



HAL
open science

The Univariate Marginal Distribution Algorithm Copes Well With Deception and Epistasis

Benjamin Doerr, Martin S Krejca

► **To cite this version:**

Benjamin Doerr, Martin S Krejca. The Univariate Marginal Distribution Algorithm Copes Well With Deception and Epistasis. *Evolutionary Computation*, In press, 29 (4), pp.543–563. 10.1162/evco_a_00293 . hal-03277754

HAL Id: hal-03277754

<https://hal.science/hal-03277754v1>

Submitted on 5 Jul 2021

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

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

The Univariate Marginal Distribution Algorithm Copes Well With Deception and Epistasis

Benjamin Doerr

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

doerr@lix.polytechnique.fr

Martin S. Krejca

Sorbonne Université, CNRS, LIP6, Paris, France*

martin.krejca@lip.fr

Abstract

In their recent work, Lehre and Nguyen (FOGA 2019) show that the univariate marginal distribution algorithm (UMDA) needs time exponential in the parent populations size to optimize the DECEPTIVELEADINGBLOCKS (DLB) problem. They conclude from this result that univariate EDAs have difficulties with deception and epistasis.

In this work, we show that this negative finding is caused by the choice of the parameters of the UMDA. When the population sizes are chosen large enough to prevent genetic drift, then the UMDA optimizes the DLB problem with high probability with at most $\lambda(\frac{n}{2} + 2e \ln n)$ fitness evaluations. Since an offspring population size λ of order $n \log n$ can prevent genetic drift, the UMDA can solve the DLB problem with $O(n^2 \log n)$ fitness evaluations. In contrast, for classic evolutionary algorithms no better run time guarantee than $O(n^3)$ is known (which we prove to be tight for the $(1 + 1)$ EA), so our result rather suggests that the UMDA can cope well with deception and epistasis.

From a broader perspective, our result shows that the UMDA can cope better with local optima than many classic evolutionary algorithms; such a result was previously known only for the compact genetic algorithm. Together with the lower bound of Lehre and Nguyen, our result for the first time rigorously proves that running EDAs in the regime with genetic drift can lead to drastic performance losses.

Keywords

Estimation-of-distribution algorithm, univariate marginal distribution algorithm, run time analysis, epistasis, theory.

1 Introduction

Estimation-of-distribution algorithms (EDAs) are randomized search heuristics that evolve a probabilistic model of the search space in an iterative manner. Typically starting with the uniform distribution, an EDA takes samples from its current model and then adjusts it such that better solutions have a higher probability of being generated in the next iteration. This method of refinement leads to gradually better solutions and performs well on many practical problems, often outperforming competing approaches (Pelikan et al., 2015).

*Most part of this work was carried out at the Hasso Plattner Institute, University of Potsdam, Potsdam, Germany.

Theoretical analyses of EDAs also often suggest an advantage of EDAs when compared to evolutionary algorithms (EAs); for an in-depth survey of run time results for EDAs, please refer to the article by [Krejca and Witt \(2020b\)](#). With respect to simple unimodal functions, EDAs seem to be comparable to EAs. For example, [Sudholt and Witt \(2019\)](#) proved that the two EDAs cGA and 2-MMAS_{ib} have an expected run time of $\Theta(n \log n)$ on the standard theory benchmark function ONEMAX (assuming optimal parameter settings; n being the problem size), which is a run time that many EAs share. The same is true for the EDA UMDA, as shown by the results of [Krejca and Witt \(2020a\)](#), [Lehre and Nguyen \(2017\)](#), and [Witt \(2019\)](#). For the benchmark function LEADINGONES, [Dang and Lehre \(2015\)](#) proved an expected run time of $O(n^2)$ for the UMDA when setting the parameters right, which is, again, a common run time bound for EAs on this function. One result suggesting that EDAs can outperform EAs on unimodal functions was given by [Doerr and Krejca \(2020b\)](#). They proposed an EDA called sig-cGA, which has an expected run time of $O(n \log n)$ on both ONEMAX and LEADINGONES – a performance not known for any classic EA or EDA.

For the class of all linear functions, EDAs perform slightly worse than EAs. The classical $(1 + 1)$ evolutionary algorithm optimizes all linear functions in time $O(n \log n)$ in expectation ([Droste et al., 2002](#)). In contrast, the conjecture of [Droste \(2006\)](#) that the cGA does not perform equally well on all linear functions was recently proven by [Witt \(2018\)](#), who showed that the cGA has an $\Omega(n^2)$ expected run time on the binary-value function. We note that the binary-value function was found harder also for classical EAs. While the $(1 + \lambda)$ evolutionary algorithm optimizes ONEMAX with $\Theta(n \lambda \log \log \lambda / \log \lambda)$ fitness evaluations, it takes $\Theta(n \lambda)$ fitness evaluations for the binary-value functions ([Doerr and Künnemann, 2015](#)).

For the multimodal JUMP _{k} benchmark function, which has local optima with a Hamming distance of k away from the global optimum, EDAs seem to drastically outperform EAs. [Hasenöhrl and Sutton \(2018\)](#) recently proved that the cGA only has a run time of $\exp(O(k + \log n))$ with high probability. [Doerr \(2019b\)](#) proved that the cGA with high probability has a run time of $O(n \log n)$ on JUMP _{k} if $k < \frac{1}{20} \ln n$, meaning that the cGA is unfazed by the gap of k separating the local from the global optimum. In contrast, common EAs have run times such as $\Theta(n^k)$ ($(1 + 1)$ EA, [Droste et al., 2002](#)), $\Omega(n^k)$ ((μ, λ) EA, [Doerr, 2020a](#)), and $\Theta(n^{k-1})$ ($(\mu + 1)$ GA, [Dang et al., 2018](#)), and only go down to smaller run times such as $O(n \log n + kn + 4^k)$ by using crossover in combination with diversity mechanisms like island models ([Dang et al., 2016](#)).

Another result in favor of EDAs was given by [Chen et al. \(2009\)](#), who introduced the SUBSTRING function and proved that the UMDA optimizes it in polynomial time, whereas the $(1 + 1)$ evolutionary algorithm has an exponential run time, both with high probability. In the SUBSTRING function, only substrings of length αn , for $\alpha \in (0, 1)$, of the global optimum are relevant to the fitness of a solution, and these substrings provide a gradient to the optimum. In the process, the $(1 + 1)$ evolutionary algorithm loses bits that are not relevant anymore for following the gradient (but relevant for the optimum). The UMDA fixes its model for correct positions while it is following the gradient and thus does not lose these bits.

The first, and so far only, result to suggest that EDAs can be drastically worse than EAs was recently stated by [Lehre and Nguyen \(2019\)](#) via the DECEPTIVELEADINGBLOCKS function (DLB for short), which they introduce and which consists of blocks of size 2 that need to be solved sequentially. Each block is deceptive in the following sense: the values 10 and 01 have the worst fitness among all four possible values of a block. The value 00 has the second best fitness, and the value 11 has the best fitness. That is, the Hamming

distance to the optimal block value 11 is not monotone with respect to increasing fitness, and the value 00 represents a local optimum for a block. This is why DLB is considered deceptive.

Lehre and Nguyen (2019) prove that many common EAs optimize DLB within $O(n^3)$ fitness evaluations in expectation, whereas the UMDA has a run time of $e^{\Omega(\mu)}$ (where μ is an algorithm-specific parameter that often is chosen as a small power of n) for a large regime of parameters. Only for extreme parameter values $\lambda = \Omega(\mu^2)$ (where λ is another algorithm-specific parameter that is often chosen in a similar order of magnitude to μ), the authors prove an expected run time of $O(n^3)$ also for the UMDA.

In this paper, we prove that the UMDA is, in fact, able to optimize DLB in time $O(n^2 \log n)$ with high probability if its parameters are chosen more carefully (Theorem 3). Note that our result is better than any of the run times proven in the paper by Lehre and Nguyen (2019). We achieve this run time by choosing the parameters of the UMDA such that its model is unlikely to degenerate during the run time (Lemma 2). Here by *degenerate* we mean that the sampling frequencies approach the boundary values 0 and 1 where this is not justified by the objective function. This leads to a probabilistic model that is strongly concentrated around a single search point. This effect is often called *genetic drift* (Sudholt and Witt, 2019). While it appears natural to choose the parameters of an EDA as to prevent genetic drift, it also has been proven that genetic drift can lead to a complicated run time landscape and inferior performance (see Lengler et al. (2018) for the cGA).

In contrast to our setting, for their exponential lower bound, Lehre and Nguyen (2019) use parameters that lead to genetic drift. Once the probabilistic model is sufficiently degenerated, that is, the frequencies of a block in the first half are $O(1/n)$, the progress of the UMDA is so slow that even to leave the local optima of DLB (that have a better search point in Hamming distance two only), the EDA takes time exponential in μ .

Since the UMDA shows a good performance in the (more natural) regime without genetic drift and was shown inferior only in the regime with genetic drift, we disagree with the statement of Lehre and Nguyen (2019) that there are “inherent limitations of univariate EDAs against deception and epistasis”.

