

Significance-based Estimation-of-Distribution Algorithms

Benjamin Doerr and Martin S. Krejca

Abstract—*Estimation-of-distribution algorithms* (EDAs) are randomized search heuristics that create a probabilistic model of the solution space, which is updated iteratively, based on the quality of the solutions sampled according to the model. As previous works show, this iteration-based perspective can lead to erratic updates of the model, in particular, to bit-frequencies approaching a random boundary value. In order to overcome this problem, we propose a new EDA based on the classic compact genetic algorithm (cGA) that takes into account a longer history of samples and updates its model only with respect to information which it classifies as statistically significant. We prove that this *significance-based compact genetic algorithm* (sig-cGA) optimizes the commonly regarded benchmark functions ONEMAX, LEADINGONES, and BINVAL all in quasilinear time, a result shown for no other EDA or evolutionary algorithm so far. For the recently proposed scGA – an EDA that tries to prevent erratic model updates by imposing a bias to the uniformly distributed model – we prove that it optimizes ONEMAX only in a time exponential in its hypothetical population size. Similarly, we show that the convex search algorithm cannot optimize ONEMAX in polynomial time.

Index Terms—theory; estimation-of-distribution algorithm; run time analysis

I. INTRODUCTION

Estimation-of-distribution algorithms (EDAs; [2]) are nature-inspired heuristics, similar to *evolutionary algorithms* (EAs). In contrast to EAs, which maintain an explicit set of solutions, EDAs optimize a function by evolving a probabilistic model of the solution space. Iteratively, an EDA uses its probabilistic model in order to generate samples and make observations from them. It then updates its model based on these observations, where an algorithm-specific parameter determines how strong the changes to the model in each iteration are.

For an EDA to succeed in optimization, it is important that its model is changed over time in a way that better solutions are sampled more frequently. However, due to the randomness in sampling, the model should not be changed too drastically

B. Doerr is with the Laboratoire d'Informatique (LIX), CNRS, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

M. S. Krejca is with the Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

This work was supported by a public grant as part of the Investissement d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH, in a joint call with Gaspard Monge Program for optimization, operations research and their interactions with data sciences. This publication is based upon work from COST Action CA15140, supported by COST.

This paper extends our results published at GECCO'18 [1].

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

in a single iteration in order to prevent wrong updates from having a long-lasting impact.

The theoretical analysis of EDAs has recently gained momentum (see, e.g., the survey [3]) and has clearly shown that this trade-off between convergence speed and accumulation of erratic updates can be delicate and non-trivial to understand. Among the most relevant works, Sudholt and Witt [4] and Krejca and Witt [5] prove lower bounds of the expected run times of three common EDAs on the benchmark function ONEMAX. In simple words, these bounds show that if the update parameter for the model is too large, the model converges too quickly and very likely to a wrong model; consequently, it then takes a long time to find the optimum. On the other hand, if the parameter is too small, then the model converges to the correct model but does so slowly. More formally, Sudholt and Witt [4] prove a lower bound of $\Omega(K\sqrt{n} + n \log n)$ for the 2-MMAS_{ib} and the cGA, where $1/K$ is the step size of the algorithm, and Krejca and Witt [5] prove a lower bound of $\Omega(\lambda\sqrt{n} + n \log n)$ for the UMDA, where λ is the population size of the algorithm. These results show that choosing the parameter with a value of $\omega(\sqrt{n} \log n)$ has no benefit. Further, it has been recently shown by Lengler et al. [6] that the run time of the cGA on ONEMAX is $\Omega(K^{1/3}n + n \log n)$ for $K = O(\sqrt{n}/(\log n \log \log n))$. Together with the results from Sudholt and Witt [4], this implies a bimodal behavior in the run time with respect to K if $K = \Omega(\log n) \cap O(\sqrt{n} \log n)$, showing that the run time is sensitive to the parameter choice.

