



HAL
open science

An exponential lower bound for the runtime of the compact genetic algorithm on jump functions

Benjamin Doerr

► **To cite this version:**

Benjamin Doerr. An exponential lower bound for the runtime of the compact genetic algorithm on jump functions. FOGA '19: Foundations of Genetic Algorithms XV, 2019, Potsdam Germany, France. pp.25-33, 10.1145/3299904.3340304 . hal-04484767

HAL Id: hal-04484767

<https://hal.science/hal-04484767>

Submitted on 4 Apr 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.

An Exponential Lower Bound for the Runtime of the cGA on Jump Functions*

Benjamin Doerr
École Polytechnique
CNRS
Laboratoire d'Informatique (LIX)
Palaiseau
France

June 26, 2019

Abstract

In the first runtime analysis of an estimation-of-distribution algorithm (EDA) on the multi-modal jump function class, Hasenöhr and Sutton (GECCO 2018) proved that the runtime of the compact genetic algorithm with suitable parameter choice on jump functions with high probability is at most polynomial (in the dimension) if the jump size is at most logarithmic (in the dimension), and is at most exponential in the jump size if the jump size is super-logarithmic. The exponential runtime guarantee was achieved with a hypothetical population size that is also exponential in the jump size. Consequently, this setting cannot lead to a better runtime.

In this work, we show that any choice of the hypothetical population size leads to a runtime that, with high probability, is at least exponential in the jump size. This result might be the first non-trivial exponential lower bound for EDAs that holds for arbitrary parameter settings.

*Notice required by the publisher of the final version of this work: (c) by the author. This is the author's version of the submitted version of this work. It is posted here for your personal use. Not for redistribution. The definitive version will be published in "Benjamin Doerr. 2019. An Exponential Lower Bound for the Runtime of the Compact Genetic Algorithm on Jump Functions. In Foundations of Genetic Algorithms XV (FOGA 19), August 27–29, 2019, Potsdam, Germany. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3299904.3340304>."

1 Introduction

Due to the inherent highly complex stochastic processes, the mathematical analysis of estimation-of-distribution algorithms (EDAs) is still in its early childhood. Whereas for classic evolutionary algorithms many deep analysis exist, see, e.g., [NW10, AD11, Jan13], for EDAs even some of the most basic problems are not fully understood, such as the runtime of the compact genetic algorithm (cGA) on the ONEMAX benchmark function [Dro06, SW16, LSW18]. We direct the reader to the recent survey [KW18] for a complete picture of the state of the art in mathematical analyses of EDAs.

Given this state of the art, it is not surprising that the first runtime analysis of an EDA on a multi-unimodal objective function appeared only very recently. In their GECCO 2018 paper, Hasenöhrl and Sutton [HS18] analyze the optimization time of the cGA on the jump function class. Jump functions are simple unimodal functions except that they have a valley of low fitness of scalable size k around the global optimum. Hasenöhrl and Sutton show [HS18, Theorem 3.3] that, for a sufficiently large constant C and any constant $\varepsilon > 0$, the cGA with hypothetical population size at least $\mu \geq \max\{Cne^{4k}, n^{3.5+\varepsilon}\}$ with probability $1 - o(1)$ finds the optimum of any jump function with jump size at most $k = o(n)$ in $O(\mu n^{1.5} \log n + e^{4k})$ generations (which is also the number of fitness evaluations, since the cGA evaluates two search points in each iteration).

We note that this runtime guarantee is polynomial (in the problem dimension n) when $k = O(\log n)$ and exponential (in k) otherwise. Both parts of the result are remarkable when recalling that most classical evolutionary algorithms need time $\Omega(n^k)$.

For the polynomial part, the upper bound of order $\mu n^{1.5} \log n$, which is $n^{5+\varepsilon}$ when choosing μ optimally, was for $k < \frac{1}{20} \ln n$ recently [Doe19] improved to $O(\mu\sqrt{n})$ for all $\mu = \Omega(\sqrt{n} \log n) \cap \text{poly}(n)$, which is $O(n \log n)$ for the optimal choice $\mu = \Theta(\sqrt{n} \log n)$. Note that $n \log n$ is the asymptotic runtime of many evolutionary algorithms, including the cGA with good parameter choices [Dro06, SW16, LSW18], on the simple unimodal ONEMAX problem. Hence this result shows that the cGA does not suffer significantly from the valley of low fitness around the optimum which is characteristic for jump functions (as long as this valley is not too wide, that is, $k < \frac{1}{20} \ln n$). If we are willing to believe that $\Omega(n \log n)$ is also a lower bound for the runtime of the cGA on these jump functions (which given the results for ONEMAX appears very plausible, but which seems hard to prove, see Section 4), then the result in [Doe19] determines the precise asymptotic runtime of the cGA with optimal parameter choice for $k < \frac{1}{20} \ln n$.

What is left open by these two previous works is how good the exponential upper bound (for k super-logarithmic in n) is. Since Hasenöhrl and Sutton prove their exponential runtime guarantee only for a hypothetical population size $\mu = \Omega(ne^{4k})$, it is clear that they, in this setting, cannot have a sub-exponential runtime (for the sake of completeness, we shall make this elementary argument precise in Lemma 1). So the question remains if, by choosing a smaller hypothetical population size, one could have obtained a better runtime guarantee.

Our main result is a negative answer to this question. In Theorem 6 we show that, regardless of the hypothetical population size, the runtime of the cGA on a jump function with jump size k is at least exponential in k with high probability. Interestingly, not only our result is a uniform lower bound independent of the hypothetical population size, but our proof is also “uniform” in the sense that it does neither need case distinctions w.r.t. the hypothetical population size nor w.r.t. different reasons for the lower bound. Here we recall that the existing runtime analyses, see, e.g., again [Dro06, SW16, LSW18] find two reasons why an EDA can be inefficient. (i) The hypothetical population size is large and consequently it takes long to move the frequencies into the direction of the optimum. (ii) The hypothetical population size is small and thus, in the absence of a strong fitness signal, the random walk of the frequencies brings some frequencies close to the boundaries of the frequency spectrum; from there they are hard to move back into the game.

We avoid such potentially tedious case distinctions via an elegant drift argument on the sum of the frequencies. Ignoring some technicalities here, we show that, regardless of the hypothetical population size, the frequency sum overshoots a value of $n - \frac{1}{4}k$ only after an expected number of $2^{\Omega(k)}$ iterations. However, in an iterations where the frequency sum is below $n - \frac{1}{4}k$, the optimum is sampled only with probability $2^{-\Omega(k)}$. These two results give our main result.

As a **side result**, we show in Section 4 that a result like “ONEMAX is an easiest function with a unique global optimum for the cGA”, if true at all, cannot be proven along the same lines as the corresponding results for many mutation-based algorithms. This in particular explains why we and the previous works on this topic have not shown an $\Omega(n \log n)$ lower bound for the runtime of the cGA on jump functions.

2 Preliminaries

2.1 The Compact Genetic Algorithm

The *compact genetic algorithm* (cGA) is an estimation-of-distribution algorithm (EDA) proposed by Harik, Lobo, and Goldberg [HLG99] for the maximization of pseudo-Boolean functions $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{R}$. Being a univariate EDA, it develops a probabilistic model described by a frequency vector $f \in [0, 1]^n$. This frequency vector describes a probability distribution on the search space $\{0, 1\}^n$. If $X = (X_1, \dots, X_n) \in \{0, 1\}^n$ is a search point sampled according to this distribution—we write

$$X \sim \text{Sample}(f)$$

to indicate this—then we have $\Pr[X_i = 1] = f_i$ independently for all $i \in [1..n] := \{1, \dots, n\}$. In other words, the probability that X equals some fixed search point y is

$$\Pr[X = y] = \prod_{i:y_i=1} f_i \prod_{i:y_i=0} (1 - f_i).$$

In each iteration, the cGA updates this probabilistic model as follows. It samples two search points $x^1, x^2 \sim \text{Sample}(f)$, computes the fitness of both, and defines $(y^1, y^2) = (x^1, x^2)$ when x^1 is at least as fit as x^2 and $(y^1, y^2) = (x^2, x^1)$ otherwise. Consequently, y^1 is the rather better search point of the two. We then define a preliminary model by $f' := f + \frac{1}{\mu}(y^1 - y^2)$. This definition ensures that, when y^1 and y^2 differ in some bit position i , the i -th preliminary frequency moves by a step of $\frac{1}{\mu}$ into the direction of y_i^1 , which we hope to be the right direction since y^1 is the better of the two search points. The *hypothetical populations size* μ is used to control how strong this update is.