In addition to the improved run time, we derive our result using only tools commonly used in the analysis of EDAs and EAs, whereas the proof of the polynomial run time of $O(n^3)$ for the UMDA with uncommon parameter settings (Lehre and Nguyen, 2019) uses the level-based population method (Lehre, 2011; Dang and Lehre, 2016; Corus et al., 2018; Doerr and Kötzing, 2019), which is an advanced tool that can be hard to use. We are thus optimistic that our analysis method can be useful also in other run time analyses of EDAs.

We recall that the previous work (Lehre and Nguyen, 2019) only proved upper bounds for the run time of EAs on DLB, namely of order $O(n^3)$ unless overly large population sizes are used. To support our claim that the UMDA shows a better performance on DLB than EAs, we rigorously prove a lower bound of $\Omega(n^3)$ for the run time of the $(1+1)$ EA on DLB (Theorem 4). More precisely, we determine a precise expression for the expected run time of this algorithm on DLB, which is asymptotically equal to $(1 \pm o(1)) \frac{e-1}{4} n^3$. In addition, we prove that a run time of $\Omega(n^3)$ holds with overwhelming probability (Theorem 5).

Last, we complement our theoretical result with an empirical comparison of the UMDA to various other evolutionary algorithms. The outcome of these experiments suggests that the UMDA outperforms the competing approaches while also having a smaller variance (Figure 1). Further, we compare the UMDA to the EDA MIMIC (De

(Bonet et al., 1996), which is similar to the UMDA but uses a more sophisticated probabilistic model that is capable of capturing dependencies among bit positions. Our comparison shows that, for the same parameter regime of both algorithms, the UMDA and the MIMIC behave asymptotically equally, with the MIMIC being slightly faster. This, again, highlights that the UMDA is well suited to optimize DLB.

The remainder of this paper is structured as follows: in Section 2, we introduce our notation, formally define DLB and the UMDA, and we state the tools we use in our analysis. Section 3 contains our main result (Theorem 3) and discusses its proof informally before stating the different lemmas used to prove it. In Section 4, we conduct a tight run time analysis of the $(1 + 1)$ EA on DLB. In Section 5, we discuss our empirical results. Last, we conclude this paper in Section 6.

This paper extends our conference version (Doerr and Krejca, 2020c) in two major ways: (i) we prove lower bounds of $\Omega(n^3)$ for the run time of the $(1 + 1)$ EA on DLB, which hold in expectation (Theorem 4) and with overwhelming probability (Theorem 5). (ii) Our empirical analysis contains more EAs, the MIMIC (Figure 1), and also considers the impact of the UMDA’s population size (Figure 2).

2 Preliminaries

We are concerned with the run time analysis of algorithms optimizing pseudo-Boolean functions, that is, functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$, where $n \in \mathbb{N}^+$ denotes the dimension of the problem. Given a pseudo-Boolean function f and a bit string x , we refer to f as a *fitness function*, to x as an *individual*, and to $f(x)$ as the *fitness of x* .

For $n_1, n_2 \in \mathbb{N} := \{0, 1, 2, \dots\}$, we define $[n_1..n_2] = [n_1, n_2] \cap \mathbb{N}$, and for an $n \in \mathbb{N}$, we define $[n] = [1..n]$. From now on, if not stated otherwise, the variable n always denotes the problem size. For a vector x of length n , we denote its component at index $i \in [n]$ by x_i and, for an index set $I \subseteq [n]$, we denote the subvector of length $|I|$ consisting only of the components at indices in I by x_I . Further, let $|x|_1$ denote the number of 1s of x and $|x|_0$ its number of 0s.

DeceptiveLeadingBlocks.

The pseudo-Boolean function DECEPTIVELEADINGBLOCKS (abbreviated as DLB) was introduced by Lehre and Nguyen (2019) as a deceptive version of the well known benchmark function LEADINGONES. In DLB, an individual x of length n is divided into blocks of equal size 2. Each block consists of a trap, where the fitness of each block is determined by the number of 0s (minus 1), except that a block of all 1s has the best fitness of 2. The overall fitness of x is then determined by the longest prefix of blocks with fitness 2 plus the fitness of the following block. Note that in order for the chunking of DLB to make sense, it needs to hold that 2 divides n . In the following, we always assume this implicitly.

We now provide a formal definition of DLB. To this end, we first introduce the function DECEPTIVEBLOCK: $\{0, 1\}^2 \rightarrow [0..2]$ (abbreviated as DB), which determines the fitness of a block (of size 2). For all $x \in \{0, 1\}^2$, we have

$$\text{DB}(x) = \begin{cases} 2 & \text{if } |x|_1 = 2, \\ |x|_0 - 1 & \text{else.} \end{cases}$$

Further, we define the function PREFIX: $\{0, 1\}^n \rightarrow [0..n]$, which determines the longest prefix of x with blocks of fitness 2. For a logic formula P , let $[P]$ denote the Iverson

bracket, that is, $[P]$ is 1 if P is true and 0 otherwise. We define, for all $x \in \{0, 1\}^n$,

$$\text{PREFIX}(x) = \sum_{i=1}^{n/2} [\forall j \leq i: \text{DB}(x_{\{2i-1, 2i\}}) = 2].$$

DLB is now defined as follows for all $x \in \{0, 1\}^n$:

$$\text{DLB}(x) = \begin{cases} n & \text{if PREFIX}(x) = n, \\ \sum_{i=1}^{\text{PREFIX}(x)+1} \text{DB}(x_{\{2i-1, 2i\}}) & \text{else.} \end{cases}$$

The univariate marginal distribution algorithm.

Our algorithm of interest is the UMDA (Mühlenbein and Paaß, 1996; Algorithm 1) with parameters $\mu, \lambda \in \mathbb{N}^+$, $\mu \leq \lambda$. It maintains a vector p (frequency vector) of probabilities (frequencies) of length n as its probabilistic model. This vector is used to sample an individual $x \in \{0, 1\}^n$, which we denote as $x \sim \text{sample}(p)$, such that, for all $y \in \{0, 1\}^n$,

$$\Pr[x = y] = \prod_{\substack{i=1; \\ y_i=1}}^n p_i \prod_{\substack{i=1; \\ y_i=0}}^n (1 - p_i).$$

The UMDA updates this vector iteratively in the following way: first, λ individuals are sampled. Then, among these λ individuals, a subset of μ with the highest fitness is chosen (breaking ties uniformly at random), and, for each index $i \in [n]$, the frequency p_i is set to the relative number of 1s at position i among the μ best individuals. Last, if a frequency p_i is below $\frac{1}{n}$, it is increased to $\frac{1}{n}$, and, analogously, frequencies above $1 - \frac{1}{n}$ are set to $1 - \frac{1}{n}$. Capping into the interval $[\frac{1}{n}, 1 - \frac{1}{n}]$ circumvents frequencies from being stuck at the extremal values 0 or 1. Last, we denote the frequency vector of iteration $t \in \mathbb{N}$ with $p^{(t)}$.

Algorithm 1: The UMDA (Mühlenbein and Paaß, 1996) with parameters μ and λ , $\mu \leq \lambda$, maximizing a fitness function $f: \{0, 1\}^n \rightarrow \mathbb{R}$ with $n \geq 2$

```

1  $t \leftarrow 0$ ;
2  $p^{(t)} \leftarrow (\frac{1}{2})_{i \in [n]}$ ; // vector of length  $n$  with all entries being 1/2
3 repeat // iteration  $t$ 
4   for  $i \in [\lambda]$  do  $x^{(i)} \sim \text{sample}(p^{(t)})$ ;
5   let  $y^{(1)}, \dots, y^{(\mu)}$  denote the  $\mu$  best individuals out of  $x^{(1)}, \dots, x^{(\lambda)}$  (breaking
   ties uniformly at random);
6   for  $i \in [n]$  do  $p_i^{(t+1)} \leftarrow \frac{1}{\mu} \sum_{j=1}^{\mu} y_i^{(j)}$ ;
7   restrict  $p^{(t+1)}$  to the interval  $[\frac{1}{n}, 1 - \frac{1}{n}]$ ;
8    $t \leftarrow t + 1$ ;
9 until termination criterion met;

```

Run time analysis.

When analyzing the run time of the UMDA optimizing a fitness function f , we are interested in the number T of fitness function evaluations until an optimum of f is sampled for the first time. Since the UMDA is a randomized algorithm, this run time T is a random variable. Note that the run time of the UMDA is at most λ times the

number of iterations until an optimum is sampled for the first time, and it is at least λ times this number minus $(\lambda - 1)$.

In the area of run time analysis of randomized search heuristics, it is common to give bounds for the expected value of the run time of the algorithm under investigation. This is uncritical when the run time is concentrated around its expectation, as often observed for classical evolutionary algorithms. For EDAs, it has been argued, among others by Doerr (2019b), that it is preferable to give bounds that hold with high probability. This is what we shall aim at in this work as well. Of course, it would be even better to give estimates in a distributional sense, e.g., via stochastic domination by another distribution, as argued for by Doerr (2019a), but this appears to be difficult for EDAs, among others, because of the very different behavior in the regimes with and without strong genetic drift.