Friedrich et al. [7] also discuss the problem of how to choose the update strength. They consider a class of EDAs optimizing functions over bit strings of length n that all current theoretical results fall into, named n -Bernoulli- λ -EDA. The model of such EDAs uses one variable per bit of a bit string, resulting in a probability vector τ of length n called the *frequency vector*. In each iteration, a bit string x is sampled bit-wise independently and independent of any other sample such that bit x_i is 1 with probability (*frequency*) τ_i and 0 otherwise.

Friedrich et al. [7] consider two different properties of such EDAs, namely *balanced* and *stable*. Intuitively, a *balanced* EDA does not change a frequency τ_i in expectation if the fitness function has no preference for 0s or 1s at position i . A *stable* EDA keeps a frequency, in such a scenario, close to $1/2$. Friedrich et al. [7] then prove that an n -Bernoulli- λ -EDA cannot be both balanced and stable. They also prove that all commonly theoretically analyzed EDAs are balanced. This means that the frequencies will always move toward 0 or 1, even if there is no bias from the objective function.

Motivated by these results, Friedrich et al. [7] propose an EDA (called *scGA*) that is stable (but not balanced) by

introducing an artificial bias into the update process that should counteract the bias of a balanced EDA. However, we prove that this approach fails badly on the standard benchmark function ONEMAX (Thm. 9). We note that a similar bias towards the middle frequency of $1/2$ was proven for a binary differential evolution algorithm by Zheng et al. [8]. Similar to the situation of the scGA, their run time results (partially relying on mean-field assumptions) indicate that LEADINGONES is optimized in a number of generations that is linear in the problem size n . This gives an $O(n \log n)$ number of function evaluations when using a logarithmic population size (and smaller population sizes are provably not successful). For ONEMAX, the results are less conclusive, but they indicate a run time exponential in the population size can occur.

The results of Friedrich et al. [7], Sudholt and Witt [4], Krejca and Witt [5], and Lengler et al. [6] draw the following picture: for a balanced EDA, there exists some inherent noise in the update. Thus, if the parameter responsible for the update of the probabilistic model is large and the speed of convergence high, the algorithm only uses a few samples before it converges. During this time, the noise introduced by the balance-property may not be overcome, resulting in the probabilistic model converging to an incorrect one, as the algorithms are not stable. Hence, the parameter has to be chosen sufficiently small in order to guarantee convergence to the correct model, resulting in a slower optimization time.

As we shall argue in this work, the reason for this dilemma is that EDAs only use information from a single iteration when performing an update. Thus, the decision of whether and how a frequency should be changed has to be made on the spot, which may result in harmful decisions.

To overcome these difficulties, we propose a conceptually new EDA that has some access to the search history and updates the model only if there is sufficient reason. The *significance-based compact genetic algorithm* (sig-cGA) stores for each position the history of bits of good solutions so far. If it detects that either statistically significantly more 1s than 0s or vice versa were sampled, it changes the corresponding frequency, otherwise not. Thus, the sig-cGA only performs an update when it has proof that it makes sense. This sets it apart from the other EDAs analyzed so far.

We prove that the sig-cGA is able to optimize LEADINGONES, ONEMAX, and BINVAL in $O(n \log n)$ function evaluations in expectation and with high probability (Thms. 5 and 8 and Cor. 6), which has not been proven before for any other EDA or classical EA (for further details, see Table I).

We also observe that the analysis for LEADINGONES can easily be modified to also show an $O(n \log n)$ run time for the *binary value* function BINVAL, which is a linear function with exponentially growing coefficients. This result is interesting in that it indicates that the sig-cGA has asymptotically the same run time on BINVAL and ONEMAX. In contrast, for the classic cGA it is known [25] that the run times on ONEMAX and BINVAL differ significantly.

We then show that two previously regarded algorithms which solve LEADINGONES in $O(n \log n)$ time behave poorly on ONEMAX. The run time of the scGA proposed in [7] is $\Omega(2^{\Theta(\min\{n, 1/\rho\})})$ (Thm. 9), where $1/\rho$ is an algorithm-

specific parameter controlling the strength of the model update and denotes the hypothetical population size of the algorithm. For the convex search algorithm (CSA) proposed in [18], we prove that the run time, even when adding suitable restart schemes, is asymptotically larger than any polynomial (Thm. 10). These results, together with the large number of existing results, suggest that none of the previously known algorithms performs exceptionally well on both ONEMAX and LEADINGONES.