To avoid a premature convergence, we ensure that the new frequency vector is in $[\frac{1}{n}, 1 - \frac{1}{n}]^n$ by capping too small or too large values at the corresponding boundaries. More precisely, for all $\ell \leq u$ and all $r \in \mathbb{R}$ we define

$$\text{minmax}(\ell, r, u) := \max\{\ell, \min\{r, u\}\} = \begin{cases} \ell & \text{if } r < \ell \\ r & \text{if } r \in [\ell, u] \\ u & \text{if } r > u \end{cases}$$

and we lift this notation to vectors by reading it component-wise. Now the new frequency vector is $\text{minmax}(\frac{1}{n}\mathbf{1}_n, f', (1 - \frac{1}{n})\mathbf{1}_n)$.

This iterative frequency development is pursued until some termination criterion is met. Since we aim at analyzing the time (number of iterations)

it takes to sample the optimal solution (this is what we call the *runtime* of the cGA), we do not specify a termination criterion and pretend that the algorithm runs forever.

The pseudo-code for the cGA is given in Algorithm 1. We shall use the notation given there frequently in our proofs. For the frequency vector f_t obtained at the end of iteration t , we denote its i -th component by $f_{i,t}$ or, when there is no risk of ambiguity, by f_{it} .

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

```

1  $t \leftarrow 0$ ;
2  $f_t = (\frac{1}{2}, \dots, \frac{1}{2}) \in [0, 1]^n$ ;
3 repeat
4    $x^1 \leftarrow \text{Sample}(f_t)$ ;
5    $x^2 \leftarrow \text{Sample}(f_t)$ ;
6   if  $\mathcal{F}(x^1) \geq \mathcal{F}(x^2)$  then  $(y^1, y^2) \leftarrow (x^1, x^2)$  else  $(y^1, y^2) \leftarrow (x^2, x^1)$ ;
7    $f'_{t+1} \leftarrow f_t + \frac{1}{\mu}(y^1 - y^2)$ ;
8    $f_{t+1} \leftarrow \text{minmax}(\frac{1}{n}\mathbf{1}_n, f'_{t+1}, (1 - \frac{1}{n})\mathbf{1}_n)$ ;
9    $t \leftarrow t + 1$ ;
10 until forever;
```

Well-behaved frequency assumption: For the hypothetical population size μ , we take the common assumption that any two frequencies that can occur in a run of the cGA differ by a multiple of $\frac{1}{\mu}$. We call this the *well-behaved frequency assumption*. This assumption was implicitly already made in [HLG99] by using even μ in all experiments (note that the hypothetical population size is denoted by n in [HLG99]). This assumption was made explicit in [Dro06] by requiring μ to be even. Both works do not use the frequencies boundaries $\frac{1}{n}$ and $1 - \frac{1}{n}$, so an even value for μ ensures well-behaved frequencies.

For the case with frequency boundaries, the well-behaved frequency assumption is equivalent to $(1 - \frac{2}{n})$ being an even multiple of the update step size $\frac{1}{\mu}$. In this case, $n_\mu = (1 - \frac{2}{n})\mu \in 2\mathbb{N}$ and the set of frequencies that can occur is

$$F := F_\mu := \{\frac{1}{n} + \frac{i}{\mu} \mid i \in [0..n_\mu]\}.$$

This assumption was made, e.g., in the proof of Theorem 2 in [SW16] and in the paper [LSW18] (see the paragraph following Lemma 2.1).

A trivial lower bound: We finish this subsection on the cGA with the following very elementary remark, which shows that the cGA with hypothetical population size μ with probability $1 - \exp(-\Omega(n))$ has a runtime

of at least $\min\{\mu/4, \exp(\Theta(n))\}$ on any $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{R}$ with a unique global optimum. This shows, in particular, that the cGA with the parameter value $\mu = \exp(\Omega(k))$ used to optimize jump functions with gap size $k \in \omega(\log n) \cap o(n)$ in time $\exp(O(k))$ in [HS18] cannot have a runtime better than exponential in k .

Lemma 1. *Let $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{R}$ have a unique global optimum. The probability that the cGA generates the optimum of f in $T = \min\{\mu/4, (1.3)^n\}$ iterations is at most $\exp(-\Omega(n))$.*

Proof. By the definition of the cGA, the frequency vector f used in iteration $t = 1, 2, 3, \dots$ satisfies $f \in [\frac{1}{2} - \frac{t-1}{\mu}, \frac{1}{2} + \frac{t-1}{\mu}]^n$. Consequently, the probability that a fixed one of the two search points which are generated in this iteration is the optimum, is at most $(\frac{1}{2} + \frac{t-1}{\mu})^n$. For $t \leq \mu/4$, this is at most $(3/4)^n$. Hence by a simple union bound, the probability that the optimum is generated in the first $T = \min\{\mu/4, (1.3)^n\}$ iterations, is at most $2T(3/4)^n = \exp(-\Omega(n))$. \square

2.2 Related Work

In all results described in this section, we shall assume that the hypothetical population size is at most polynomial in the problem size n , that is, that there is a constant c such that $\mu \leq n^c$.

The first to conduct a rigorous runtime analysis for the cGA was Droste in his seminal work [Dro06]. He regarded the cGA without frequency boundaries, that is, he just took $f_{t+1} := f'_{t+1}$ in our notation. He showed that this algorithm with $\mu \geq n^{1/2+\varepsilon}$, $\varepsilon > 0$ any positive constant, finds the optimum of the ONEMAX function defined by

$$\text{ONEMAX}(x) = \|x\|_1 = \sum_{i=1}^n x_i$$

for all $x \in \{0, 1\}^n$ with probability at least $1/2$ in $O(\mu\sqrt{n})$ iterations [Dro06, Theorem 8].

Droste also showed that this cGA for any objective function \mathcal{F} with unique optimum has an expected runtime of $\Omega(\mu\sqrt{n})$ when conditioning on no premature convergence [Dro06, Theorem 6]. It is easy to see that his proof of the lower bound can be extended to the cGA with frequency boundaries, that is, to Algorithm 1. For this, it suffices to deduce from his drift argument the result that the first time $T_{n/4}$ that the frequency distance $D = \sum_{i=1}^n (1 - f_{it})$ is less than $n/4$ satisfies $E[T_{n/4}] \geq \mu\sqrt{n}\frac{\sqrt{2}}{4}$. Since the

probability to sample the optimum from a frequency distance of at least $n/4$ is at most

$$\begin{aligned} \prod_{i=1}^n f_{it} &= \prod_{i=1}^n (1 - (1 - f_{it})) \leq \prod_{i=1}^n \exp(-(1 - f_{it})) \\ &= \exp\left(-\sum_{i=1}^n (1 - f_{it})\right) \leq \exp(-n/4), \end{aligned}$$

the algorithm with high probability does not find the optimum before time $T_{n/4}$.