Probabilistic tools.

We use the following results in our analysis. In order to prove statements on random variables that hold with high probability, we use the following commonly known Chernoff bound.

Theorem 1 (Chernoff bound; Doerr, 2020b, Theorem 10.5, Hoeffding, 1963). *Let $k \in \mathbb{N}$, $\delta \in [0, 1]$, and let X be the sum of k independent random variables, each taking values in $[0, 1]$. Then*

$$\Pr[X \leq (1 - \delta)E[X]] \leq \exp\left(-\frac{\delta^2 E[X]}{2}\right).$$

The next lemma tells us that, for a random X following a binomial law, the probability of exceeding $E[X]$ is bounded from above by roughly the term with the highest probability.

Lemma 1 (Doerr, 2020b, Eq. (10.62)). *Let $k \in \mathbb{N}$, $p \in [0, 1]$, $X \sim \text{Bin}(k, p)$, and let $m \in [E[X] + 1..k]$. Then*

$$\Pr[X \geq m] \leq \frac{m(1-p)}{m - E[X]} \cdot \Pr[X = m].$$

We use Lemma 1 for the converse case, that is, in order to bound the probability that a binomially distributed random variable is smaller than its expected value.

Corollary 1. *Let $k \in \mathbb{N}$, $p \in [0, 1]$, $X \sim \text{Bin}(k, p)$, and let $m \in [0..E[X] - 1]$. Then*

$$\Pr[X \leq m] \leq \frac{(k-m)p}{E[X] - m} \cdot \Pr[X = m].$$

Proof. Let $\bar{X} := k - X$, and let $\bar{m} := k - m$. Note that $\bar{X} \sim \text{Bin}(k, 1 - p)$ with $E[\bar{X}] = k - E[X]$ and that $\bar{m} \in [E[\bar{X}] + 1..k]$. With Lemma 1, we compute

$$\Pr[X \leq m] = \Pr[\bar{X} \geq \bar{m}] \leq \frac{\bar{m}p}{\bar{m} - E[\bar{X}]} \cdot \Pr[\bar{X} = \bar{m}] = \frac{(k-m)p}{E[X] - m} \cdot \Pr[X = m],$$

which proves the claim. □

Last, the following theorem deals with a *neutral* bit in a fitness function f , that is, a position $i \in [n]$ such that bit values at i do not contribute to the fitness value at all. The following theorem by Doerr and Zheng (2020) states that if the UMDA optimizes such an f , then the frequency at position i stays close to its initial value $\frac{1}{2}$ for $\Omega(\mu)$ iterations. We go more into detail about how this relates to DLB at the beginning of Section 3.

Theorem 2 (Doerr and Zheng, 2020, Theorem 2). *Consider the UMDA optimizing a fitness function f with a neutral bit $i \in [n]$. Then, for all $d > 0$ and all $t \in \mathbb{N}$, we have*

$$\Pr[\forall t' \in [0..t]: |p_i^{(t')} - \frac{1}{2}| < d] \geq 1 - 2 \exp\left(-\frac{d^2 \mu}{2t}\right).$$

3 Run Time Analysis of the UMDA

In the following, we prove that the UMDA optimizes DLB efficiently, which is the following theorem.

Theorem 3. *Let $\delta, \varepsilon, \zeta \in (0, 1)$ be constants, and let $c_\mu = 16/\varepsilon^2$ and $c_\lambda = (1 - \zeta)(1 - \delta^2)(1 - \varepsilon)^4/(16e)$. Consider the UMDA optimizing DLB with $\mu \geq c_\mu n \ln n$ and $\mu/\lambda \leq c_\lambda$. Then the UMDA samples the optimum after $\lambda(\frac{n}{2} + 2e \ln n)$ fitness function evaluations with a probability of at least $1 - 9n^{-1}$.*

Before we present the proof, we sketch its main ideas and introduce important notation. We show that the frequencies of the UMDA are set to $1 - \frac{1}{n}$ block-wise from left to right with high probability. We formalize this concept by defining that a block $i \in [\frac{n}{2}]$ is *critical* (in iteration t) if and only if $p_{2i-1}^{(t)} + p_{2i}^{(t)} < 2 - \frac{2}{n}$ and, for each index $j \in [2i - 2]$, the frequency $p_j^{(t+1)}$ is at $1 - \frac{1}{n}$. Intuitively, a critical block is the first block whose frequencies are not at their maximum value. We prove that a critical block is optimized within a single iteration with high probability if we assume that its frequencies are not below $(1 - \varepsilon)/2$, for $\varepsilon \in (0, 1)$ being a constant.

In order to assure that the frequencies of each block are at least $(1 - \varepsilon)/2$ until it becomes critical, we show that most of the frequencies right of the critical block are not impacted by the fitness function. We call such frequencies *neutral*. More formally, a frequency p_i is neutral in iteration t if and only if the probability to have a 1 at position i in each of the μ *selected* individuals equals $p_i^{(t)}$. Note that since we assume that $\mu = \Omega(n \log n)$, by Theorem 2, the impact of the genetic drift on neutral frequencies is low with high probability.

We know which frequencies are neutral and which are not by the following key observation: consider a population of λ individuals of the UMDA during iteration t ; only the first (leftmost) block that has strictly fewer than μ 1s is relevant for selection, since the fitness of individuals that do not have a 11 in this block cannot be changed by bits to the right anymore. We call this block *selection-relevant*. Note that this is a notion that depends on the random offspring population in iteration t , whereas the notion *critical* depends only on $p^{(t)}$.

The consequences of a selection-relevant block are as follows: if block $i \in [\frac{n}{2}]$ is selection-relevant, then all frequencies in blocks left of i are set to $1 - \frac{1}{n}$, since there are at least μ individuals with 11s. All blocks right of i have *no* impact on the selection process: if an individual has no 11 in block i , its fitness is already fully determined by all of its bits up to block i by the definition of DLB. If an individual has a 11 in block i , it is definitely chosen during selection, since there are fewer than μ such individuals and since its fitness is better than that of all of the other individuals that do not have a 11 in block i . Thus, its bits at positions in blocks right of i are irrelevant for selection. Overall, since the bits in blocks right of i do not matter, the frequencies right of block i get no signal from the fitness function and are thus neutral (see Lemma 2).

Regarding block i itself, all of the individuals with 11s are chosen, since they have the best fitness. Nonetheless, individuals with a 00, 01, or 10 can also be chosen, where an individual with a 00 in block i is preferred, as a 00 has the second best fitness after

a 11. Since the fitness for a 10 or 01 is the same, selecting individuals with such blocks does not impact the number of 1s at the positions in block i in expectation. However, if more 00s than 11s are sampled for block i , it can happen that the frequencies of block i are decreased. Since we assume that $\mu = \Omega(n \log n)$, the frequency is sufficiently high before the update and the frequencies of block i do not decrease by much with high probability (see Lemma 3). Since, in the next iteration, block i is the critical block, it is then optimized within a single iteration (see Lemma 4), and we do not need to worry about its frequencies decreasing again.

Neutral frequencies.

We now prove that the frequencies right of the selection-relevant block do not decrease by too much within the first n iterations.

Lemma 2. *Let $\varepsilon \in (0, 1)$ be a constant. Consider the UMDA with $\lambda \geq \mu \geq (16n/\varepsilon^2) \ln n$ optimizing DLB. Let $t \leq n$ be the first iteration such that block $i \in [\frac{n}{2}]$ becomes selection-relevant for the first time. Then, with a probability of at least $1 - 2n^{-1}$, all frequencies at the positions $[2i + 1..n]$ are at least $(1 - \varepsilon)/2$ within the first t iterations.*

Proof. Let $j \in [2i + 1..n]$ denote the index of a frequency right of block i . Note that by the assumption that t is the first iteration such that block i becomes selection-relevant it follows that, for all $t' \leq t$, the frequency $p_j^{(t')}$ is neutral, as we discussed above.

Since $p_j^{(t')}$ is neutral for all $t' \leq t$, by Theorem 2 with $d = \frac{\varepsilon}{2}$, we see that the probability that p_j leaves the interval $((1 - \varepsilon)/2, (1 + \varepsilon)/2)$ within the first $t \leq n$ iterations is at most $2 \exp(-\varepsilon^2 \mu / (8t)) \leq 2 \exp(-\varepsilon^2 \mu / (8n)) \leq 2n^{-2}$, where we used our bound on μ .

Applying a union bound over all $n - 2i \leq n$ neutral frequencies yields that at least one frequency leaves the interval $((1 - \varepsilon)/2, (1 + \varepsilon)/2)$ within the first t iterations with a probability of at most $2n^{-1}$, as desired. \square

Update of the selection-relevant block.

As mentioned at the beginning of the section, while frequencies right of the selection-relevant block do not drop below $(1 - \varepsilon)/2$ with high probability (by Lemma 2), the frequencies of the selection-relevant block can drop below $(1 - \varepsilon)/2$, as the following example shows.