These results, the positive ones for the sig-cGA using a longer history of the search process and the negative ones for other algorithms not exploiting a longer history, suggest that a fruitful direction for the future development of the field of evolutionary computation (EC; not restricted to theory) is the search for algorithms that enrich the classic generational approaches with mechanisms that profit from regarding more than one generation. We discuss this in more detail in the conclusions of this paper. We note that, from a practical point of view, our algorithm not only showed a performance not seen so far with other algorithms, it is also easier to use since, unlike with most other EDAs, the delicate choice of the update strength is obsolete.

This paper extends our previous results on the sig-cGA [1] by proving an upper bound of the sig-cGA on BINVAL (Cor. 6) and a lower bound of the CSA on ONEMAX (Thm. 10). Due to space constraints, some proofs can be found in the supplementary material.

II. PRELIMINARIES

In this paper, we consider the maximization of pseudo-Boolean functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$, where n is a positive integer (fixed for the remainder of this work). We call f a *fitness function*, an element $x \in \{0, 1\}^n$ an *individual*, and, for an $i \in [n] := [1, n] \cap \mathbb{N}$, we denote the i th bit of x by x_i . When talking about *run time*, we always mean the number of fitness function evaluations of an algorithm until an optimum is sampled for the first time.

In our analysis, we regard the two classic benchmark functions ONEMAX (OM) and LEADINGONES (LO) defined by

$$\text{OM}(x) = \sum_{i \in [n]} x_i \quad \text{and} \quad (1)$$

$$\text{LO}(x) = \sum_{i \in [n]} \prod_{j \in [i]} x_j. \quad (2)$$

Intuitively, OM returns the number of 1s of an individual, whereas LO returns the longest sequence of consecutive 1s, starting from the left. Note that the all-1s bit string is the unique global optimum for both functions.

In Table I, we state the asymptotic run times of many algorithms on these two functions. We note that (i) the black-box complexity of OM is $\Theta(n/\log n)$, see [26], [27], and (ii) the black-box complexity of LO is $\Theta(n \log \log n)$, see [28], however, all black-box algorithms witnessing these run times are highly artificial. Consequently, $\Theta(n \log n)$ appears to be the best run time to aim for these two problems.

For our calculations, we shall regularly use the following well-known variance-based additive Chernoff bounds (see, e.g., the respective Chernoff bound in [29]).

Table I

EXPECTED RUN TIMES (NUMBER OF FITNESS EVALUATIONS) OF VARIOUS ALGORITHMS UNTIL THEY FIRST FIND AN OPTIMUM FOR THE TWO FUNCTIONS OM (EQ. (1)) AND LO (EQ. (2)). FOR OPTIMAL PARAMETER SETTINGS, MANY ALGORITHMS HAVE A RUN TIME OF $\Theta(n \log n)$ FOR OM AND OF $\Theta(n^2)$ FOR LO. WE NOTE THAT THE $(1 + (\lambda, \lambda))$ GA HAS AN $o(n \log n)$ RUN TIME ON OM (AND EVEN LINEAR RUN TIME WITH A DYNAMIC PARAMETER CHOICE), BUT WE DO NOT SEE WHY IT SHOULD HAVE A PERFORMANCE BETTER THAN QUADRATIC ON LO.