Ten years after Droste’s work, Sudholt and Witt [SW16] showed that the $O(\mu\sqrt{n})$ upper bound also holds for the cGA with frequency boundaries. There (but the same should be true for the cGA without boundaries) a hypothetical population size of $\mu = \Omega(\sqrt{n} \log n)$ suffices (recall that Droste required $\mu = \Omega(n^{1/2+\epsilon})$). The technically biggest progress with respect to upper bounds most likely lies in the fact that the analysis in [SW16] also holds for the expected optimization time, which means that it also includes the rare case that frequencies reach the lower boundary. We refer to Section 2.3 of [Doe19] for a detailed discussion of the relation of expectations and tail bounds for runtimes of EDAs, including a method to transform EDAs with with-high-probability guarantees into those with guarantees on the expected runtime). Sudholt and Witt also show that the cGA with frequency boundaries with high probability (and thus also in expectation) needs at least $\Omega(\mu\sqrt{n} + n \log n)$ iterations to optimize ONEMAX. While the $\mu\sqrt{n}$ lower bound could have been also obtained with methods similar to Droste’s, the innocent-looking $\Omega(n \log n)$ bound is surprisingly difficult to prove.

Not much is known for hypothetical population sizes below the order of \sqrt{n} . It is clear that then the frequencies will reach the lower boundary of the frequency range, so working with a non-trivial lower boundary like $\frac{1}{n}$ is necessary to prevent premature convergence. The recent lower bound $\Omega(\mu^{1/3}n)$ valid for $\mu = O(\frac{\sqrt{n}}{\log n \log \log n})$ of [LSW18] indicates that already a little below the \sqrt{n} regime significantly larger runtimes occur, but with no upper bounds this regime remains largely not understood.

We refer the reader to the recent survey [KW18] for more results on the runtime of the cGA on classic unimodal test functions like LEADINGONES and BINVAL. Interestingly, nothing was known for non-unimodal functions before the recent work of Hasenöhl and Sutton [HS18] on jump functions, which we discussed already in the introduction.

The general topic of lower bounds on runtimes of EDAs remains largely little understood. Apart from the lower bounds for the cGA on ONEMAX

discussed above, the following is known. Krejca and Witt [KW17] prove a lower bound for the UMDA on ONEMAX, which is of a similar flavor as the lower bound for the cGA of Sudholt and Witt [SW16]: For $\lambda = (1 + \beta)\mu$, where $\beta > 0$ is a constant, and λ polynomially bounded in n , the expected runtime of the UMDA on ONEMAX is $\Omega(\mu\sqrt{n} + n \log n)$. For the binary value function BINVAL, Droste [Dro06] and Witt [Wit18] together give a lower bound of $\Omega(\min\{n^2, Kn\})$ for the runtime of the cGA. Apart from these sparse results, we are not aware of any lower bounds for EDAs. Of course, the black-box complexity of the problem is a lower bound for any black-box algorithm, hence also for EDAs, but these bounds are often lower than the true complexity of a given algorithm. For example, the black-box complexities of ONEMAX, LEADINGONES, and jump functions with jump size $k \leq \frac{1}{2}n - n^\varepsilon$, $\varepsilon > 0$ any constant, are $\Theta(\frac{n}{\log n})$ [DJW06, AW09], $\Theta(n \log \log n)$ [AAD+13], and $\Theta(\frac{n}{\log n})$ [BDK16], respectively.

To round off the picture, we briefly describe some typical runtimes of evolutionary algorithms on jump functions. We recall that the n -dimensional jump function with jump size $k \geq 1$ is defined by

$$\text{JUMP}_{nk}(x) = \begin{cases} \|x\|_1 + k & \text{if } \|x\|_1 \in [0..n - k] \cup \{n\}, \\ n - \|x\|_1 & \text{if } \|x\|_1 \in [n - k + 1..n - 1]. \end{cases}$$

Hence for $k = 1$, we have a fitness landscape isomorphic to the one of ONEMAX, but for larger values of k there is a fitness valley (“gap”)

$$G_{nk} := \{x \in \{0, 1\}^n \mid n - k < \|x\|_1 < n\}$$

consisting of the $k - 1$ highest sub-optimal fitness levels of the ONEMAX function. This valley is hard to cross for evolutionary algorithms using standard-bit mutation with mutation rate $\frac{1}{n}$ since with very high probability they need to generate the optimum from one of the local optima, which in a single application of the mutation operator happens only with probability less than n^{-k} . For this reason, e.g., the classic $(\mu + \lambda)$ and (μ, λ) EAs all have a runtime of at least n^k . This was proven formally for the $(1 + 1)$ EA in the classic paper [DJW02], but the argument just given proves the n^k lower bound equally well for all $(\mu + \lambda)$ and (μ, λ) EAs. By using larger mutation rates or a heavy-tailed mutation operator, a $k^{\Theta(k)}$ runtime improvement can be obtained [DLMN17], but the runtime remains $\Omega(n^k)$ for k constant.

Asymptotically better runtimes can be achieved when using crossover, though this is harder than expected. The first work in this direction [JW02], among other results, could show that a simple $(\mu + 1)$ genetic algorithm using uniform crossover with rate $p_c = O(1/kn)$ obtains an $O(\mu n^2 k^3 + 2^{2k} p_c^{-1})$

runtime when the population size is at least $\mu = \Omega(k \log n)$. A shortcoming of this result, already noted by the authors, is that it only applies to uncommonly small crossover rates. Using a different algorithm that first always applies crossover and then mutation, a runtime of $O(n^{k-1} \log n)$ was achieved by Dang et al. [DFK⁺18, Theorem 2]. For $k \geq 3$, the logarithmic factor in the runtime can be removed by using a higher mutation rate. With additional diversity mechanisms, the runtime can be further reduced up to $O(n \log n + 4^k)$, see [DFK⁺16]. In the light of this last result, the insight stemming from the previous work [HS18] and ours is that the cGA apparently without further modifications supplies the necessary diversity to obtain a runtime of $O(n \log n + 2^{O(k)})$.

Finally, we note that runtimes of $O(n \binom{n}{k})$ and $O(k \log(n) \binom{n}{k})$ were shown for the $(1 + 1)$ IA^{hyp} and the $(1 + 1)$ Fast-IA artificial immune systems, respectively [COY17, COY18].

2.3 Preliminaries

We now collect a few elementary tools that will be used on our analysis. The first is well-known and the next two are from [Doe19], so it is only the last one which we could not find in the literature.

The following estimate seems well-known (e.g., it was used in [JJW05] without proof or reference). Gießen and Witt [GW17, Lemma 3] give a proof via estimates of binomial coefficients and the binomial identity. A more elementary proof can be found in [Doe18b, Lemma 10.37].

Lemma 2. *Let $X \sim \text{Bin}(n, p)$. Let $k \in [0..n]$. Then*

$$\Pr[X \geq k] \leq \binom{n}{k} p^k.$$

The next estimate was essentially proven in the extended version of [Doe19, Lemma 1]. The only small difference here is that in the first inequality, we used a slightly different Chernoff bound, which also allows deviation parameters Δ which are greater than one. We therefore omit the proof.

Lemma 3. *Let $f \in [0, 1]^n$, $D := n - \|f\|_1$, $D^- \leq D \leq D^+$, $x \sim \text{Sample}(f)$, and $d(x) := n - \|x\|_1$. Then for all $\Delta \geq 0$ and $\delta \in [0, 1]$, we have*

$$\begin{aligned} \Pr[d(x) \geq (1 + \Delta)D^+] &\leq \exp(-\tfrac{1}{3} \min\{\Delta^2, \Delta\} D^+), \\ \Pr[d(x) \leq (1 - \delta)D^-] &\leq \exp(-\tfrac{1}{2} \delta^2 D^-). \end{aligned}$$

To estimate the influence from capping the frequencies into the interval $[\frac{1}{n}, 1 - \frac{1}{n}]$, the following elementary result was shown in the extended version of [Doe19].