Example 1. *Consider the UMDA with $\mu = 0.05\lambda \geq c \ln n$, for a sufficiently large constant c , optimizing DLB. Consider an iteration t and assume that block $i = \frac{n}{2} - 1 - o(n)$ is critical. Assume that the frequencies in blocks i and $i + 1$ are all at $2/5$. Since the prefix of $2(i - 1)$ leading 1s is sampled with probability $(1 - \frac{1}{n})^{2(i-1)} \geq (1 - \frac{1}{n})^{n-1} \geq 1/e$, the offspring population in iteration t has roughly $((2/5)^2/e)\lambda \approx 0.058\lambda > \mu$ individuals with at least $2i$ leading 1s in expectation. By Theorem 1, this also holds with high probability. Thus, the frequencies in block i are set to $1 - \frac{1}{n}$ with high probability.*

The expected number of individuals with at least $2i + 2$ leading 1s is roughly $((2/5)^4/e)\lambda \approx 0.0095\lambda$, and the expected number of individuals with $2i$ leading 1s followed by a 00 is roughly $((2/5)^2 \cdot (3/5)^2/e)\lambda \approx 0.02\lambda$. In total, we expect approximately $0.0295\lambda < \mu$ individuals with $2i$ leading 1s followed by either a 11 or a 00. Again, by Theorem 1, these numbers occur with high probability. Note that this implies that block $i + 1$ is selection-relevant with high probability.

Consider block $i + 1$. For selection, we choose all 0.0295λ individuals with $2i$ leading 1s followed by either a 11 or a 00 (which are sampled with high probability). For the remaining $\mu - 0.0295\lambda = 0.0205\lambda$ selected individuals with $2i$ leading 1s, we expect

half of them, that is, 0.01025λ individuals to have a 10. Thus, with high probability, the frequency $p_{2i+1}^{(t+1)}$ is set to roughly $(0.0095 + 0.01025)\lambda/\mu = 0.395$, which is less than $0.4 = p_{2i+1}^{(t)}$. Thus, this frequency decreased.

The next lemma shows that such frequencies do not drop too low, however.

Lemma 3. *Let $\varepsilon, \delta \in (0, 1)$ be constants, and let c be a sufficiently large constant. Consider the UMDA with $\lambda \geq \mu \geq c \ln n$ optimizing DLB. Further, consider an iteration t such that block $i \in [2.. \frac{n}{2}]$ is selection-relevant, and assume that its frequencies $p_{2i-1}^{(t)}$ and $p_{2i}^{(t)}$ are at least $(1 - \varepsilon)/2$ when sampling the population. Then the frequencies $p_{2i-1}^{(t+1)}$ and $p_{2i}^{(t+1)}$ are at least $(1 - \delta)(1 - \varepsilon)^2/4$ with a probability of at least $1 - 4n^{-2}$.*

Proof. Let for the moment i be arbitrary. Let k denote the number of individuals with a prefix of at least $2i - 2$ leading 1s and let X be the random variable that counts how many of these have 1s also in positions $2i - 1$ and i . Then X follows a binomial law with k trials and with a success probability of $p_{2i-1}^{(t)}p_{2i}^{(t)} =: \tilde{p} \geq (1 - \varepsilon)^2/4$. Since in the following we condition on block i being selection-relevant, it follows that $k \geq \mu$ and $X < \mu$. We now bound the probability that at least $(1 - \delta)\tilde{p}\mu =: m$ individuals have $2i$ leading 1s, that is, we bound $\Pr[X \geq m \mid X < \mu]$.

Elementary calculations show that

$$\begin{aligned} \Pr[X \geq m \mid X < \mu] &= 1 - \Pr[X < m \mid X < \mu] = 1 - \frac{\Pr[X < m, X < \mu]}{\Pr[X < \mu]} \\ &= 1 - \frac{\Pr[X < m]}{\Pr[X < \mu]}. \end{aligned} \quad (1)$$

To show a lower bound for (1), consider separately the two cases that $\mathbb{E}[X] < \mu$ and $\mathbb{E}[X] \geq \mu$.

Case 1: $\mathbb{E}[X] < \mu$. We first bound the numerator of the subtrahend in (1). Since $m/(1 - \delta) = \tilde{p}\mu \leq \tilde{p}k = \mathbb{E}[X]$, we have $\Pr[X < m] \leq \Pr[X < (1 - \delta)\mathbb{E}[X]]$. By [Theorem 1](#), by $\mathbb{E}[X] \geq \tilde{p}\mu$, and by our assumption that $\mu \geq c \ln n$, choosing c sufficiently large, we have

$$\Pr[X < (1 - \delta)\mathbb{E}[X]] \leq \exp\left(-\frac{\delta^2\mathbb{E}[X]}{2}\right) \leq \exp\left(-\frac{\delta^2\tilde{p}\mu}{2}\right) \leq n^{-2}.$$

For bounding the denominator, we note that $\tilde{p} \leq 1 - \frac{1}{n}$ and use the fact that a binomially distributed random variable with a success probability of at most $1 - \frac{1}{n}$ is below its expectation with a probability of at least $\frac{1}{4}$ ([Doerr, 2020b](#), Lemma 10.20 (b)). This yields

$$\Pr[X < \mu] \geq \Pr[X < \mathbb{E}[X]] \geq \frac{1}{4}.$$

Combining these bounds, we obtain $\Pr[X \geq m \mid X < \mu] \geq 1 - 4n^{-2}$ for this case.

Case 2: $\mathbb{E}[X] \geq \mu > m$. We bound the subtrahend from (1) from above. By basic estimations and by [Corollary 1](#), we see that

$$\frac{\Pr[X < m]}{\Pr[X < \mu]} \leq \frac{\Pr[X \leq m - 1]}{\Pr[X = \mu - 1]} \leq \frac{(k - m + 1)\tilde{p}}{\mathbb{E}[X] - m + 1} \cdot \frac{\Pr[X = m - 1]}{\Pr[X = \mu - 1]}. \quad (2)$$

We bound the first factor of (2) as follows, recalling that $(1 - \delta)\tilde{p}\mu = m$ and noting that $m \geq 1$ for sufficiently large values of n :

$$\begin{aligned} \frac{(k - m + 1)\tilde{p}}{\mathbb{E}[X] - m + 1} &\leq \frac{\mathbb{E}[X]}{\mathbb{E}[X] - m} = 1 + \frac{m}{\mathbb{E}[X] - m} \leq 1 + \frac{m}{\mu - m} \leq 1 + \frac{m}{m/\tilde{p} - m} \\ &= 1 + \frac{\tilde{p}}{1 - \tilde{p}} \leq 1 + n - 1 = n, \end{aligned}$$

where the last inequality uses that $\tilde{p} \leq (1 - \frac{1}{n})^2 \leq 1 - \frac{1}{n}$.

For the second factor of (2), we compute

$$\begin{aligned} \frac{\Pr[X = m - 1]}{\Pr[X = \mu - 1]} &= \frac{\binom{k}{m-1}\tilde{p}^{m-1}(1 - \tilde{p})^{k-m+1}}{\binom{k}{\mu-1}\tilde{p}^{\mu-1}(1 - \tilde{p})^{k-\mu+1}} \\ &= \frac{(\mu - 1)!(k - \mu + 1)!}{(m - 1)!(k - m + 1)!} \cdot \left(\frac{1 - \tilde{p}}{\tilde{p}}\right)^{\mu - m}. \end{aligned} \quad (3)$$

Since $\tilde{p} \geq (1 - \varepsilon)^2/4$, we see that $(1 - \tilde{p})/\tilde{p} \leq 4/(1 - \varepsilon)^2$.

For the first factor of (3), let $p^* := (1 - \delta)\tilde{p}$, thus $\mu p^* = m$. Noting that, for all $a, b \in \mathbb{R}$ with $a < b$, the function $j \mapsto (a + j)(b - j)$ is maximal for $j = (b - a)/2$, we first bound

$$\begin{aligned} \frac{(\mu - 1)!}{(m - 1)!} &= \prod_{j=0}^{\mu-m-1} (\mu - 1 - j) = \left(\prod_{j=0}^{\lfloor(\mu-m-1)/2\rfloor} (\mu - 1 - j)\right) \left(\prod_{j=\lceil(\mu-m)/2\rceil}^{\mu-m-1} (\mu - 1 - j)\right) \\ &\leq \left(\prod_{j=0}^{\lfloor(\mu-m-1)/2\rfloor} (\mu - 1 - j)\right) \left(\prod_{j=0}^{\lfloor(\mu-m-1)/2\rfloor} (m + j)\right) \\ &= \prod_{j=0}^{\lfloor(\mu-m-1)/2\rfloor} ((\mu - 1 - j)(m + j)) \leq \left(\frac{\mu + m}{2}\right)^{\mu-m} \leq \left(\frac{\mu}{2}(1 + p^*)\right)^{\mu-m}. \end{aligned}$$

