabilistic model with learning rate ρ . An alternative parameterization, which we shall also prefer to ease the comparison with previous works, is to have as parameters (besides the learning rate ρ) the sample size λ and the selection pressure η , which define μ via $\mu = \lceil \eta\lambda \rceil$. Similar to the other EDAs regarded in this work, we also use the artificial margins $\{1/n, 1 - 1/n\}$ to prevent a premature convergence. The pseudocode of this algorithm is given in Algorithm 5.

Algorithm 5 The algorithm PBIL (with learning rate ρ , sample size λ , and selection pressure η) to maximize a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$

```

1:  $p^0 = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2}) \in [0, 1]^n$ 
2: for  $g = 1, 2, \dots$  do
   %%Sample  $\lambda$  individuals  $X_1^g, \dots, X_\lambda^g$ 
3:   for  $i = 1, 2, \dots, \lambda$  do
4:     for  $j = 1, 2, \dots, n$  do
5:        $X_{i,j}^g \leftarrow 1$  with probability  $p_j^{g-1}$  and  $X_{i,j}^g \leftarrow 0$  with probability  $1 - p_j^{g-1}$ 
6:     end for
7:   end for
   %%Update of the frequency vector
8:   Let  $\tilde{X}_1^g, \dots, \tilde{X}_\mu^g$  be the best  $\mu = \lceil \eta\lambda \rceil$  individuals (ties broken randomly)
9:    $p' = \frac{\rho}{\mu} \sum_{i=1}^{\mu} \tilde{X}_i^g + (1 - \rho)p^{g-1}$ 
10:   $p^g = \min\{\max\{\frac{1}{n}, p'\}, 1 - \frac{1}{n}\}$ 
11: end for

```

Recall that (Doerr and Zheng, 2020b, Theorem 3) proved that for PBIL, the frequency of a neutral bit moves out of the interval $(c\frac{\rho}{\mu}, 1 - c\frac{\rho}{\mu})$ for a constant $c \in (\frac{1}{2}, \frac{1}{\sqrt{2}})$ in at most $\frac{16}{2-1/c} \frac{\mu}{\rho^2}$ generations in expectation, that is, at most $\frac{16}{2-1/c} \frac{\mu\lambda}{\rho^2}$ fitness evaluations in expectation. With the notation of the selection pressure $\eta = \mu/\lambda$ and by Markov's inequality, the

3. We point out a possible tiny difference between PBIL and CE. According to the algorithm description of the CE algorithm in (Costa et al., 2007) and in the textbook (Rubinstein and Kroese, 2004, Algorithm 2.4.1), the CE algorithm selects all individuals with fitness at least the fitness of the μ -th best solution for the model update, whereas PBIL selects exactly μ best solutions breaking possible ties at random. However, the code provided in (Rubinstein and Kroese, 2004, Page 275) and also the recent pseudocode in (Wu et al., 2017, Algorithm 1) both have a fixed cutoff (and they do not discuss the problem of tie-breaking; more precisely, the tie-breaking is determined by how the sorting routine breaks the ties; since we are talking about random samples, it is clear anyway that the tie-breaking is not important). Given this state of the art, we prefer to think of the CE algorithm as also working with a fixed cutoff, and thus say that PBIL and CE are identical, as also said in (Krejca and Witt, 2020b).

probability that a boundary is reached in $b\lambda^2, b > \frac{16\eta}{(2-1/c)\rho^2}$, fitness evaluations, is at least $1 - \frac{16\eta}{(2-1/c)\rho^2 b}$. Hence, it fits into the framework of the smart-restart mechanism as discussed in Section 3. We thus obtain a smart-restart PBIL by letting the smart-restart mechanism control the parameter λ and keeping the parameters ρ and η fixed.

6.2 Other Restart Strategies

In Section 1, we mentioned several other generic strategies to remove the population size as a parameter of an algorithm. The two most interesting restart strategies among these shall be included in our experimental investigation, namely the classical strategy from (Harik and Lobo, 1999) (we only consider the first strategy of that work as the other strategy is a parallel-run strategy) and the well-cited strategy of Auger and Hansen (2005). Both restart strategies are not originally designed for an EDA, hence we need to adapt them to the PBIL.