Lemma 4. *Let $P = 2\frac{1}{n}(1 - \frac{1}{n})$. Let $t \geq 0$. Using the notation given in Algorithm 1, consider iteration $t+1$ of a run of the cGA started with a fixed frequency vector $f_t \in [\frac{1}{n}, 1 - \frac{1}{n}]^n$.*

(i) *Let $L = \{i \in [1..n] \mid f_{it} = \frac{1}{n}\}$, $\ell = |L|$, and $M = \{i \in L \mid x_i^1 \neq x_i^2\}$. Then $|M| \sim \text{Bin}(\ell, P)$ and*

$$\|f_{t+1}\|_1 - \|f'_{t+1}\|_1 \preceq \|(f_{t+1})_{|L}\|_1 - \|(f'_{t+1})_{|L}\|_1 \preceq \frac{1}{\mu}|M| \preceq \frac{1}{\mu} \text{Bin}(n, \frac{2}{n}).$$

(ii) *Let $L = \{i \in [1..n] \mid f_{it} = 1 - \frac{1}{n}\}$, $\ell = |L|$, and $M = \{i \in L \mid x_i^1 \neq x_i^2\}$. Then $|M| \sim \text{Bin}(\ell, P)$ and*

$$\|f'_{t+1}\|_1 - \|f_{t+1}\|_1 \preceq \|(f'_{t+1})_{|L}\|_1 - \|(f_{t+1})_{|L}\|_1 \preceq \frac{1}{\mu}|M| \preceq \frac{1}{\mu} \text{Bin}(n, \frac{2}{n}).$$

To argue that the cGA makes at least some small progress, we shall use the following blunt estimate for the probability that two bit strings $x, y \sim \text{Sample}(f)$ sampled from the same product distribution have a different distance from the all-ones string (and, by symmetry, from any other string, but this is a statement which we do not need here).

Lemma 5. *Let $n \in \mathbb{N}$, $m \in [\frac{n}{2}, n]$, and $f \in [\frac{1}{n}, 1 - \frac{1}{n}]^m$. Let $x^1, x^2 \sim \text{Sample}(f)$ be independent. Then $\Pr[\|x^1\|_1 \neq \|x^2\|_1] \geq \frac{1}{16}$.*

Proof. For all $v \in \mathbb{R}^m$ and $a, b \in [1..m]$ with $a \leq b$ we use the abbreviation $v_{[a..b]} := \sum_{i=a}^b v_i$. By symmetry, we can assume that $f_{[1..m]} \leq \frac{m}{2}$. Without loss of generality, we may further assume that $f_i \leq f_{i+1}$ for all $i \in [1..m-1]$. We have $f_{\lfloor m/4 \rfloor} \leq \frac{2}{3}$ as otherwise

$$f_{[1..m]} \geq f_{[\lfloor m/4 \rfloor + 1..n]} > \frac{2}{3}(n - \lfloor m/4 \rfloor) \geq \frac{2}{3} \cdot \frac{3}{4}m = \frac{m}{2},$$

contradicting our assumption.

Let ℓ be minimal such that $S = f_{[1..\ell]} \geq \frac{1}{8}$. Since $\ell \leq \frac{n}{8} \leq \frac{m}{4}$, we have $f_\ell \leq \frac{2}{3}$ and thus $S \leq \frac{1}{8} + \frac{2}{3} = \frac{19}{24}$.

For $j \in \{0, 1\}$ let $q_j = \Pr[x_{[1..\ell]}^1 = j] = \Pr[x_{[1..\ell]}^2 = j]$. We compute

$$\begin{aligned} \Pr[\|x^1\|_1 \neq \|x^2\|_1] &\geq \Pr[x_{[1..\ell]}^1 = x_{[1..\ell]}^2 \wedge x_{[\ell+1..n]}^1 \neq x_{[\ell+1..n]}^2] \\ &\quad + \Pr[x_{[1..\ell]}^1 \neq x_{[1..\ell]}^2 \wedge x_{[\ell+1..n]}^1 = x_{[\ell+1..n]}^2] \\ &= \Pr[x_{[1..\ell]}^1 = x_{[1..\ell]}^2] \Pr[x_{[\ell+1..n]}^1 \neq x_{[\ell+1..n]}^2] \end{aligned}$$

$$\begin{aligned}
& + \Pr[x_{[1..\ell]}^1 \neq x_{[1..\ell]}^2] \Pr[x_{[\ell+1..n]}^1 = x_{[\ell+1..n]}^2] \\
& \geq \min\{\Pr[x_{[1..\ell]}^1 = x_{[1..\ell]}^2], \Pr[x_{[1..\ell]}^1 \neq x_{[1..\ell]}^2]\} \\
& \geq \min\{q_0^2 + q_1^2, 2q_0q_1\} = 2q_0q_1,
\end{aligned}$$

the latter by the inequality of the arithmetic and geometric mean. Using Bernoulli's inequality, we estimate coarsely

$$\begin{aligned}
q_0 &= \prod_{i=1}^{\ell} (1 - f_i) \geq 1 - f_{[1..\ell]}, \\
q_1 &= \sum_{i=1}^{\ell} f_i \prod_{j \in [1..\ell] \setminus \{i\}} (1 - f_j) \geq f_{[1..\ell]} (1 - f_{[1..\ell]}).
\end{aligned}$$

From the concavity of $z \mapsto z(1-z)^2$ in $[0, 1]$, we obtain

$$\begin{aligned}
2q_0q_1 &\geq 2 \min\{z(1-z)^2 \mid z \in [\frac{1}{8}, \frac{19}{24}]\} \\
&= 2 \min\{z(1-z)^2 \mid z \in \{\frac{1}{8}, \frac{19}{24}\}\} = 2 \frac{19}{24} (\frac{5}{25})^2 \geq \frac{1}{16}.
\end{aligned}$$

□

3 Main Result

We now state precisely our main result, explain the central proof ideas, and state the formal proof.

Theorem 6. *There are constants $\alpha_1, \alpha_2 > 0$ such that for any n sufficiently large and any $k \in [1..n]$, regardless of the hypothetical population size μ , the runtime of the cGA on JUMP_{nk} with probability $1 - \exp(-\alpha_1 k)$ is at least $\exp(\alpha_2 k)$. In particular, the expected runtime is exponential in k .*

To prove this result, we will regard the stochastic process $D_t := n - \|f_t\|_1$, that is, the difference of the sum of the frequencies from the ideal value n . Our general argument is that this process with probability $1 - \exp(-\Omega(k))$ stays above $\frac{1}{4}k$ for $\exp(\Omega(k))$ iterations. In each iteration where $D_t \geq \frac{1}{4}k$, the probability that the optimum is sampled, is only $\exp(-\Omega(k))$. Hence there is a $T = \exp(\Omega(k))$ such that with probability $1 - \exp(-\Omega(k))$, the optimum is not sampled in the first T iterations.

The heart of the proof is an analysis of the process (D_t) . It is intuitively clear that once the process is below k , then often the two search points sampled in one iteration both lie in the gap region, which gives a positive drift (that is, a decrease of the average frequency). To turn this drift away from

the target (a small D_t value) into an exponential lower bound for the runtime, we consider the process $Y_t = \exp(c \min\{\frac{1}{2}k - D_t, \frac{1}{4}k\})$, that is, an exponential rescaling of D_t . Such a rescaling has recently also been used in [ADY19]. We note that the usual way to prove exponential lower bounds is the negative drift theorem of Oliveto and Witt [OW12]. We did not immediately see how to use it for our purposes, though, since in our process we do not have very strong bounds on the one-step differences. E.g., when $D_t = \frac{1}{2}k$, then the underlying frequency vector may be such that $D_{t+1} \geq D_t + \sqrt{k}$ happens with constant probability.