Substituting this into the first factor of (3), we bound

$$\begin{aligned} \frac{(\mu - 1)!(k - \mu + 1)!}{(m - 1)!(k - m + 1)!} &\leq \left(\frac{\mu}{2}(1 + p^*)\right)^{\mu-m} \cdot \frac{(k - \mu + 1)!}{(k - m + 1)!} \\ &= \frac{\left(\frac{\mu}{2}(1 + p^*)\right)^{\mu-m}}{\prod_{j=0}^{\mu-m-1} (k - m + 1 - j)} = \prod_{j=0}^{\mu-m-1} \frac{\mu(1 + p^*)}{2(k - m + 1 - j)}. \end{aligned}$$

By noting that $k\tilde{p} = \mathbb{E}[X] \geq \mu$, we bound the above estimate further:

$$\begin{aligned} \prod_{j=0}^{\mu-m-1} \frac{\mu(1 + p^*)}{2(k - m + 1 - j)} &\leq \prod_{j=0}^{\mu-m-1} \frac{\mu(1 + p^*)}{2(\mu/\tilde{p} - m + 1 - j)} \\ &\leq \left(\frac{\mu(1 + p^*)}{2\mu(1/\tilde{p} - 1) + 2}\right)^{\mu-m} \leq \left(\frac{\tilde{p}(1 + p^*)}{2(1 - \tilde{p})}\right)^{\mu-m}. \end{aligned}$$

Substituting both bounds into (3) and recalling that $m = \mu p^*$, we obtain

$$\frac{\Pr[X = m - 1]}{\Pr[X = \mu - 1]} \leq \left(\frac{1 + p^*}{2}\right)^{\mu(1-p^*)} = \exp\left(-\mu(1 - p^*) \ln\left(\frac{2}{1 + p^*}\right)\right).$$

Finally, substituting this back into our bound of (2), using our assumption that $\mu \geq c \ln n$ and noting that p^* is constant, choosing c sufficiently large, we obtain

$$\frac{\Pr[X < m]}{\Pr[X < \mu]} \leq n \exp\left(-\mu(1-p^*) \ln\left(\frac{2}{1+p^*}\right)\right) \leq n^{-2}.$$

Concluding the proof. In both cases, we see that the number of 1s in block i is at least $m = (1-\delta)\tilde{p}\mu \geq ((1-\delta)(1-\varepsilon)^2/4)\mu$ with a probability of at least $1-4n^{-2}$. Since each 11 contributes to the new values of p_{2i-1} and p_{2i} , after the update, both frequencies are at least $(1-\delta)(1-\varepsilon)^2/4$ with the same probability bound, as we claimed. \square

Optimizing the critical block.

Our next lemma considers the critical block $i \in [\frac{n}{2}]$ of an iteration t . It shows that, with high probability, for all $j \in [2i]$, we have that $p_j^{(t+1)} = 1 - \frac{1}{n}$. Informally, this means that (i) all frequencies left of the critical block remain at $1 - \frac{1}{n}$, and (ii) the frequencies of the critical block are increased to $1 - \frac{1}{n}$.

Lemma 4. *Let $\delta, \varepsilon, \zeta \in (0, 1)$ be constants and let $q = (1-\delta)^2(1-\varepsilon)^4/16$. Consider the UMDA optimizing DLB with $\lambda \geq (4/\zeta^2) \ln n$ and $\mu/\lambda \leq (1-\zeta)q/e$, and consider an iteration t such that block $i \in [\frac{n}{2}]$ is critical and that $p_{2i-1}^{(t)}$ and $p_{2i}^{(t)}$ are at least \sqrt{q} . Then, with a probability of at least $1-n^{-2}$, at least μ offspring are generated with at least $2i$ leading 1s. In other words, the selection-relevant block of iteration t is at a position in $[i+1.. \frac{n}{2}]$.*

Proof. Let X denote the number of individuals that have at least $2i$ leading 1s. Since block i is critical, each frequency at a position $j \in [2i-2]$ is at $1 - \frac{1}{n}$. Thus, the probability that all of these frequencies sample a 1 for a single individual is $(1 - \frac{1}{n})^{2i-2} \geq (1 - \frac{1}{n})^{n-1} \geq 1/e$. Further, since the frequencies $p_{2i-1}^{(t)}$ and $p_{2i}^{(t)}$ are at least \sqrt{q} , the probability to sample a 11 at these positions is at least q . Hence, we have $E[X] \geq q\lambda/e$.

We now apply [Theorem 1](#) to show that it is unlikely that fewer than μ individuals from the current iteration have fewer than $2i$ leading 1s. Using our bounds on μ and λ , we compute

$$\Pr[X < \mu] \leq \Pr\left[X \leq (1-\zeta)\frac{q}{e}\lambda\right] \leq \Pr[X \leq (1-\zeta)E[X]] \leq e^{-\frac{\zeta^2\lambda}{2}} \leq n^{-2}.$$

Thus, with a probability of at least $1-n^{-2}$, at least μ individuals have at least $2i$ leading 1s. This concludes the proof. \square

The run time of the UMDA on DLB.

We now prove our main result.

Proof of [Theorem 3](#). We prove that the UMDA samples the optimum after $\frac{n}{2} + 2e \ln n$ iterations with a probability of at least $1-9n^{-1}$. Since it samples λ individuals each iteration, the theorem follows.

Due to [Lemma 2](#) and $\mu \geq c_\mu n \ln n$, within the first n iterations, with a probability of at least $1-2n^{-1}$, no frequency drops below $(1-\varepsilon)/2$ while its block has not been selection-relevant yet.

By [Lemma 3](#), since $\mu = \omega(\log n)$, with a probability of at least $1-4n^{-2}$, once a block becomes selection-relevant, its frequencies do not drop below $(1-\delta)(1-\varepsilon)^2/4$ for the next iteration. By a union bound, this does not fail for n consecutive times with a

probability of at least $1 - 4n^{-1}$. Note that a selection-relevant block becomes critical in the next iteration.

Consider a critical block $i \in [\frac{n}{2}]$. By Lemma 4, since $\lambda = \omega(\log n)$, with a probability of at least $1 - n^{-2}$, all frequencies at positions in $[2i]$ are immediately set to $1 - \frac{1}{n}$ in the next iteration, and the selection-relevant block has an index of at least $i + 1$, thus, moving to the right. Applying a union bound for the first n iterations of the UMDA and noting that each frequency belongs to a selection-relevant block at most once shows that all frequencies are at $1 - \frac{1}{n}$ after the first $\frac{n}{2}$ iterations, since each block contains two frequencies, and stay there for at least $\frac{n}{2}$ additional iterations with a probability of at least $1 - 2n^{-1}$.

Consequently, after the first $\frac{n}{2}$ iterations, the optimum is sampled in each iteration with a probability of $(1 - \frac{1}{n})^n \geq 1/(2e)$. Thus, after $2e \ln n$ additional iterations, the optimum is sampled with a probability of at least $1 - (1 - 1/(2e))^{2e \ln n} \geq 1 - n^{-1}$.

Overall, by applying a union bound over all failure probabilities above, the UMDA needs at most $\frac{n}{2} + 2e \ln n$ iterations to sample the optimum with a probability of at least $1 - 9n^{-1}$. \square

4 Run Time Analysis of the (1 + 1) EA

In order to see how well the (1 + 1) EA compares to the UMDA on DLB, we prove a precise run time in the order of $\Omega(n^3)$ (see Theorem 4). As our result for the UMDA only proves a run time bound that holds with high probability, for a fairer comparison, we also prove that the (1 + 1) EA needs $\Omega(n^3)$ with overwhelming probability when optimizing DLB (see Theorem 5).

Our bound of order $\Omega(n^3)$ shows that the corresponding $O(n^3)$ bound from Lehre and Nguyen (2019) is tight apart from constant factors and lower order terms. We note that Lehre and Nguyen (2019) have also shown upper bounds for the run time of other evolutionary algorithms, again typically of order $O(n^3)$. We conjecture that for these setting an $\Omega(n^3)$ lower bound is valid as well, as our empirical results in Section 5 suggest, but we do not prove this here.

To make the following result precise, we recall that the (1 + 1) EA for the maximization of $f: \{0, 1\}^n \rightarrow \mathbb{R}$ is the simple EA which (i) starts with an individual x chosen uniformly at random from $\{0, 1\}^n$ and then (ii) in each iteration creates from the current solution x an offspring y via flipping each bit independently with probability $\frac{1}{n}$ and replaces x by y if and only if $f(y) \geq f(x)$.

The following result determines precisely the run time of the (1 + 1) EA on DLB. The proof, not surprisingly, takes some ideas from the precise analyses of the run time of the (1 + 1) EA on LEADINGONES by Böttcher et al. (2010) and Doerr (2019a).

Theorem 4. *In expectation, the (1 + 1) EA samples the optimum of DLB after*

$$\frac{1}{4}n^2 \left(n - \frac{1}{2} \right) \frac{\left(1 + \frac{1}{n-1}\right)^n - 1}{1 + \frac{1}{2(n-1)}} = (1 + o(1)) \frac{e - 1}{4} n^3$$

fitness function evaluations.