The first strategy of Harik and Lobo (1999) restarts a crossover-based genetic algorithm when all individuals have become identical. As a reasonable analog for PBIL, we take the criterion that all entries of the probabilistic model p^g are at the boundaries or close to them. More precisely, when using the frequency margins $\{1/n, 1 - 1/n\}$, the restart criterion is that for all $i = 1, \dots, n$, we have $p_i^g \in \{1/n, 1 - 1/n\}$. For PBIL without margins, we conduct a restart when $p_i^g \in (0, 1/n^2) \cup (1 - 1/n^2, 1)$ for all $i = 1, \dots, n$. Note that we cannot set all $p_i^g \in \{0, 1\}$ as the criterion for PBIL without margins, since the frequencies in a run of PBIL never reach (but approach) 0 or 1 for sufficiently large n .

The strategy of Auger and Hansen (2005) contains five criteria to trigger a restart with a larger population size for the CMA-ES. Four of the five relate to the covariance matrix or the evolution path and thus are specific to CMA-ES. For these, we did not find a natural analog for EDAs. Hence in our analysis for the PBIL, we only discuss their criterion related to the fitness. Here a restart is triggered if the range of the best fitnesses among the last $L = 10 + \lceil 30n/\lambda \rceil$ iterations is zero or the range of these best fitnesses and the best fitnesses in the current iteration is below a predefined threshold ϵ . Algorithm 6 gives the pseudocode for these two restart strategies (adapted to PBIL). We use HL and AH as shorthands for the two strategies.

Algorithm 6 Two adapted restart mechanisms from (Harik and Lobo, 1999; Auger and Hansen, 2005) with update factor U applied to an algorithm \mathcal{A} with parameter λ for the maximization of a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

- 1: **for** $\ell = 1, 2, \dots$ **do**
 - 2: Run \mathcal{A} with parameter value $\lambda_\ell = 2U^{\ell-1}$ for arbitrarily long time on the maximization problem f until one of the following situations happens.
 - For HL, at a certain iteration g , for all $i = 1, \dots, n$, $p_i^g \in \{1/n, 1 - 1/n\}$ for \mathcal{A} with margins or $p_i^g \in (0, 1/n^2) \cup (1 - 1/n^2, 1)$ for \mathcal{A} without margins
 - For AH, at a certain iteration g , $\max V - \min V = 0$ or $\max V' - \min V' < \epsilon$ where $V = \{v^{g-L}, \dots, v^{g-1}\}$, $V' = V \cup V^g$, L is the predefined memory size, v^t is the best fitness value at generation t , and V^g is the samples at the generation g
 - 3: **end for**
-

We note a central difference between our restart strategy (Algorithm 4) and the two from Algorithm 6. Our smart-restart strategy stops each parameter trial when a prespecified computational budget (based on a general understanding of genetic drift in EDAs) is reached. In contrast, HL and AH try to detect on the fly when a situation is reached from which further progress appears difficult.

6.3 Experiments on Two Combinatorial Problems

We empirically test the smart-restart PBIL on two combinatorial problems, namely the max-cut problem and the bipartition problem in the settings used in the textbook (Rubinstein and Kroese, 2004). We recall that our focus is to understand how effective the smart-restart scheme is in finding an efficient value for the sample size λ . For that reason, we conduct experiments with the PBIL/CE algorithm as proposed in (Rubinstein and Kroese, 2004) and with our smart-restart version of it, but we do not regard other EDAs. We also regard the two other restart mechanisms discussed above.

6.3.1 OPTIMIZATION PROBLEMS

Now we briefly introduce the two problems regarded in this section. Both were used in the textbook (Rubinstein and Kroese, 2004) to demonstrate the power of the CE algorithm.

In the *max-cut problem*, the input consists of an undirected graph $G = (V, E)$ together with edge weights $w : E \rightarrow \mathbb{R}$. The target is to find a partition (V_1, V_2) of the node set V such that the sum of the weights of the edges from V_1 to V_2 is maximized. The max-cut problem is NP-complete and APX-hard. The best known approximation algorithm is a 0.878-approximation algorithm based on semi-definite programming by Goemans and Williamson (1995).

In the *bipartition problem*, we are given the same input data, but now in addition the sizes of V_1 and V_2 are prescribed. Again, the target is to maximize the sum of the weights of the edges from V_1 to V_2 . This problem is again NP-complete and APX-hard. An approximation algorithm with minimally weaker approximation ratio than the Goemans-Williamson algorithm for the max-cut problem was given by Austrin et al. (2016).