The main technical work now is showing that the process Y_t has at most a constant drift, more precisely, that $E[Y_{t+1} - Y_t] \leq 2$ whenever $Y_t < Y_{\max}$. The difficulty is hidden in a small detail. When $D_t \in [\frac{1}{4}k, \frac{3}{4}k]$, and this is the most interesting case, then we have $\|f'_{t+1}\|_1 \geq \|f_t\|_1$ whenever the two search points sampled lie in the gap region, and hence with probability $1 - \exp(-\Omega(k))$; from Lemma 5 we obtain, in addition, a true increase, that is, $\|f'_{t+1}\|_1 \geq \|f_t\|_1 + \frac{1}{\mu}$, with constant probability. Hence the true difficulty arises from the capping of the frequencies into the interval $[\frac{1}{n}, 1 - \frac{1}{n}]$. This appears to be a minor problem, among others, because only a capping at the lower bound $\frac{1}{n}$ can have an adverse effect on our process, and there are at most $O(k)$ frequencies sufficiently close to the lower boundary. Things become difficult due to the exponential scaling, which can let rare event still have a significant influence on the expected change of the process.

We now make these arguments precise and prove Theorem 6.

Proof. Since we are aiming at an asymptotic statement, we can assume in the following that n is sufficiently large.

Since it will ease the presentation when we can assume that $k \geq w(n)$ for some function $w : \mathbb{N} \rightarrow \mathbb{N}$ with $\lim_{n \rightarrow \infty} w(n) = \infty$, let us first give a basic argument for the case of small k .

We first note that with probability $f_{it}^2 + (1 - f_{it})^2 \geq \frac{1}{2}$, the two search points x^1 and x^2 generated in the t -th iteration agree on the i -th bit, which in particular implies that $f_{i,t+1} = f_{it}$. Hence with probability at least 2^{-T} , this happens for the first T iterations, and thus $f_{it} = \frac{1}{2}$ for all $t \in [0..T]$. Let us call such a bit position i *sleepy*.

Note that the events of being sleepy are independent for all $i \in [1..n]$. Hence, taking $T = \lfloor \frac{1}{2} \log_2 n \rfloor$, we see that the number X of sleepy positions has an expectation of $E[X] \geq n2^{-T} \geq \sqrt{n}$, and by a simple Chernoff bound, we have $\Pr[X \geq \frac{1}{2}\sqrt{n}] \geq 1 - \exp(-\Omega(\sqrt{n}))$.

Conditional on having at least $\frac{1}{2}\sqrt{n}$ sleepy bit positions, the probability that a particular search point sampled in the first T iterations is the optimum is at most $2^{-\frac{1}{2}\sqrt{n}}$. By a simple union bound argument, the probability

that at least one of the search points generated in the first T iterations is the optimum is at most $2T2^{-\frac{1}{2}\sqrt{n}} = \exp(-\Omega(\sqrt{n}))$. In summary, we have that with probability at least $1 - \exp(-\Omega(\sqrt{n}))$, the runtime of the cGA on any function with unique optimum (and in particular any jump function) is greater than $T = \frac{1}{2} \log_2 n$. This implies the claim of this theorem for any $k \leq C \log \log n$, where C is a sufficiently small constant, and, as discussed above, n is sufficiently large.

With this, we can now safely assume that $k = \omega(1)$. For the case that $k \geq \frac{n}{320}$, we will need slightly modified calculations. To keep this proof readable, we hide (but treat nevertheless) this case as follows. We consider a run of the cGA on a jump function $\text{JUMP}_{nk'}$ with $k' \in [1..n] \cap \omega(1)$ arbitrary and we let $k := \min\{k', \frac{n}{320}\}$.

Let $X_t := \|f_t\|_1 = \sum_{i=1}^n f_{it}$ be the sum of the frequencies at the end of iteration t . Since we are mostly interested in the case where X_t is close to the maximum value, we also define $D_t = n - X_t$.

Our intuition (which will be made precise) is that the process (D_t) finds it hard to go significantly below k because there we will typically sample individuals in the gap, which lead to a decrease of the sum of frequencies (when the two individuals have different distances from the optimum). To obtain an exponential lower bound on the runtime, we suitably rescale the process by defining, for a sufficiently small constant c ,

$$Y_t = \min\{\exp(c(\frac{1}{2}k - D_t)), \exp(\frac{1}{4}ck)\} = \exp(c \min\{\frac{1}{2}k - D_t, \frac{1}{4}k\}).$$

Observe that Y_t attains its maximal value $Y_{\max} = \exp(\frac{1}{4}ck)$ precisely when $D_t \leq \frac{1}{4}k$. Also, $Y_t \leq 1$ for $D_t \geq \frac{1}{2}k$.

To argue that we have $D_t > \frac{1}{4}k$ for a long time, we now show that the drift $E[Y_{t+1} - Y_t \mid Y_t < \exp(\frac{1}{4}ck)]$ is at most constant. To this aim, we condition on a fixed value of f_t , which also determines X_t and D_t . We treat separately the two cases that $D_t \geq \frac{3}{4}k$ and that $\frac{3}{4}k > D_t > \frac{1}{4}k$.

Case 1: Assume first that $D_t \geq \frac{3}{4}k$. By Lemma 3, with probability $1 - \exp(-\Omega(k))$, the two search points x^1, x^2 sampled in iteration $t + 1$ both satisfy $|\|x^i\|_1 - X_t| < \frac{1}{6}(D_t - \frac{1}{2}k)$. Let us call this event A . In this case, we argue as follows. Let $\{y^1, y^2\} = \{x^1, x^2\}$ such that $f(y^1) \geq f(y^2)$. Then

$$\begin{aligned} \|f'_{t+1}\|_1 &= \|f_t + \frac{1}{\mu}(y^1 - y^2)\|_1 \leq \|f_t\|_1 + \frac{1}{\mu}2\frac{1}{6}(D_t - \frac{1}{2}k) \\ &\leq \|f_t\|_1 + 2\frac{1}{6}(D_t - \frac{1}{2}k) = n - D_t + 2\frac{1}{6}(D_t - \frac{1}{2}k) \\ &= n - \frac{2}{3}D_t - \frac{1}{6}k \leq n - \frac{2}{3} \cdot \frac{3}{4}k - \frac{1}{6}k \leq n - \frac{2}{3}k. \end{aligned}$$

We still need to consider the possibility that $f_{i,t+1} > f'_{i,t+1}$ for some $i \in [1..n]$. By Lemma 4, not conditioning on A , we have that $\|f_{t+1}\|_1 - \|f'_{t+1}\|_1 \preceq \frac{1}{\mu} \text{Bin}(\ell, P) \preceq \text{Bin}(\ell, P)$ for some $\ell \in [1..n]$ and $P = 2\frac{1}{n}(1 - \frac{1}{n})$.

Let us call B the event that $\|f_{t+1}\|_1 - \|f'_{t+1}\|_1 < \frac{1}{6}k$. Note that $A \cap B$ implies $\|f_{t+1}\|_1 < n - \frac{1}{2}k$ and thus $Y_{t+1} \leq 1$. By Lemma 2 and the estimate $\binom{a}{b} \leq (\frac{ea}{b})^b$, we have $\Pr[B \geq \frac{1}{6}k] \leq \binom{\ell}{\frac{1}{6}k} P^{k/6} \leq (\frac{12e\ell}{kn})^{k/6} \leq k^{-\Omega(k)}$.

We conclude that the event $A \cap B$ holds with probability $1 - \exp(-\Omega(k))$; in this case $Y_t, Y_{t+1} \leq 1$. In all other cases, we bluntly estimate $Y_{t+1} - Y_t \leq Y_{\max}$. This gives $E[Y_{t+1} - Y_t] \leq (1 - \exp(-\Omega(k))) \cdot 1 + \exp(-\Omega(k))Y_{\max}$. By choosing the constant c sufficiently small and taking n sufficiently large, we have $E[Y_{t+1} - Y_t] \leq 2$.