Algorithm	OM	constraints	LO	constraints
(1 + 1) EA	$\Theta(n \log n)$ [9]	none	$\Theta(n^2)$ [9]	none
$(\mu + 1)$ EA	$\Theta(\mu n + n \log n)$ [10]	$\mu = O(\text{poly}(n))$	$\Theta(\mu n \log n + n^2)$ [10]	$\mu = O(\text{poly}(n))$
$(1 + \lambda)$ EA	$\Theta\left(n \log n + \frac{\lambda n \log \log \lambda}{\log \lambda}\right)$ [11], [12] ^a	$\lambda = O(n^{1-\varepsilon})$	$\Theta(n^2 + \lambda n)$ [11]	$\lambda = O(\text{poly}(n))$
$(\mu + \lambda)$ EA	$\Theta\left(\frac{n \log n}{\lambda} + \frac{n}{\lambda/\mu} + \frac{n \log^+ \log^+ \lambda/\mu}{\log^+ \lambda/\mu}\right)$ [16]	$\log^+ x := \max\{1, \log x\}$	$\Omega\left(n^2 + \frac{\lambda n}{\log(\lambda/n)}\right)$ [13]	–
$(1 + (\lambda, \lambda))$ GA	$\Theta\left(\max\left\{\frac{n \log n}{\lambda}, \frac{n \lambda \log \log \lambda}{\log \lambda}\right\}\right)$ [17]	$p = \frac{\lambda}{n}, c = \frac{1}{\lambda}$	unknown	–
CSA	$\Omega(n^c)$ [Thm. 10]	$c > 0$	$O(n \log n)$ [18]	$\mu \geq 8 \ln((4n + 6)n)$, restarts
UMDA/PBIL ^b	$\Omega(\lambda \sqrt{n} + n \log n)$ [5] $O(\lambda n)$ [21], [19]	$\mu = \Theta(\lambda)$ $\mu = \Omega(\log n) \cap O(\sqrt{n})$, $\lambda = \Omega(\mu)$ or $\mu = \Omega(\sqrt{n} \log n)$, $\mu = \Theta(\lambda)$ or $\mu = \Omega(\log n) \cap o(n)$, $\mu = \Theta(\lambda)$	$O(n \lambda \log \lambda + n^2)$ [19], [20]	$\lambda = \Omega(\log n)$, $\mu = \Theta(\lambda)$
cGA/2-MMAS _{ib}	$\Omega\left(\frac{\sqrt{n}}{\rho} + n \log n\right)$ [4] $O\left(\frac{\sqrt{n}}{\rho}\right)$ [4]	$\frac{1}{\rho} = O(\text{poly}(n))$ $\frac{1}{\rho} = \Omega(\sqrt{n} \log n) \cap O(\text{poly}(n))$	unknown	–
1-ANT	$O(n^2)$ [22] ^c	$\rho = \Omega(n^{-1+\varepsilon})$	$O(n^2 \cdot (6e)^{1/(n\rho)})$ [23] $2^{\Omega(\min\{n, 1/(n\rho)\})}$ [23]	none none
MMAS*	$O\left(\frac{n \log n}{\rho}\right)$ [24]	$\rho = O(1)$	$O\left(n^2 + \frac{n \log n}{\rho}\right)$ [24] $\Omega\left(n^2 + \frac{n}{-\rho \log(2\rho)}\right)$ [24]	$\rho = O(1)$ $\rho = 1/\text{poly}(n)$
scGA	$\Omega(\min\{2^{\Theta(n)}, 2^{c/\rho}\})$ [Thm. 9]	$1/\rho = \Omega(\log n)$, $a = \Theta(\rho)$, $d = \Theta(1)$, $c > 0$	$O(n \log n)$ [7]	$1/\rho = \Theta(\log n)$, $a = O(\rho)$, $d = \Theta(1)$
sig-cGA (Alg. 1)	$O(n \log n)$ [Thm. 8]	$\varepsilon > 12$	$O(n \log n)$ [Thm. 5]	$\varepsilon > 12$

^aBetter run time bounds for the $(1 + \lambda)$ EA are known if the mutation rate is (i) fitness-dependent [13], (ii) self-adjusting [14], or (iii) self-adaptive [15].

^bThe results shown for PBIL are the results of UMDA if not mentioned otherwise, since the latter is a special case of the former.

^cFor $\rho \geq (n - 2)/(3n - 2)$, the algorithm simulates the $(1 + 1)$ EA and has a run time of $\Theta(n \log n)$.