Both problems can easily be modeled as pseudo-Boolean optimization problems. For both, Rubinstein and Kroese (2004) propose synthetic problem instances with clear structures so that the unique optimal solution is known in advance. The precise details are not important to understand the remainder, so we omit the details and refer the interested reader to (Rubinstein and Kroese, 2004, Pages 46-49) for the max-cut problem and (Rubinstein and Kroese, 2004, Pages 145-147) for the bipartition problem.

6.3.2 EXPERIMENTAL SETTINGS

In all our experiments, we use the problems as proposed and modeled in Rubinstein and Kroese (2004). In our implementation of the core PBIL, we use the Matlab code provided in (Rubinstein and Kroese, 2004, Pages 274-276). This code is formulated for the max-cut problem. For the bipartition problem, in addition we use our own implementation of the *random partition generation algorithm* (Rubinstein and Kroese, 2004, Algorithm 4.6.1) (since no code for this algorithm is given in Rubinstein and Kroese (2004)). This algorithm is used to sample a partition with prescribed sizes of the partition classes and given

marginal distributions, which is the only difference between the PBIL for the max-cut and the partition problem.

We note that the Matlab code in (Rubinstein and Kroese, 2004, Pages 274-276) does not use the artificial margins. We speculate that their absence did not create problems because of the relatively large population sizes used in their experiments. As discussed in Section 2.1.1, the artificial margins $\{1/n, 1-1/n\}$ are usually utilized to avoid the frequencies reaching the absorbing boundaries 0 or 1. To obtain a complete picture, we shall conduct experiments both with and without margins.

We use the following settings, again taken from (Rubinstein and Kroese, 2004).

- **Max-cut problem:** We use a problem size of $n = 400$. For the PBIL, the learning rate is $\rho = 1$, the selection pressure is $\eta = 0.1$, and the sample size is $\lambda = 1,000$. These are the settings from (Rubinstein and Kroese, 2004, Table 2.3), where we note that the notation there writes (α, ρ, N) where we use (ρ, η, λ) .
- **Bipartition problem:** We use a problem size of $n = 600$. For the algorithm parameters, we use $\rho = 0.7$, $\eta = 0.01$, and $\lambda = 6,000$ as in (Rubinstein and Kroese, 2004, Table 4.7).⁴

In the experiments with the smart-restart PBIL, we used the PBIL kernel as above (with the same values for ρ and η), and only the value of λ was set via the smart-restart mechanism. The following shows the hyperparameter choices for the smart-restart PBIL, which were identical for both optimization problems.

- **Update factor U :** We used the same (natural) value $U = 2$ as in our experiments with the smart-restart cGA.
- **Budget factor b :** We used the two factors $b = 96 \frac{\eta}{\rho^2}$ and $b = 6 \frac{\eta}{\rho^2 \ln n}$ (depending on the values of η and ρ as used in the PBIL kernel). That is, $b = 96 \cdot \frac{0.1}{1^2} = 9.6$ and $6 \frac{0.1}{1^2 \ln n} = 0.6 / \ln n$ as $\eta = 0.1$ and $\rho = 1$ for PBIL on the max-cut problem (Rubinstein and Kroese, 2004, Table 2.3), and $b = 96 \frac{0.01}{0.7^2} = \frac{96}{49}$ and $6 \frac{0.01}{0.7^2 \ln n} = \frac{6}{49 \ln n}$ as $\eta = 0.01$ and $\rho = 0.7$ for the bipartition problem (Rubinstein and Kroese, 2004, Table 4.7). The motivation for these choices is as follows. We recall that in our experiments on the smart-restart cGA, we chose the budget factor $b = 16$ so that the probability of detecting the genetic drift is at least $1/2$. We chose $b = 1/\ln n$ since the order $\Theta(1/\log n)$ allows a union bound over all n frequencies, and we chose the precise value $b = 1/\ln(n)$ based on preliminary experiments. For smart-restart PBIL, to ensure a detection probability of at least $1/2$, we set $b = \frac{32\eta}{(2-1/c)\rho^2}$ as discussed in Section 6.1. For the $c \in (1/2, 1/\sqrt{2})$, we chose $c = 3/5$ as roughly the middle point in this interval. This explains our choice $b = \frac{32\eta}{(2-5/3)\rho^2} = 96 \frac{\eta}{\rho^2}$. Since this first value of b is $6 \frac{\eta}{\rho^2}$ times the 16, the first value for smart-restart cGA, for reasons of comparability we set the second value of b also $6 \frac{\eta}{\rho^2}$ times $1/\ln n$ (the second value of b for the smart-restart cGA), that is, $6 \frac{\eta}{\rho^2 \ln n}$.