Case 2: Assume now that $\frac{3}{4}k > D_t > \frac{1}{4}k$. Let x^1, x^2 be the two search points sampled in iteration $t+1$. By Lemma 3 again, we have $k > n - \|x^i\|_1 > 0$ with probability $1 - \exp(-\Omega(k))$ for both $i \in \{1, 2\}$. Let us call this event A . Note that if A holds, then both offspring lie in the gap region. Consequently, $\|y^1\|_1 \leq \|y^2\|_1$ and thus $\|f'_{t+1}\|_1 \leq \|f_t\|_1$.

Let $L = \{i \in [1..n] \mid f_{it} = \frac{1}{n}\}$, $\ell = |L|$, and $M = \{i \in L \mid x_i^1 \neq x_i^2\}$ as in Lemma 4. Note that by definition, $D_t \geq (1 - \frac{1}{n})\ell$, hence from $D_t < \frac{3}{4}k$ and $n \geq 4$ we obtain $\ell < k$.

Let B_0 be the event that $|M| = 0$, that is, $x_{|L}^1 = x_{|L}^2$. Note that in this case, $\|f_{t+1}\|_1 = \|f'_{t+1}\|_1 \leq \|f_t + \frac{1}{\mu}(y^1 - y^2)\|_1$. By Lemma 4, Bernoulli's inequality, and $\ell \leq k$, we have

$$\Pr[B_0] = (1 - 2\frac{1}{n}(1 - \frac{1}{n}))^\ell \geq 1 - \frac{2\ell}{n} \geq 1 - \frac{2k}{n}.$$

Since $\ell < k \leq \frac{n}{320} < \frac{n}{2}$, by Lemma 5, we have $x_{[n] \setminus L}^1 \neq x_{[n] \setminus L}^2$ with probability at least $\frac{1}{16}$. This event, called C in the following, is independent of B_0 . We have

$$\Pr[A \cap B_0 \cap C] \geq \Pr[B_0 \cap C] - \Pr[\overline{A}] \geq (1 - \frac{2k}{n})\frac{1}{16} - \exp(-\Omega(k)).$$

If $A \cap B_0 \cap C$ holds, then $\|f_{t+1}\|_1 \leq \|f_t\|_1 - \frac{1}{\mu}$. If $A \cap B_0 \cap \overline{C}$ holds, then we still have $\|f_{t+1}\|_1 \leq \|f_t\|_1$.

Let us now, for $j \in [1..\ell]$, denote by B_j the event that $|M| = j$, that is, that $x_{|L}^1$ and $x_{|L}^2$ differ in exactly j bits. By Lemma 4 again, we have $\Pr[B_j] = \Pr[\text{Bin}(\ell, P) = j]$.

The event $A \cap B_j$ implies $\|f_{t+1}\|_1 \leq \|f'_t\|_1 + \frac{j}{\mu} \leq \|f_t\|_1 + \frac{j}{\mu}$ and occurs with probability $\Pr[A \cap B_j] \leq \Pr[B_j] = \Pr[\text{Bin}(\ell, P) = j]$.

Taking these observations together, we compute

$$\begin{aligned} E[Y_{t+1}] &= \Pr[\overline{A}] E[Y_{t+1} \mid \overline{A}] \\ &\quad + \sum_{j=1}^{\ell} \Pr[A \cap B_j] E[Y_{t+1} \mid A \cap B_j] \end{aligned}$$

$$\begin{aligned}
& + \Pr[A \cap B_0 \cap \overline{C}] E[Y_{t+1} \mid A \cap B_0 \cap \overline{C}] \\
& + \Pr[A \cap B_0 \cap C] E[Y_{t+1} \mid A \cap B_0 \cap C] \\
\leq & \exp(-\Omega(k)) Y_{\max} \\
& + \sum_{j=1}^{\ell} \Pr[\text{Bin}(\ell, P) = j] Y_t \exp\left(\frac{cj}{\mu}\right) \\
& + \Pr[\text{Bin}(\ell, P) = 0] Y_t \\
& - \left(\frac{1}{16} \left(1 - \frac{2k}{n}\right) - \exp(-\Omega(k))\right) Y_t \exp\left(-\frac{c}{\mu}\right).
\end{aligned} \tag{1}$$

We note that the second and third term amount to $Y_t E[\exp(\frac{cZ}{\mu})]$, where $Z \sim \text{Bin}(\ell, P)$. Writing $Z = \sum_{i=1}^{\ell} Z_i$ as a sum of ℓ independent binary random variables with $\Pr[Z_i = 1] = P$, we obtain

$$E[\exp(\frac{cZ}{\mu})] = \prod_{i=1}^{\ell} E[\exp(\frac{cZ_i}{\mu})] = (1 - P + P \exp(\frac{c}{\mu}))^{\ell}.$$

By assuming $c \leq 1$ and using the elementary estimate $e^x \leq 1 + 2x$ valid for $x \in [0, 1]$, see, e.g., Lemma 4.2(b) in [Doe18b], we have $1 - P + P \exp(\frac{c}{\mu}) \leq 1 + 2P(\frac{c}{\mu})$. Hence with $P \leq \frac{2}{n}$, $c \leq 1$, $\mu \geq 1$, and $\ell \leq \frac{n}{320}$, we obtain

$$E[\exp(\frac{cZ}{\mu})] \leq (1 + 2P(\frac{c}{\mu}))^{\ell} \leq \exp(2P(\frac{c}{\mu})\ell) \leq \exp(\frac{4\ell}{n}) \leq 1 + \frac{8\ell}{n}$$

again by using $e^x \leq 1 + 2x$.

In the first term of (1), we again assume that c is sufficiently small to ensure that $\exp(-\Omega(k)) Y_{\max} = \exp(-\Omega(k)) \exp(\frac{1}{4}ck) \leq 1$. Recalling that $k \leq \frac{n}{320}$, we finally estimate in the last term

$$\frac{1}{16} \left(1 - \frac{2k}{n}\right) - \exp(-\Omega(k)) \exp\left(-\frac{c}{\mu}\right) \geq \frac{1}{20} \cdot \frac{1}{2} = \frac{1}{40}.$$

With these estimates and using, again, $\ell \leq k \leq \frac{n}{320}$, we obtain $E[Y_{t+1}] \leq 1 + (1 + \frac{8\ell}{n})Y_t - \frac{1}{40}Y_t \leq 1 + Y_t$ and thus $E[Y_{t+1} - Y_t] \leq 1$.

In summary, we have now shown that the process Y_t satisfies $E[Y_{t+1} - Y_t \mid Y_t < Y_{\max}] \leq 2$. We note that $Y_0 \leq 1$ with probability one. For the sake of the argument, let us artificially modify the process from the point on when it has reached a state of at least Y_{\max} . So we define (\tilde{Y}_t) by setting $\tilde{Y}_t = Y_t$, if $Y_t < Y_{\max}$ or if $Y_t \geq Y_{\max}$ and $Y_{t-1} < Y_{\max}$, and $\tilde{Y}_t = \tilde{Y}_{t-1}$ otherwise. In other words, (\tilde{Y}_t) is a copy of (Y_t) until it reaches a state of at least Y_{\max} and then does not move anymore. With this trick, we have $E[\tilde{Y}_{t+1} - \tilde{Y}_t] \leq 2$ for all t .