Theorem 1 (Variance-based Additive Chernoff Bounds). *Let X_1, \dots, X_n be independent random variables such that, for all $i \in [n]$, $E[X_i] - 1 \leq X_i \leq E[X_i] + 1$. Further, let $X = \sum_{i=1}^n X_i$ and $\sigma^2 = \sum_{i=1}^n \text{Var}[X_i] = \text{Var}[X]$. Then, for all $\lambda \geq 0$, abbreviating $m = \min\{\lambda^2/\sigma^2, \lambda\}$,*

$$\Pr[X \geq E[X] + \lambda] \leq e^{-\frac{1}{3}m} \text{ and } \Pr[X \leq E[X] - \lambda] \leq e^{-\frac{1}{3}m}.$$

We say that an event A occurs with high probability (w.h.p.) if there is a $c = \Omega(1)$ such that $\Pr[A] \geq (1 - n^{-c})$.

Last, we use the \circ operator to denote string concatenation. For a bit string $H \in \{0, 1\}^*$, let $|H|$ denote its length, $\|H\|_0$ its number of 0s, $\|H\|_1$ its number of 1s, and, for a $k \in [|H|]$, let $H[k]$ denote the *last* k bits in H . In addition to that, \emptyset denotes the empty string.

III. THE SIGNIFICANCE-BASED COMPACT GENETIC ALGORITHM

Before we present our algorithm sig-cGA in detail in Section III-A, we provide more information about the *compact genetic algorithm* (cGA [30]), which the sig-cGA as well as the scGA are based on.

The cGA is a univariate EDA, that is, it assumes independence of the bits in the search space. As such, it keeps a vector of probabilities $(\tau_i)_{i \in [n]}$ (the *frequency vector*). In

Algorithm 1: The sig-cGA with parameter ε and significance function sig_ε (eq. (3)) optimizing f

```

1  $t \leftarrow 0$ ;
2 for  $i \in [n]$  do  $\tau_i^{(t)} \leftarrow \frac{1}{2}$  and  $H_i \leftarrow \emptyset$ ;
3 repeat
4    $x, y \leftarrow$  offspring sampled with respect to  $\tau^{(t)}$ ;
5    $x \leftarrow$  winner of  $x$  and  $y$  with respect to  $f$ ;
6   for  $i \in [n]$  do
7      $H_i \leftarrow H_i \circ x_i$ ;
8     if  $\text{sig}_\varepsilon(\tau_i^{(t)}, H_i) = \text{up}$  then  $\tau_i^{(t+1)} \leftarrow 1 - 1/n$ ;
9     else if  $\text{sig}_\varepsilon(\tau_i^{(t)}, H_i) = \text{down}$  then
10       $\tau_i^{(t+1)} \leftarrow 1/n$ ;
11     else  $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)}$ ;
12     if  $\tau_i^{(t+1)} \neq \tau_i^{(t)}$  then  $H_i \leftarrow \emptyset$ ;
12    $t \leftarrow t + 1$ ;
13 until termination criterion met;
```

each iteration, two individuals (*offspring*) are sampled in the following way with respect to τ : for an individual $x \in \{0, 1\}^n$, we have $x_i = 1$ with probability τ_i , and $x_i = 0$ with probability $1 - \tau_i$, independently of any τ_j with $j \neq i$.

After sampling, the frequency vector is updated with respect

to a fitness-based ranking of the offspring. The process of choosing how the offspring are ranked is called *selection*. Let x and y denote both offspring of the cGA during an iteration. Given a fitness function f , we rank x above y if $f(x) > f(y)$ (as we maximize), and we rank y above x if $f(y) > f(x)$. If $f(x) = f(y)$, we rank them randomly. The higher-ranked individual is called the *winner*, the other individual the *loser*. Assume that x is the winner. The cGA then changes a frequency τ_i then with respect to the difference $x_i - y_i$ by a value of ρ (where $1/\rho$ is usually referred to as population size). Hence, no update is performed if the bit values are identical, and the frequency is moved to the bit value of the winner. In order to prevent a frequency τ_i getting stuck at 0 or 1,¹ the cGA usually caps its frequency to the range $[1/n, 1 - 1/n]$, as is common practice. This way, a frequency can get close to 0 or 1, but it is always possible to sample 0s and 1s.