4. (Rubinstein and Kroese, 2004, Table 4.7) does not specify the value of ρ , so we used $\rho = 0.7$ as specified in the preceding table (Rubinstein and Kroese, 2004, Table 4.6).

All algorithms were terminated only when the optimum is found.

In the experiments with the two restart mechanisms discussed in Section 6.2, we used the same PBIL kernel with the same (η, ρ) and initial $\lambda = 2$ as in the experiments with the smart-restart PBIL. For the AH restart strategy, the memory size $L = 10 + \lceil 30n/\lambda \rceil$ and the threshold $\epsilon = 10^{-12}$ are set as in the original paper (Auger and Hansen, 2005) (their notation for ϵ is `TolFun`). Since in our experiments the runs with AH (when using frequency margins) took extremely long, in our experiments we terminate the algorithm the first time the number of fitness evaluations exceeds 1.5×10^7 for the max-cut problem, and 3×10^6 for the bipartition problem. We note that in our experiments, all 20 runs of our smart-restart PBIL (with or without margins) reached the optimum within these limits.

6.3.3 EXPERIMENTAL RESULTS AND ANALYSES

Figure 5 shows the runtimes (measured by the number of fitness evaluations) of the original PBIL with margins, our smart-restart PBIL with margins, and HL with margins, utilizing the settings described above. We did not plot AH with margins as it fails to reach the global optimum in all 20 runs. We easily see a good performance of our smart-restart PBIL. For example, in terms of the median runtime (the red line in each box), the smart-restart PBIL with both values for b has a smaller runtime than the PBIL with the population size suggested in (Rubinstein and Kroese, 2004) for both problems, and better than the HL for the max-cut problem. For the bipartition problem, the smart-restart PBIL with smaller b shows a clear superiority over the HL. We notice the large variances for $b = 96 \frac{\eta}{\rho^2}$ for the bipartition problem and $b = 6 \frac{\eta}{\rho^2 \ln n}$ for the max-cut problem, but we have no explanation for these. These variances, however, do not change our general impression that the smart-restart approach is generally preferable over using fixed parameters with the values suggested in (Rubinstein and Kroese, 2004).

To understand the influence of using frequency margins, we collect the median results among 20 independent runs of the different algorithms without margins in Table 1. For comparison, we also record the results from the runs with margins from Figure 5. From Table 1, we deduce that the (common) usage of the margins is usually beneficial also for this problem setting, with the exception of the AH strategy. For AH, a restart is triggered when for a certain time interval no or only small fitness changes are observed. Here, apparently, margins are not helpful as they prevent the samples from becoming too similar. We note that also without margins, our smart-restart PBIL succeeds in finding the optimum, albeit with larger runtimes. The runtime increase is most drastic for the larger budget factor as here more time is spent in a run, hence more time is lost, when a frequency due to genetic drift has reached the wrong boundary value.

When not using frequency margins, the ranking of the four algorithms is the same for both combinatorial problems: The smart-restart strategy with small b performs best, then HL, AH, and the smart-restart strategy with the large b . This is somewhat intuitive when recalling the ideas behind these strategies. The smart-restart strategy with small b tries to estimate the first time when some frequency reaches a boundary due to genetic drift. HL restarts when all frequencies have actually reached the boundaries. AH restarts when only small fitness changes happen in fixed-length intervals of iterations (where we recall that the original AH strategy contained four more criteria that were specific to the CMA-ES). The