A simple induction and the initial condition $\tilde{Y}_0 \leq 1$ shows that $E[\tilde{Y}_t] \leq 2t+1$ for all t . In particular, for $T = \exp(\frac{1}{8}ck)$, we have $E[Y_T] \leq 2 \exp(\frac{1}{8}ck) + 1$ and, by Markov's inequality,

$$\Pr[\tilde{Y}_T \geq Y_{\max}] \leq \frac{2 \exp(\frac{1}{8}ck) + 1}{Y_{\max}} = \exp(-(1 - o(1))\frac{1}{8}ck).$$

Hence with probability $1 - \exp(-(1 - o(1))\frac{1}{8}ck)$, we have $\tilde{Y}_T < Y_{\max}$. We now condition on this event. By construction of (\tilde{Y}_t) , we have $Y_t < Y_{\max}$, equivalently $D_t > \frac{1}{4}k$, for all $t \in [0..T]$. If $D_t > \frac{1}{4}k$, then the probability that a sample generated in this iteration is the optimum, is at most $\prod_{i=1}^n f_{it} = \prod_{i=1}^n (1 - (1 - f_{it})) \leq \prod_{i=1}^n \exp(-(1 - f_{it})) = \exp(-(n - \|f_t\|_1)) = \exp(-D_t) \leq \exp(-\frac{1}{4}k)$. Assuming $c \leq 1$ again, we see that the probability that the optimum is generated in one of the first T iterations, is at most $2T \exp(-\frac{1}{4}k) = 2 \exp(\frac{1}{8}ck) \exp(-\frac{1}{4}k) = \exp(-(1 - o(1))\frac{1}{8}k)$. This shows the claim. \square

4 An $\Omega(n \log n)$ Lower Bound?

With the exponential lower bound proven in the previous section, the runtime of the cGA on jump functions is well understood, except that the innocent looking lower bound $\Omega(n \log n)$, matching the corresponding upper bound for $k < \frac{1}{20} \ln n$, is still missing. Since Sudholt and Witt [SW16] have proven an $\Omega(n \log n)$ lower bound for the simple unimodal function ONEMAX, which for many EAs is known to be one of the easiest functions with unique global optimum [DJW12, Sud13, Wit13, Doe18a], it would be very surprising if this lower bound would not hold for jump functions as well.

Unfortunately, we do not see any easy way to prove such a lower bound. We strongly believe that the proof of [SW16] can be extended to also include jump functions, but since this proof is truly complicated, we shy away from taking such an effort to prove a result that would be that little surprising. We instead argue here why the usual ‘‘ONEMAX is the easiest case’’ argument fails. While we would not say that it is not a valuable research goal to extend the proof of [SW16] to jump functions, we would much prefer if someone could prove a general $\Omega(n \log n)$ lower bound for all functions with unique global optimum (or disprove this statement).

The true reason why ONEMAX is the easiest optimization problem for many evolutionary algorithms, implicit in all such proofs and explicit in [Doe18a], is that when comparing a run of the evolutionary algorithm on ONEMAX and some other function \mathcal{F} with unique global optimum, then at all times the Hamming distance between the current-best solution and the

optimum in the ONEMAX process is stochastically dominated by the one of the other process. This follows by induction and a coupling argument from the following key insight (here formulated for the $(1 + 1)$ EA only).

Lemma 7. *Let $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{R}$ be some function with unique global optimum x^* and let ONEMAX be the n -dimensional ONEMAX function with unique global optimum $y^* = (1, \dots, 1)$. Let $x, y \in \{0, 1\}^n$ such that $H(x, x^*) \geq H(y, y^*)$, where $H(\cdot, \cdot)$ denotes the Hamming distance. Consider one iteration of the $(1 + 1)$ EA optimizing \mathcal{F} , started with x as parent individual, and denote by x' the parent in the next iteration. Define y' analogously for ONEMAX and y . Then $H(x', x^*) \succeq H(y', y^*)$.*

As a side remark, note that the lemma applied in the special case $\mathcal{F} = \text{ONEMAX}$ shows that the intuitive rule “the closer a search point is to the optimum, the shorter is the optimization time when starting from this search point” holds for optimizing ONEMAX via the $(1 + 1)$ EA.

We now show that a statement like Lemma 7 is not true for the cGA. Since the states of a run of the cGA are the frequency vectors f , the natural extension of the Hamming distance quality measure above is the ℓ_1 -distance $d(f, x^*) = \|f - x^*\|_1 = \sum_{i=1}^n |f_i - x_i^*|$.

Consider now a run of the cGA on an n -dimensional (n even for simplicity) jump function \mathcal{F} with jump size $k \leq n/4$. Consider one iteration starting with the frequency vector $f = \frac{1}{2}\mathbf{1}_n$. For comparison, consider one iteration of the cGA optimizing ONEMAX, starting with a frequency vector $g \in [0, 1]^n$ such that half the entries of g are equal to $\frac{1}{n} + \frac{1}{\mu}$ and the other half equals $1 - \frac{1}{n} - \frac{1}{\mu}$. Let us take $\mu = n$ for simplicity. Note that both \mathcal{F} and ONEMAX have the same unique global optimum $x^* = y^* = (1, \dots, 1)$.

We obviously have $d(f, x^*) \geq d(g, y^*)$, since both numbers are equal to $\frac{n}{2}$. Let f', g' be the frequency vectors after one iteration. Since with probability $1 - \exp(-\Omega(n))$, both search points sampled in the jump process have between $\frac{n}{4}$ and $\frac{3}{4n}$ ones, their jump fitnesses equal their ONEMAX fitnesses. Consequently, we may apply Lemma 5 from [Dro06] and see that $E[d(f', x^*)] \leq \frac{n}{2} - \Omega(\sqrt{n})$. For the ONEMAX process, however, denoting the two search points generated in this iteration by x^1 and x^2 , we see

$$\begin{aligned} E[\|g - g'\|_1]^2 &\leq E[\|g - g'\|_1^2] = E\left[\left(\sum_{i=1}^n (x_i^1 - x_i^2)\right)^2\right] \\ &= \text{Var}\left[\sum_{i=1}^n (x_i^1 - x_i^2)\right] = O(1) \end{aligned}$$

and hence $E[d(g', y^*)] \geq d(g', y) - O(1) = \frac{n}{2} - O(1)$. Consequently, we cannot have $d(f', x^*) \succeq d(g', y^*)$.

We note that a second imaginable domination result is also not true. Assume, for simplicity, that we optimize a function \mathcal{F} with unique global maximum equal to $(1, \dots, 1)$ and the function ONEMAX via the cGA with same parameter setting. If $f \leq g$ (component-wise), and f' is the frequency vector resulting from one iteration optimizing \mathcal{F} starting with f and g' is the frequency vector resulting from one iteration optimizing ONEMAX starting with g , then in general we do not have $f'_i \preceq g'_i$ for all $i \in [1..n]$.

As counter-example, let $f = (\frac{1}{2}, \frac{1}{n}, \dots, \frac{1}{n})$, but now $g = \frac{1}{2}\mathbf{1}_n$. Clearly, $f \leq g$. We now consider the results of one iteration of the cGA, always with the ONEMAX function as objective. When performing one iteration of the cGA on ONEMAX started with f , and denoting the two samples by x^1 and x^2 and their quality difference in all but the first bit by $\Delta = \|x^1_{[2..n]}\|_1 - \|x^2_{[2..n]}\|_1$, then the resulting frequency vector f' satisfies

$$\begin{aligned} \Pr[f'_1 = \frac{1}{2} + \frac{1}{\mu}] &= \Pr[x^1_1 \neq x^2_1] (\frac{1}{2} \Pr[\Delta \notin \{-1, 0\}] + \Pr[\Delta \in \{-1, 0\}]) \\ &= \Pr[x^1_1 \neq x^2_1] (\frac{1}{2} + \frac{1}{2} \Pr[\Delta \in \{-1, 0\}]). \end{aligned} \quad (2)$$

Since $\Pr[\Delta \in \{-1, 0\}] \geq \Pr[\|x^1_{[2..n]}\|_1 = \|x^2_{[2..n]}\|_1 = 0] = (1 - \frac{1}{n})^{2(n-1)} \geq 1/e^2$, we have $\Pr[f'_1 = \frac{1}{2} + \frac{1}{\mu}] \geq \frac{1}{4} + \frac{1}{4e^2}$.