Proof. Let $\ell \in [0..n/2 - 1]$ and $b \in \{0, 1\}$. Let $x \in \{0, 1\}^n$ such that $\text{PREFIX}(x) = \ell$ and $\text{DB}(x_{2\ell+1, 2\ell+2}) = b$, in other words, such that $x_i = 1$ for all $i \leq 2\ell$ and the contribution of the $\ell + 1$ -st block is $\text{DB}(x_{2\ell+1, 2\ell+2}) = b$. Consider a run of the (1 + 1) EA on DLB, started with the search point x instead of a random initial solution. Let X be the random variable describing the first time that a solution with PREFIX-value greater than ℓ is found.

We first observe that X is independent of x_i , $i > 2\ell + 2$. In the case that $b = 0$, also X is independent of whether $(x_{2\ell+1, 2\ell+2}) = (1, 0)$ or $(x_{2\ell+1, 2\ell+2}) = (0, 1)$. Hence we use the notation $X_{\ell, b} := X$ without ambiguity later in this proof.

We compute $E[X]$. If $b = 1$, then X follows a geometric law with success probability $p = (1 - \frac{1}{n})^{2\ell} n^{-2}$; hence

$$E[X_{\ell, 1}] = \frac{1}{p} = \left(1 - \frac{1}{n}\right)^{-2\ell} n^2.$$

If $b = 0$, in principle we could also precisely describe the distribution of X , but since this is more complicated and we only regard expected run times in this proof, we avoid this and take the more simple route to only determine the expectation. From x , the $(1 + 1)$ EA in one iteration finds a search point with PREFIX-value greater than ℓ with probability $(1 - \frac{1}{n})^{2\ell+1} n^{-1}$, since it has to flip the only 0 in the first $2\ell + 2$ positions to a 1. With the same probability, it finds a search point with (unchanged) PREFIX-value equal to ℓ and (higher) b -value 1. Otherwise it finds no true improvement and stays with x or an equally good search point. In summary, we have

$$\begin{aligned} E[X_{\ell, 0}] &= 1 + \left(1 - \frac{1}{n}\right)^{2\ell+1} n^{-1} \cdot 0 + \left(1 - \frac{1}{n}\right)^{2\ell+1} n^{-1} E[X_{\ell, 1}] \\ &\quad + \left(1 - 2\left(1 - \frac{1}{n}\right)^{2\ell+1} n^{-1}\right) E[X_{\ell, 0}], \end{aligned}$$

hence

$$E[X_{\ell, 0}] = \frac{1}{2} \left(1 - \frac{1}{n}\right)^{-2\ell-1} n + \frac{1}{2} E[X_{\ell, 1}].$$

Let T_ℓ denote the run time when starting with a random solution x with $\text{PREFIX}(x) \geq \ell$, that is, such that the first 2ℓ bits of x are all 1s and the remaining bits are random. Note that when during the subsequent optimization the prefix value increases, then this results in a random search point with prefix value $\ell + 1$. From the above, we know that

$$\begin{aligned} E[T_\ell] &= \frac{1}{4}(E[X_{\ell, 1}] + E[T_{\ell+1}]) + \frac{1}{2}(E[X_{\ell, 0}] + E[T_{\ell+1}]) + \frac{1}{4}(E[T_{\ell+1}]) \\ &= E[T_{\ell+1}] + \frac{1}{4}E[X_{\ell, 1}] + \frac{1}{2}E[X_{\ell, 0}]. \end{aligned}$$

Noting that $E[T_{n/2}] = 0$, we can write $E[T_\ell] = \sum_{j=\ell}^{n/2-1} (E[T_j] - E[T_{j+1}])$. In particular, the run time $T := T_0$ of the $(1 + 1)$ EA with random initialization satisfies

$$\begin{aligned} E[T] &= \sum_{\ell=0}^{n/2-1} \left(\frac{1}{2}E[X_{\ell, 0}] + \frac{1}{4}E[X_{\ell, 1}] \right) = \sum_{\ell=0}^{n/2-1} \left(\frac{1}{4} \left(1 - \frac{1}{n}\right)^{-2\ell-1} n + \frac{1}{2}E[X_{\ell, 1}] \right) \\ &= \sum_{\ell=0}^{n/2-1} \frac{1}{4} \left(1 - \frac{1}{n}\right)^{-2\ell-1} n + \sum_{\ell=0}^{n/2-1} \frac{1}{2} \left(1 - \frac{1}{n}\right)^{-2\ell} n^2. \end{aligned}$$

Noting that $\sum_{\ell=0}^{n/2-1} \left(1 - \frac{1}{n}\right)^{-2\ell} = \frac{1}{2}(n-1) \frac{(1 + \frac{1}{n-1})^{n-1}}{1 + \frac{1}{2(n-1)}}$, we obtain

$$E[T] = \frac{1}{4} n^2 \left(n - \frac{1}{2} \right) \frac{(1 + \frac{1}{n-1})^{n-1} - 1}{1 + \frac{1}{2(n-1)}}. \quad \square$$

The result above on the expected run time does not rule out that the typical run time of the $(1 + 1)$ EA is much better, e.g., that it is quadratic with high probability, we add a short proof showing that with overwhelming probability, the run time of the $(1 + 1)$ EA on DLB is at least cubic. We note that with deeper methods, namely martingale concentration inequalities allowing unbounded martingale differences, such as the ones used by [Fan et al. \(2015\)](#) or [Kötzing \(2016\)](#), one could show more precise statements including that already lower-order deviations from the expectation almost surely do not happen. However, we feel that such a result, while surely interesting, is not important enough to justify the effort. Therefore, we only show the following weaker statement.

Theorem 5. *The probability that the $(1 + 1)$ EA samples the optimum of DLB in at most $\frac{n^3}{16e}$ iterations, is at most $\exp(-\frac{n}{64e})$.*

Proof. Consider a run of the $(1 + 1)$ EA on DLB. Let $(x_t)_{t \in \mathbb{N}}$ be the sequence of search points generated in this run, that is, x_0 is the random initial search point and x_t is generated in iteration t by mutating the current-best search point (breaking ties towards later generated ones). Let $t_i := \min\{t \in \mathbb{N} \mid \text{PREFIX}(x_t) \geq i\}$ be the first time a search point with prefix value at least i is generated. Note that $T := t_{n/2}$ is the run time of the algorithm (when we do not count the evaluation of the initial individual) and $t_0 = 0$.

Similar to what is well-known about the optimization of the LEADINGONES benchmark, we observe that if $t = t_i$, then x_t is a search point with $2i$ leading 1s and the remaining bits independently and uniformly distributed in $\{0, 1\}$. Consequently, $d_i := t_{i+1} - t_i$ is independent of t_0, \dots, t_i . We note that with probability $1/4$, the search point x_t has 0s in positions $2i + 1$ and $2i + 2$. In this case, the parent individual in the remaining run of the algorithm agrees with this x_t in the first $2i + 2$ bit positions until in iteration t_{i+1} a better solution is found. Such a better solution is found with probability $p_i = (1 - 1/n)^{2i} n^{-2} \leq n^{-2}$. With probability $(1 - n^{-2})^{n^2 - 1} \geq 1/e$, this takes at least n^2 iterations. In summary, we see that, regardless of d_0, \dots, d_{i-1} , we have that d_i is at least n^2 with probability at least $1/(4e)$. This conditional independence is enough to apply Chernoff-type concentration inequalities on the indicator random variables X_i of the events $\{d_i \geq n^2\}$, see Lemma 11 in [Doerr and Johannsen \(2010\)](#) or Section 1.10.2 in [Doerr \(2020b\)](#). Consequently, for $X = \sum_{i=0}^{n/2-1} X_i$ the multiplicative Chernoff bound (Theorem 1) gives

$$\Pr[X \leq \frac{n}{16e}] = \Pr[X \leq (1 - \frac{1}{2})E[X]] \leq \exp(-\frac{1}{2}(\frac{1}{2})^2 \frac{n}{8e}) = \exp(-\frac{n}{64e}).$$

Since $X > n/(16e)$ implies that the run time $T = \sum_{i=0}^{n/2-1} d_i$ is larger than $n^3/(16e)$, we have shown our claim. \square

5 Experiments

In their paper, [Lehre and Nguyen \(2019\)](#) analyze the run time of many EAs on DLB. For an optimal choice of parameters, they prove an expected run time of $O(n^3)$ for all considered algorithms.

Since these are only upper bounds and since we showed in [Section 4](#) an $\Omega(n^3)$ lower bound only for the $(1 + 1)$ EA, it is not clear how well the other algorithms actually perform against the UMDA, which has a run time in the order of $n^2 \ln n$ ([Theorem 3](#)) for appropriate parameters. Thus, we provide some empirical results in [Figure 1](#) on how well these algorithms compare against each other.

Considered algorithms. Besides the UMDA, we analyze the run time of the following EAs: the $(1 + 1)$ EA, the (μ, λ) EA, and the (μ, λ) GA with uniform crossover, which are most of the EAs that [Lehre and Nguyen \(2019\)](#) consider in their paper. Lehre and Nguyen also analyze the $(1 + \lambda)$ EA and the $(\mu + 1)$ EA. However, in our preliminary experiments with $\mu = \lceil \ln n \rceil$ and $\lambda = \lceil \sqrt{n} \rceil$, they were always slower than the $(1 + 1)$ EA, so we do not include these algorithms in [Figure 1](#).