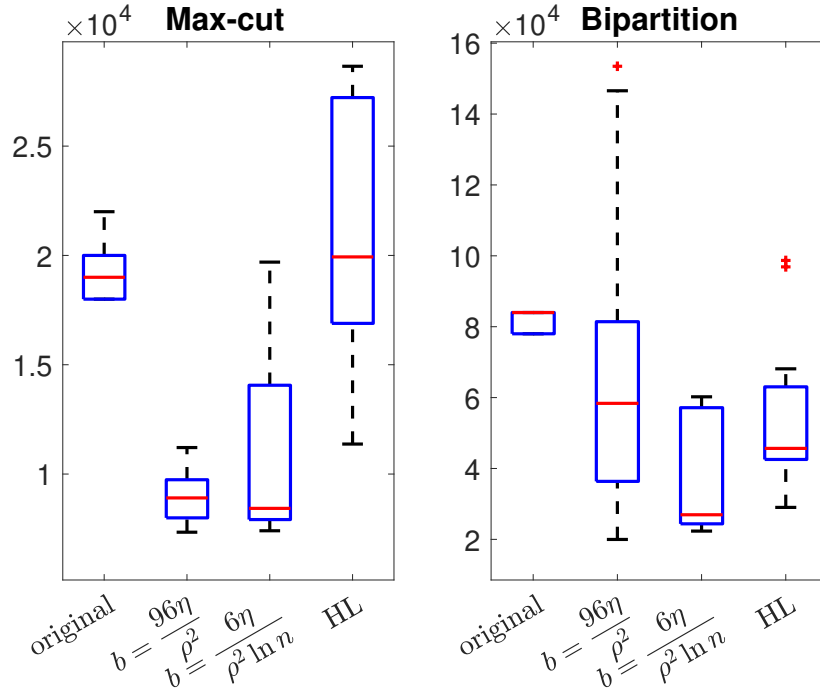


Figure 5: Boxplots of runtimes of the original PBIL (with margins), the smart-restart PBIL (with margins) with two budget factors ($b = 96\frac{\eta}{\rho^2}$ and $6\frac{\eta}{\rho^2 \ln n}$), and HL-restart PBIL (with margins) on the max-cut problem ($n = 400$) (Rubinstein and Kroese, 2004, Table 2.3) and bipartition problem ($n = 600$) (Rubinstein and Kroese, 2004, Table 4.7) in 20 independent runs.

smart-restart strategy with large b tries to estimate the time when a fixed frequency (hence also the typical frequency) reaches the boundaries. In this light, it is natural that the smart-restart strategy with a small budget factor performs better than the one with a large factor as it suffices that a single frequency is stuck at the wrong boundary to prevent finding the optimum. The fact that the HL strategy of checking whether all frequencies have reached a boundary performs relatively well, could suggest that the estimates of the smart-restart strategy with a small budget are still relatively conservative, and that possibly a restart could have been triggered even earlier. It is not totally surprising that the general fitness-dependent AH strategy has a harder stand than the strategies based on the mechanics of EDAs (but, as said, the original AH strategy was designed for the CMA-ES and included criteria exploiting the mechanics of the CMA-ES).

Similar to the results in Section 5, we observe that our smart-restart PBIL with the smaller budget clearly achieves very good runtimes, no matter with or without margins.

7. Conclusion

Choosing the parameters that control the genetic drift of estimation-of-distribution algorithms is one of the key difficulties for the practical usage of EDAs. To overcome this

Table 1: The median runtime for original, smart-restart, HL-restart, and AH-restart PBIL with and without margins on the max-cut and bipartition problems in 20 independent runs. The minimal median runtimes are in the bold font. 5 of 20 runs for the original PBIL without margins on the max-cut failed to find the optimum due to reaching the wrong boundaries. All 20 runs of the AH-restart PBIL with margins cannot reach the global optimum within the maximal number of the function evaluations.

Problem	Margins	PBIL	$b = 96 \frac{\eta}{\rho^2}$	$b = 6 \frac{\eta}{\rho^2 \ln n}$	HL	AH
Max-cut	With	19,000	8,908	8,430	19,932	$\geq 1.5 \times 10^7$
	Without	20,000	849,292	19,182	20,268	119,900
Bipartition	With	84,000	58,376	26,942	45,665	$\geq 3 \times 10^6$
	Without	84,000	702,344	26,430	61,838	221,161

difficulty, we proposed a smart-restart mechanism that removes this parameter from the algorithm. Our mechanism is a simple restart strategy with exponentially growing population size, but different from previous works it sets a prior fitness evaluation budget for each population size based on a recent quantitative analysis estimating when genetic drift is likely to occur.