When starting the iteration with g , the resulting frequency vector g' satisfies an equation analogous to (2), but now Δ is the difference of two binomial distributions with parameters $n - 1$ and $\frac{1}{2}$. Hence, we have $\Pr[\Delta \in \{-1, 0\}] = O(n^{-1/2})$, see, e.g., [Doe18b, Lemma 4.13] for this elementary estimate, and thus $\Pr[g'_1 = \frac{1}{2} + \frac{1}{\mu}] = \frac{1}{4} + o(1)$, disproving that $f'_1 \preceq g'_1$.

In summary, the richer mechanism of building a probabilistic model of the search space in the cGA (as opposed to using a population in EAs) makes it hard to argue that ONEMAX is the easiest function for the cGA. This, in particular, has the consequence that lower bounds for the runtime of the cGA on ONEMAX cannot be easily extended to other functions with a unique global optimum.

5 Conclusion

The main result of this work is an $\exp(\Omega(k))$ lower bound for the runtime of the cGA on jump functions with jump size k , regardless of the hypothetical population size μ . This in particular shows that the result of Hasenöhrl and Sutton [HS18] cannot be improved by running the cGA with a hypothetical population size that is sub-exponential in k .

What is noteworthy in our proof is that it does not require a distinction between the different cases that frequencies reach boundary values or not (as in, e.g., the highly technical lower bound proof for ONEMAX in [SW16]). It seems to be an interesting direction for future research to find out to what extent such an approach can be used also for other lower bound analyses.

As a side result, we observed that two natural domination arguments that could help showing that ONEMAX is the easiest function for the cGA are not true. For this reason, the natural lower bound of $\Omega(n \log n)$ remains unproven. Proving it, or even better, proving that $\Omega(n \log n)$ is a lower bound for the runtime of the cGA on any function $\mathcal{F} : \{0, 1\}^n \rightarrow \mathbb{R}$ with a unique global optimum, remain challenging open problems.

References

- [AAD⁺13] Peyman Afshani, Manindra Agrawal, Benjamin Doerr, Carola Doerr, Kasper Green Larsen, and Kurt Mehlhorn. The query complexity of finding a hidden permutation. In *Space-Efficient Data Structures, Streams, and Algorithms*, volume 8066 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2013. Full version available online at <http://eccc.hpi-web.de/report/2012/087/>.
- [AD11] Anne Auger and Benjamin Doerr, editors. *Theory of Randomized Search Heuristics*. World Scientific Publishing, 2011.
- [ADY19] Denis Antipov, Benjamin Doerr, and Quentin Yang. The efficiency threshold for the offspring population size of the (μ, λ) EA. In *Genetic and Evolutionary Computation Conference, GECCO 2019*. ACM, 2019. To appear.
- [AW09] Gautham Anil and R. Paul Wiegand. Black-box search by elimination of fitness functions. In *Foundations of Genetic Algorithms, FOGA 2009*, pages 67–78. ACM, 2009.
- [BDK16] Maxim Buzdalov, Benjamin Doerr, and Mikhail Kever. The unrestricted black-box complexity of jump functions. *Evolutionary Computation*, 24:719–744, 2016.
- [COY17] Dogan Corus, Pietro Simone Oliveto, and Donya Yazdani. On the runtime analysis of the Opt-IA artificial immune system. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 83–90. ACM, 2017.

- [COY18] Dogan Corus, Pietro Simone Oliveto, and Donya Yazdani. Fast artificial immune systems. In *Parallel Problem Solving from Nature, PPSN 2018*, pages 67–78. Springer, 2018.
- [DFK⁺16] Duc-Cuong Dang, Tobias Friedrich, Timo Kötzing, Martin S. Krejca, Per Kristian Lehre, Pietro Simone 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.
- [DFK⁺18] Duc-Cuong Dang, Tobias Friedrich, Timo Kötzing, Martin S. Krejca, Per Kristian Lehre, Pietro Simone Oliveto, Dirk Sudholt, and Andrew M. Sutton. Escaping local optima using crossover with emergent diversity. *IEEE Transactions on Evolutionary Computation*, 22:484–497, 2018.
- [DJW02] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.
- [DJW06] Stefan Droste, Thomas Jansen, and Ingo Wegener. Upper and lower bounds for randomized search heuristics in black-box optimization. *Theory of Computing Systems*, 39:525–544, 2006.
- [DJW12] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64:673–697, 2012.
- [DLMN17] 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.
- [Doe18a] Benjamin Doerr. Better runtime guarantees via stochastic domination. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2018*, pages 1–17. Springer, 2018.
- [Doe18b] Benjamin Doerr. Probabilistic tools for the analysis of randomized optimization heuristics. *CoRR*, abs/1801.06733, 2018.
- [Doe19] Benjamin Doerr. A tight runtime analysis for the cGA on jump functions—EDAs can cross fitness valleys at no extra cost. In *Genetic and Evolutionary Computation Conference, GECCO 2019*. ACM, 2019. To appear. Available at <https://arxiv.org/abs/1903.10983>.

- [Dro06] Stefan Droste. A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing*, 5:257–283, 2006.
- [GW17] Christian Gießen and Carsten Witt. The interplay of population size and mutation probability in the $(1 + \lambda)$ EA on OneMax. *Algorithmica*, 78:587–609, 2017.
- [HLG99] Georges R. Harik, Fernando G. Lobo, and David E. Goldberg. The compact genetic algorithm. *IEEE Transactions on Evolutionary Computation*, 3:287–297, 1999.
- [HS18] Václav Hasenöhrl and Andrew M. Sutton. On the runtime dynamics of the compact genetic algorithm on jump functions. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 967–974. ACM, 2018.
- [Jan13] Thomas Jansen. *Analyzing Evolutionary Algorithms - The Computer Science Perspective*. Natural Computing Series. Springer, 2013.
- [JJW05] Thomas Jansen, Kenneth A. De Jong, and Ingo Wegener. On the choice of the offspring population size in evolutionary algorithms. *Evolutionary Computation*, 13:413–440, 2005.
- [JW02] Thomas Jansen and Ingo Wegener. The analysis of evolutionary algorithms—a proof that crossover really can help. *Algorithmica*, 34:47–66, 2002.
- [KW17] Martin S. Krejca and Carsten Witt. Lower bounds on the run time of the univariate marginal distribution algorithm on onemax. In *Foundations of Genetic Algorithms, FOGA 2017*, pages 65–79. ACM, 2017.
- [KW18] Martin S. Krejca and Carsten Witt. Theory of estimation-of-distribution algorithms. *CoRR*, abs/1806.05392, 2018.
- [LSW18] Johannes Lengler, Dirk Sudholt, and Carsten Witt. Medium step sizes are harmful for the compact genetic algorithm. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1499–1506. ACM, 2018.
- [NW10] Frank Neumann and Carsten Witt. *Bioinspired Computation in Combinatorial Optimization – Algorithms and Their Computational Complexity*. Springer, 2010.

- [OW12] Pietro Simone Oliveto and Carsten Witt. Erratum: Simplified drift analysis for proving lower bounds in evolutionary computation. *CoRR*, abs/1211.7184, 2012.
- [Sud13] Dirk Sudholt. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Trans. Evolutionary Computation*, 17:418–435, 2013.
- [SW16] Dirk Sudholt and Carsten Witt. Update strength in EDAs and ACO: How to avoid genetic drift. In *Genetic and Evolutionary Computation Conference, GECCO 2016*, pages 61–68. ACM, 2016.
- [Wit13] Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing*, 22:294–318, 2013.
- [Wit18] Carsten Witt. Domino convergence: why one should hill-climb on linear functions. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1539–1546. ACM, 2018.