Further, we also depict the run time of the *mutual-information-maximizing input clustering* algorithm (MIMIC; [De Bonet et al., 1996](#)), which is one of the algorithms that [Lehre and Nguyen \(2019\)](#) also analyze empirically. The MIMIC is an EDA with a more sophisticated probabilistic model than the UMDA. This model is capable of capturing dependencies among the bit positions by storing a permutation π of indices and conditional probabilities. An individual is created bit by bit, following π . Each bit is sampled with respect to a conditional probability, depending on the bit sampled in the prior position. The MIMIC updates its model as follow: similar to the UMDA, the MIMIC samples each iteration λ individuals and selects the μ best. To update its richer probabilistic model, the MIMIC then searches for the position with the least (empirical) entropy, that is, the position that has the most 0s or 1s, and sets the new frequency equal to the number of 1s in the selected subpopulation. Iteratively, for all remaining positions, it determines the position with the lowest conditional (empirical) entropy and sets the conditional frequencies to the conditional numbers of 1s, where conditional is always with respect to the prior position in π . Since a precise description of this algorithm would need a significant amount of space, we refer the reader to a recent paper by [Doerr and Krejca \(2020a\)](#) for more details.

Parameter choice. For each of the EAs depicted in [Figure 1](#), we choose parameters such that the upper run time bound proven by [Lehre and Nguyen \(2019\)](#) is $O(n^3)$, that is, optimal. For both the (μ, λ) EA and the (μ, λ) GA, we choose $\mu = \lceil \ln n \rceil$ and $\lambda = 9\mu$. For the latter, we further choose uniform crossover and, in each iteration, execute it with probability $1/2$. Our parameter choices for the (μ, λ) EA stem from aiming for a constant ratio μ/λ as well as from the constraints $\lambda \geq c \ln n$ (for a sufficiently large constant c) and $\lambda \geq (1 + \delta)e^2\mu$ as stated by [Lehre and Nguyen \(2019\)](#), choosing $\delta = 0.1$. For the (μ, λ) GA, we choose the same values for μ and λ , as the algorithm uses the same selection mechanism as the (μ, λ) EA and as the values comply with the constraints stated by [Lehre and Nguyen \(2019\)](#). The crossover probability of $1/2$ is chosen as it is a constant probability that satisfies the constraint to be at most $1 - (1 + \delta)\epsilon\mu/\lambda$, where we choose $\delta = 0.1$ as before.

For the UMDA and the MIMIC, we choose $\mu = \lceil 3n \ln n \rceil$ and $\lambda = 12\mu$. Our parameter choice for the UMDA is based on a constant ratio λ/μ and on a preliminary search for the first integer factors for μ and λ such that the UMDA is successful in a reasonable time. For the MIMIC, we speculate that the similarities with the UMDA imply that these parameter values are suitable as well. In any case, with these parameter values all runs of these two algorithms were successful, that is, the optimum of DLB was found within $10n^3$ function evaluations.

Results. When interpreting the plots from [Figure 1](#), please note that both axes use a logarithmic scale. Thus, any polynomial is depicted as a linear function, where the slope depends on the exponent of the highest term. [Figure 1](#) clearly shows two separate regimes: The EAs with a larger slope and the EDAs with a smaller slope.

We see that the $(1 + 1)$ EA from $n \geq 200$ on performs the worst out of all algorithms. The EAs all have a very similar slope, which ranges, according to a regression to fit a

power law, from $\Theta(n^{2.9967})$ to $\Theta(n^{3.1})$. This indicated that all of these EAs have a run time of $\Theta(n^3)$, as we proved for the $(1+1)$ EA in [Theorem 4](#). In contrast, the UMDA and the MIMIC have a smaller slope, about $\Theta(n^{2.2957})$ and $\Theta(n^{2.3096})$, respectively, suggesting that our upper bound of $O(n^2 \log n)$ for the UMDA ([Theorem 3](#)) is close to its true run time. Further note that the variance of the EDAs is very small, suggesting that their run time holds with a higher concentration bound than what we prove for the UMDA in [Theorem 3](#).

Interestingly, the MIMIC behaves very similarly to the UMDA. However, this may be a result of the same parameter choice of both algorithms. [Lehre and Nguyen \(2019\)](#) also analyze the MIMIC empirically, with choices of μ in the orders of \sqrt{n} , $\sqrt{n} \log n$, and n . In their results, the run time of the MIMIC is slightly faster; for $n = 100$, the median of the empirical run time of the MIMIC as reported by [Lehre and Nguyen \(2019\)](#) is in the order of $2^{19} \approx 5.2 \cdot 10^5$. In [Figure 1](#), for the same value of n , the run time of the MIMIC is about $6 \cdot 10^5$, which is very close. This suggests that the parameter regime without strong genetic drift, which we consider in this paper for the UMDA, is not that important for the MIMIC. While the MIMIC seems to have a larger tolerable parameter regime, it is impressive that the UMDA has a similar performance to the MIMIC.*

5.1 Different Regimes of the UMDA's Population Size

[Figure 1](#) compares the run time of the UMDA to that of other algorithms. To this end, we chose the parameters in the regime without strong genetic drift, as proposed in [Theorem 3](#). However, [Lehre and Nguyen \(2019\)](#) prove a lower bound of $\exp(\Omega(\mu))$ for the UMDA on DLB, which can be lower than our upper bound of $O(n^2 \log n)$ if μ is sufficiently small. Thus, in [Figure 2](#), we analyze the impact of μ on the run time of the UMDA on DLB for $n = 300$.

Parameter choice. In each run, we choose $\lambda = 12\mu$, as we did before; μ ranges from 2^1 to 2^{12} in powers of 2. Note that our range for μ ends with $2^{12} = 4,096$, which is less than our choice of $\mu = \lceil 3n \ln n \rceil$ in the previous experiment, which results in a value of $\mu = 5,134$.

Results. In [Figure 2](#), we show the run times for those values of μ where at least one out of 100 runs was successful, that is, where the UMDA found the optimum of DLB within $10n^3$ function evaluations. There are two very distinct ranges of μ leading to successful runs. The first regime consists of the values 2, 4, and 8; note that the effect of genetic drift is strong in this regime. For each value of μ , the number of fitness evaluations is at least 10^7 , which is worse than the performance of the UMDA in [Figure 1](#). Interestingly, in this regime the UMDA is successful in every run. This suggests that the range of frequency values for the UMDA is coarse enough such that it is able to quickly change frequencies from $\frac{1}{n}$ to $1 - \frac{1}{n}$. If a frequency is at $\frac{1}{n}$, it takes some time until a 1 is sampled and selected, but if this occurs, the respective frequency has a decent chance of being increased to $1 - \frac{1}{n}$.

For values of μ between 2^4 to 2^9 , the range of frequency values increases, consequently leading to longer times for a frequency to be increased from $\frac{1}{n}$ to $1 - \frac{1}{n}$. The growth of the plot from the regime of smaller values of μ suggests a superpolynomial run time in μ for this behavior. Thus, it is not surprising that no run of the UMDA is successful for the medium values of μ .

*It is important to note that we only report the number of fitness evaluations in [Figure 1](#). When comparing the time to perform an update, the UMDA is faster, as the model update can be computed in time $\Theta(n)$, whereas the one for the MIMIC takes time $\Theta(n^2)$ due to the iterated search of the minimal conditional entropy.




Figure 1: A log-log plot depicting the number of fitness evaluations of various algorithms until the optimum of DLB is sampled for the first time against the input size n from 50 to 300 in steps of 50. For each value of n , 100 independent runs were started per algorithm. The lines depict the median of the 50 runs of an algorithm, and the shaded areas denote the center 50%.

Please refer to [Section 5](#) for more details.




Figure 2: A log-log plot depicting the number of fitness evaluations of the UMDA until the optimum of DLB for $n = 300$ is sampled for the first time against the parameter μ from 2^1 to 2^{12} , doubling the value in each step. For each value of n , 100 independent runs were started. We terminated each run after $10n^3 = 2.7 \cdot 10^8$ function evaluations if the UMDA did not find the optimum until then.

The number above each data point denotes the ratio of runs in which the UMDA actually found an optimum. The lines depict the median of the *successful* runs, and the shaded areas denote their center 50%.

Note that for the values of μ from 2^4 to 2^9 , no run was successful. Thus, no data points are depicted.

Please refer to [Section 5.1](#) for more details.