Under a reasonable assumption on how the runtime depends on the population size, we theoretically analyzed our scheme and observed that it can lead to asymptotically optimal runtimes for the cGA and the UMDA.

Via extensive experiments on ONEMAX, LEADINGONES, JUMP, and DECEPTIVELEADINGBLOCKS, we showed the efficiency of the smart-restart cGA, also when compared with the parallel-run cGA. The results for the original cGA with different population sizes experimentally show that the population size is crucial for the performance of the cGA and that the theoretically suggested population size can be far away from the right one.

We also applied our smart-restart mechanism to the PBIL on two combinatorial problems. Our experiments showed again the efficiency of the smart-restart PBIL compared to the PBIL with the original settings and other restart strategies.

The problem of how to cope with genetic drift, naturally, is equally interesting for multivariate EDAs such as (Bonet et al., 1996; Pelikan and Mühlenbein, 1999; Mühlenbein and Mahnig, 1999; Harik et al., 2006; Probst and Rothlauf, 2020). For these, however, our theoretical understanding is limited to very few results such as (Zhang and Mühlenbein, 2004; Lehre and Nguyen, 2019; Doerr and Krejca, 2023). In particular, a quantitative understanding of genetic drift comparable to (Doerr and Zheng, 2020b) is completely missing. Another interesting question is if dynamic choices of the population size in EDAs can be fruitful. In classic EAs, dynamic parameter choices have recently been used very successfully to overcome the difficulty of finding a suitable static parameter value, see, e.g., the survey (Doerr and Doerr, 2020). How to use such ideas for EDAs is currently not at all clear.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. 62306086), Science, Technology and Innovation Commission of Shenzhen Municipality (Grant No. GXWD20220818191018001), Guangdong Basic and Applied Basic Research Foundation (Grant No. 2019A1515110177).

This work was also supported by a public grant as part of the Investissement d’avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH.

This work has profited from many scientific discussions at the Dagstuhl Seminar 22081 “Theory of Randomized Optimization Heuristics” and the Dagstuhl Seminar 22182 “Estimation-of-Distribution Algorithms: Theory and Applications”.

References

- Denis Antipov and Benjamin Doerr. A tight runtime analysis for the $(\mu + \lambda)$ EA. *Algorithmica*, 83:1054–1095, 2021.
- Denis Antipov, Benjamin Doerr, and Vitalii Karavaev. A rigorous runtime analysis of the $(1 + (\lambda, \lambda))$ GA on jump functions. *Algorithmica*, 84:1573–1602, 2022.
- Anne Auger and Nikolaus Hansen. A restart CMA evolution strategy with increasing population size. In *IEEE Congress on Evolutionary Computation, CEC 2005*, volume 2, pages 1769–1776. IEEE, 2005.
- Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better balance by being biased: a 0.8776-approximation for max bisection. *ACM Transactions on Algorithms*, 13, 2016.
- Shumeet Baluja. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Technical report, Carnegie Mellon University, 1994.
- Shumeet Baluja and Rich Caruana. Removing the genetics from the standard genetic algorithm. In *International Conference on Machine Learning, ICML 1995*, pages 38–46. Morgan Kaufmann, 1995.
- Daniel Berend and Aryeh Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics and Probability Letters*, 83:1254–1259, 2013.
- Jeremy S. De Bonet, Charles Lee Isbell Jr., and Paul A. Viola. MIMIC: finding optima by estimating probability densities. In *Advances in Neural Information Processing Systems, NIPS 1996*, pages 424–430. MIT Press, 1996.
- 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.
- Jürgen Branke, Clemens Lode, and Jonathan L. Shapiro. Addressing sampling errors and diversity loss in UMDA. In *Genetic and Evolutionary Computation Conference, GECCO 2007*, pages 508–515. ACM, 2007.