The second successful regime of μ consists of the values 2^{10} , 2^{11} , and 2^{12} . While not all runs are successful, in particular not those for the smallest value of μ , the trend of the curve seems to be linear, which suggests a polynomial run time of the UMDA (because of the log-log scale). Further, the run time is comparable to that of the UMDA seen in [Figure 1](#). Likely, starting with $\mu = 2^{10}$, the UMDA enters the regime where the effect of genetic drift starts to diminish. For $\mu = 2^{10}$, the effect is still rather large, and oftentimes frequencies reach the lower border $\frac{1}{n}$, however, not all of the time. With increasing μ , this happens less and less. For $\mu = 2^{12}$, almost all runs are successful. In comparison, in [Figure 1](#), *all* runs of the UMDA were successful, which uses a slightly larger value of μ than 2^{12} .

Overall, the results from [Figure 2](#) show an interesting transition from the run time of the UMDA on DLB when the effect of genetic drift vanishes.

6 Conclusion

We conducted a rigorous run time analysis of the UMDA on the DECEPTIVELEADING-BLOCKS function. It shows that the algorithm with the right parameter choice finds the optimum in $O(n^2 \log n)$ fitness evaluations with high probability ([Theorem 3](#)). This result shows that the lower bound by [Lehre and Nguyen \(2019\)](#), which is exponential in μ , is not due to the UMDA being ill-suited for coping with epistasis and deception, but rather due to the choice of the algorithm’s parameters. For several EAs, [Lehre and Nguyen \(2019\)](#) showed a run time bound of $O(n^3)$ on DLB. We proved a matching lower bound for the $(1 + 1)$ EA ([Section 4](#)) and conducted experiments which suggest that also other EAs perform worse than the UMDA on DLB ([Section 5](#)). In this light, our result suggests that the UMDA can handle epistasis and deception even better than many evolutionary algorithms and that the UMDA does so similar to a more complex EDA.

Our run time analysis holds for parameter regimes that prevent genetic drift. When comparing our run time with the one shown by [Lehre and Nguyen \(2019\)](#), we obtain a strong suggestion for running EDAs in regimes of low genetic drift. In contrast to the work of [Lengler et al. \(2018\)](#) that indicates moderate performance losses due to genetic drift, here we obtain the first fully rigorous proof of such a performance loss, and in addition one that is close to exponential in n (the $\exp(\Omega(\mu))$ lower bound of [Lehre and Nguyen \(2019\)](#) holds for μ up to $o(n)$). Our proven upper and lower bound also show that the UMDA has an advantage in coping with local optima compared to EAs. Such an observation has previously only been made for the compact genetic algorithm (when optimizing jump functions, see [Hasenöhrl and Sutton \(2018\)](#); [Doerr \(2019b\)](#)).

On the technical side, our result indicates that the regime of low genetic drift admits relatively simple and natural analyses of run times of EDAs, in contrast, e.g., to the level-based methods previously used in comparable analyses, e.g., by [Dang and Lehre \(2015\)](#) and [Lehre and Nguyen \(2019\)](#).

We conjecture that our result can be generalized to a version of the DLB function with a block size of $k \leq n$.

Acknowledgments

We thank the anonymous reviewers of this paper, who provided valuable feedback that improved the paper in various aspects.

This work was supported by COST Action CA15140, by the Paris Île-de-France Region, and by a public grant as part of the Investissements d’avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH.

References

- Böttcher, S., Doerr, B., and Neumann, F. (2010). Optimal fixed and adaptive mutation rates for the LeadingOnes problem. In *Proc. of PPSN '10*, pages 1–10. Springer.
- Chen, T., Lehre, P. K., Tang, K., and Yao, X. (2009). When is an estimation of distribution algorithm better than an evolutionary algorithm? In *Proc. of CEC '09*, pages 1470–1477.
- Corus, D., Dang, D., Eremeev, A. V., and Lehre, P. K. (2018). Level-based analysis of genetic algorithms and other search processes. *IEEE Transactions on Evolutionary Computation*, 22(5):707–719.
- Dang, D., Friedrich, T., Kötzing, T., Krejca, M. S., Lehre, P. K., Oliveto, P. S., Sudholt, D., and Sutton, A. M. (2016). Escaping local optima with diversity mechanisms and crossover. In *Proc. of GECCO '16*, pages 645–652.
- Dang, D., Friedrich, T., Kötzing, T., Krejca, M. S., Lehre, P. K., Oliveto, P. S., Sudholt, D., and Sutton, A. M. (2018). Escaping local optima using crossover with emergent diversity. *IEEE Transactions on Evolutionary Computation*, 22(3):484–497.
- Dang, D. and Lehre, P. K. (2015). Simplified runtime analysis of estimation of distribution algorithms. In *Proc. of GECCO '15*, pages 513–518.
- Dang, D. and Lehre, P. K. (2016). Runtime analysis of non-elitist populations: from classical optimisation to partial information. *Algorithmica*, 75(3):428–461.
- De Bonet, J. S., Isbell, Jr., C. L., and Viola, P. A. (1996). MIMIC: finding optima by estimating probability densities. In *Proc. of NIPS '96*, pages 424–430.
- Doerr, B. (2019a). Analyzing randomized search heuristics via stochastic domination. *Theoretical Computer Science*, 773:115–137.
- Doerr, B. (2019b). A tight runtime analysis for the cGA on jump functions: EDAs can cross fitness valleys at no extra cost. In *Proc. of GECCO '19*, pages 1488–1496.
- Doerr, B. (2020a). Does comma selection help to cope with local optima? In *Proc. of GECCO '20*, pages 1304–1313.
- Doerr, B. (2020b). Probabilistic tools for the analysis of randomized optimization heuristics. In *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 1–87. Springer. Also available at <https://arxiv.org/abs/1801.06733>.
- Doerr, B. and Johannsen, D. (2010). Edge-based representation beats vertex-based representation in shortest path problems. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 759–766.
- Doerr, B. and Künnemann, M. (2015). Optimizing linear functions with the $(1+\lambda)$ evolutionary algorithm – Different asymptotic runtimes for different instances. *Theoretical Computer Science*, 561:3–23.
- Doerr, B. and Kötzing, T. (2019). Multiplicative up-drift. In *Proc. of GECCO '19*, pages 1470–1478.

- Doerr, B. and Krejca, M. S. (2020a). Bivariate estimation-of-distribution algorithms can find an exponential number of optima. In *Proc. of GECCO '20*, pages 796–804.
- Doerr, B. and Krejca, M. S. (2020b). Significance-based estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24(6):1025–1034.
- Doerr, B. and Krejca, M. S. (2020c). The univariate marginal distribution algorithm copes well with deception and epistasis. In *Proc. of EvoCOP '20*, pages 51–66.
- Doerr, B. and Zheng, W. (2020). Sharp bounds for genetic drift in estimation-of-distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 24(6):1140–1149.
- Droste, S. (2006). A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing*, 5(3):257–283.
- Droste, S., Jansen, T., and Wegener, I. (2002). On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276(1-2):51–81.
- Fan, X., Grama, I., and Liu, Q. (2015). Exponential inequalities for martingales with applications. *Electronic Journal of Probability*, 20:1–22.
- Hasenöhrl, V. and Sutton, A. M. (2018). On the runtime dynamics of the compact genetic algorithm on jump functions. In *Proc. of GECCO '18*, pages 967–974.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Kötzing, T. (2016). Concentration of first hitting times under additive drift. *Algorithmica*, 75:490–506.
- Krejca, M. S. and Witt, C. (2020a). Lower bounds on the run time of the univariate marginal distribution algorithm on OneMax. *Theoretical Computer Science*, 832:143–165.
- Krejca, M. S. and Witt, C. (2020b). Theory of estimation-of-distribution algorithms. In *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 405–442. Springer. Also available at <http://arxiv.org/abs/1806.05392>.
- Lehre, P. K. (2011). Fitness-levels for non-elitist populations. In *Proc. of GECCO '11*, pages 2075–2082.
- Lehre, P. K. and Nguyen, P. T. H. (2017). Improved runtime bounds for the univariate marginal distribution algorithm via anti-concentration. In *Proc. of GECCO '17*, pages 1383–1390.
- Lehre, P. K. and Nguyen, P. T. H. (2019). On the limitations of the univariate marginal distribution algorithm to deception and where bivariate EDAs might help. In *Proc. of FOGA '19*, pages 154–168.
- Lengler, J., Sudholt, D., and Witt, C. (2018). Medium step sizes are harmful for the compact genetic algorithm. In *Proc. of GECCO '18*, pages 1499–1506.
- Mühlenbein, H. and Paaß, G. (1996). From recombination of genes to the estimation of distributions I. Binary parameters. In *Proc. of PPSN '96*, pages 178–187.

- Pelikan, M., Hauschild, M., and Lobo, F. G. (2015). Estimation of distribution algorithms. In *Springer Handbook of Computational Intelligence*, pages 899–928. Springer.
- Sudholt, D. and Witt, C. (2019). On the choice of the update strength in estimation-of-distribution algorithms and ant colony optimization. *Algorithmica*, 81(4):1450–1489.
- Witt, C. (2018). Domino convergence: why one should hill-climb on linear functions. In *Proc. of GECCO '18*, pages 1539–1546.
- Witt, C. (2019). Upper bounds on the running time of the univariate marginal distribution algorithm on OneMax. *Algorithmica*, 81(2):632–667.