- Andre Costa, Owen Dafydd Jones, and Dirk Kroese. Convergence properties of the cross-entropy method for discrete optimization. *Operations Research Letters*, 35:573–580, 2007.
- Duc-Cuong Dang, Tobias Friedrich, Timo Kötzing, Martin S. Krejca, Per Kristian Lehre, Pietro S. Oliveto, Dirk Sudholt, and Andrew M. Sutton. Escaping local optima with diversity mechanisms and crossover. In *Genetic and Evolutionary Computation Conference, GECCO 2016*, pages 645–652. ACM, 2016.
- 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.
- Pieter-Tjerk De Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinstein. A tutorial on the cross-entropy method. *Annals of Operations Research*, 134:19–67, 2005.
- Benjamin Doerr. Analyzing randomized search heuristics via stochastic domination. *Theoretical Computer Science*, 773:115–137, 2019.
- Benjamin Doerr. The runtime of the compact genetic algorithm on Jump functions. *Algorithmica*, 83:3059–3107, 2021.
- Benjamin Doerr. Does comma selection help to cope with local optima? *Algorithmica*, 84:1659–1693, 2022.
- Benjamin Doerr and Carola Doerr. Theory of parameter control for discrete black-box optimization: provable performance gains through dynamic parameter choices. In Benjamin Doerr and Frank Neumann, editors, *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 271–321. Springer, 2020. Also available at <https://arxiv.org/abs/1804.05650>.
- Benjamin Doerr and Martin S. Krejca. Significance-based estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24:1025–1034, 2020.
- Benjamin Doerr and Martin S. Krejca. The univariate marginal distribution algorithm copes well with deception and epistasis. *Evolutionary Computation*, 29:543–563, 2021a.
- Benjamin Doerr and Martin S. Krejca. A simplified run time analysis of the univariate marginal distribution algorithm on LeadingOnes. *Theoretical Computer Science*, 851:121–128, 2021b.
- Benjamin Doerr and Martin S. Krejca. Bivariate estimation-of-distribution algorithms can find an exponential number of optima. *Theoretical Computer Science*, 971:114074, 2023.
- 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.
- 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, 2020a.

- Benjamin Doerr and Weijie Zheng. Sharp bounds for genetic drift in estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24:1140–1149, 2020b.
- 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.
- Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. Fast genetic algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 777–784. ACM, 2017.
- Benjamin Doerr, Carola Doerr, and Johannes Lengler. Self-adjusting mutation rates with provably optimal success rules. *Algorithmica*, 83:3108–3147, 2021.
- Stefan Droste. A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing*, 5:257–283, 2006.
- Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.
- 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.
- 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.
- Josselin Garnier, Leila Kallel, and Marc Schoenauer. Rigorous hitting times for binary mutations. *Evolutionary Computation*, 7:173–203, 1999.
- Christian Gießen and Timo Kötzing. Robustness of populations in stochastic environments. *Algorithmica*, 75:462–489, 2016.
- Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995.
- Brian W. Goldman and William F. Punch. Parameter-less population pyramid. In *Genetic and Evolutionary Computation Conference, GECCO 2014*, pages 785–792. ACM, 2014.
- Georges R. Harik and Fernando G. Lobo. A parameter-less genetic algorithm. In *Genetic and Evolutionary Computation Conference, GECCO 1999*, pages 258–265. Morgan Kaufmann, 1999.
- Georges R. Harik, Fernando G. Lobo, and David E. Goldberg. The compact genetic algorithm. *IEEE Transactions on Evolutionary Computation*, 3:287–297, 1999.

- Georges R. Harik, Fernando G. Lobo, and Kumara Sastry. Linkage learning via probabilistic modeling in the extended compact genetic algorithm (ECGA). In *Scalable Optimization via Probabilistic Modeling*, pages 39–61. Springer, 2006.
- Václav Hasenöhr 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.
- 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.
- 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, 2020a. Also available at <https://arxiv.org/abs/1806.05392>.
- 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, 2020b.
- Pedro Larrañaga and José Antonio Lozano, editors. *Estimation of Distribution Algorithms*. Springer, 2002.
- 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.
- Johannes Lengler, Dirk Sudholt, and Carsten Witt. The complex parameter landscape of the compact genetic algorithm. *Algorithmica*, 83:1096–1137, 2021.
- Cláudio F. Lima and Fernando G. Lobo. Parameter-less optimization with the extended compact genetic algorithm and iterated local search. In *Genetic and Evolutionary Computation Conference, GECCO 2004*, pages 1328–1339. Springer, 2004.
- Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. Simple hyper-heuristics control the neighbourhood size of randomised local search optimally for LeadingOnes. *Evolutionary Computation*, 28:437–461, 2020.
- Heinz Mühlenbein. How genetic algorithms really work: mutation and hillclimbing. In *Parallel Problem Solving from Nature, PPSN 1992*, pages 15–26. Elsevier, 1992.
- Heinz Mühlenbein and Thilo Mahnig. FDA – a scalable evolutionary algorithm for the optimization of additively decomposed functions. *Evolutionary Computation*, 7:353–376, 1999.
- 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.

- Pietro S. Oliveto, Dirk Sudholt, and Carsten Witt. Tight bounds on the expected runtime of a standard steady state genetic algorithm. *Algorithmica*, 84:1603–1658, 2022.
- Martin Pelikan and Tz-Kai Lin. Parameter-less hierarchical BOA. In *Genetic and Evolutionary Computation Conference, GECCO 2004*, pages 24–35. Springer, 2004.
- Martin Pelikan and Heinz Mühlenbein. The bivariate marginal distribution algorithm. In *Advances in Soft Computing*, pages 521–535. Springer, 1999.
- 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.
- Malte Probst and Franz Rothlauf. Harmless overfitting: Using denoising autoencoders in estimation of distribution algorithms. *Journal of Machine Learning Research*, 21:1–31, 2020.
- Amirhossein Rajabi and Carsten Witt. Self-adjusting evolutionary algorithms for multimodal optimization. *Algorithmica*, 84:1694–1723, 2022.
- Jonathan E. Rowe and Aishwaryaprajna. The benefits and limitations of voting mechanisms in evolutionary optimisation. In *Foundations of Genetic Algorithms, FOGA 2019*, pages 34–42. ACM, 2019.
- 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.
- Reuven Y. Rubinstein and Dirk P. Kroese. *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*, volume 133. Springer, 2004.
- 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, 2002a.
- Jonathan L. Shapiro. Scaling of probability-based optimization algorithms. In *Advances in Neural Information Processing Systems, NIPS 2002*, pages 383–390. MIT Press, 2002b.
- Jonathan L. Shapiro. Drift and scaling in estimation of distribution algorithms. *Evolutionary Computing*, 13:99–123, 2005.
- Jonathan L. Shapiro. Diversity loss in general estimation of distribution algorithms. In *Parallel Problem Solving from Nature, PPSN 2006*, pages 92–101. Springer, 2006.
- Thomas Stützle and Holger H. Hoos. MAX-MIN ant system. *Future Generation Computer Systems*, 16:889–914, 2000.
- Dirk Sudholt. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 17:418–435, 2013.

- Dirk Sudholt. Analysing the robustness of evolutionary algorithms to noise: refined runtime bounds and an example where noise is beneficial. *Algorithmica*, 83:976–1011, 2021.
- Dirk Sudholt and Christian Thyssen. A simple ant colony optimizer for stochastic shortest path problems. *Algorithmica*, 64:643–672, 2012.
- 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.
- Shouda Wang, Weijie Zheng, and Benjamin Doerr. Choosing the right algorithm with hints from complexity theory. In *International Joint Conference on Artificial Intelligence, IJCAI 2021*, pages 1697–1703. ijcai.org, 2021.
- Darrell Whitley, Swetha Varadarajan, Rachel Hirsch, and Anirban Mukhopadhyay. Exploration and exploitation without mutation: solving the jump function in $\Theta(n)$ time. In *Parallel Problem Solving from Nature, PPSN 2018, Part II*, pages 55–66. Springer, 2018.
- Carsten Witt. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation*, 14:65–86, 2006.
- Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing*, 22:294–318, 2013.
- Carsten Witt. Upper bounds on the running time of the univariate marginal distribution algorithm on OneMax. *Algorithmica*, 81:632–667, 2019.
- Carsten Witt. How majority-vote crossover and estimation-of-distribution algorithms cope with fitness valleys. *Theoretical Computer Science*, 940:18–42, 2023.
- Zijun Wu, Michael Kolonko, and Rolf H. Möhring. Stochastic runtime analysis of the cross-entropy algorithm. *IEEE Transactions on Evolutionary Computation*, 21:616–628, 2017.
- Qingfu Zhang and Heinz Mühlenbein. On the convergence of a class of estimation of distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 8:127–136, 2004.