Consider a position i and any two individuals x and y that are identical except for position i . Assume that $x_i > y_i$. If the probability that x is the winner of the selection is higher than y being the winner, we speak of a *bias in selection* (for 1s) at position i . Analogously, we speak of a bias for 0s if the probability that y wins is higher than the probability that x wins. Usually, a fitness function introduces a bias into the selection and thus into the update.

A. Detailed Description of the sig-cGA

Similar to the cGA, our new algorithm – the *significance-based compact genetic algorithm* (sig-cGA; Alg. 1) – also samples two offspring each iteration. However, in contrast to the cGA, it keeps a history of bit values for each position and only performs an update when a statistical significance within a history occurs. This approach better aligns with the intuitive reasoning that an update should only be performed if there is valid evidence for a different frequency being better suited for sampling good individuals.

In more detail, for each bit position $i \in [n]$, the sig-cGA keeps a history $H_i \in \{0, 1\}^*$ of all the bits sampled by the winner of each iteration since the last time τ_i changed – the last bit denoting the latest entry. Observe that if there is no bias in selection at position i , the bits sampled by τ_i follow a binomial law with $|H_i|$ tries and a success probability of τ_i . We call this our *hypothesis*. If we happen to find a sequence (starting from the latest entry) in H_i that significantly deviates from the hypothesis, we update τ_i with respect to the bit value that occurred significantly, and we reset the history. We only use the following three frequency values:

- $1/2$: starting value;
- $1/n$: significance for 0s was detected;
- $1 - 1/n$: significance for 1s was detected.

We formalize *significance* by defining the threshold s to overcome. For all $\varepsilon, \mu \in \mathbb{R}^+$, where μ is the expected value of our hypothesis and ε is an algorithm-specific parameter:

$$s(\varepsilon, \mu) = \varepsilon \max\{\sqrt{\mu \ln n}, \ln n\}.$$

¹A frequency τ_i at one of these two values results in the offspring only having the same bit value at position i . Thus, the cGA would not change τ_i anymore.

Note that $\sqrt{\mu}$ basically describes the standard deviation of our hypothesis, and the logarithmic factor increases this value such that a deviation does not happen w.h.p. The maximum ensures that we consider at least logarithmically many samples before we conclude that we found a significance, eliminating wrong updates due to small samples sizes w.h.p. The parameter ε effectively turns into the exponent of the w.h.p. bounds. Thus, a larger value of ε decreases the probability of detecting a false significance by a polynomial amount. However, it also increases the number of samples necessary in order to change a frequency. This results in a linear factor of ε in the run time. We provide more details on how ε should be chosen at the end of this subsection (after Cor. 3).

We say, for an $\varepsilon \in \mathbb{R}^+$, that a binomially distributed random variable X deviates significantly from a hypothesis $Y \sim \text{Bin}(k, \tau)$, where $k \in \mathbb{N}^+$ and $\tau \in [0, 1]$, if there exists a $c = \Omega(1)$ such that $\Pr[|X - E[Y]| \leq s(\varepsilon, E[Y])] \leq n^{-c}$.

We now state our significance function $\text{sig}_\varepsilon: \{\frac{1}{n}, \frac{1}{2}, 1 - \frac{1}{n}\} \times \{0, 1\}^* \rightarrow \{\text{up}, \text{stay}, \text{down}\}$, which scans a history for a significance. However, it does not scan the entire history but multiple subsequences of a history (always starting from the latest entry). This is done in order to quickly notice a change from an insignificant history to a significant one. Further, we only check in steps of powers of 2, as this is faster than checking each subsequence and we can be off from any length of a subsequence by a constant factor of at most 2. More formally, for all $p \in \{\frac{1}{n}, \frac{1}{2}, 1 - \frac{1}{n}\}$ and all $H \in \{0, 1\}^*$, we define, with ε being a parameter of the sig-cGA, recalling that $H[k]$ denotes the last k bits of H ,

$$\text{sig}_\varepsilon(p, H) = \begin{cases} \text{up} & \text{if } p \in \{\frac{1}{n}, \frac{1}{2}\} \wedge \exists m \in \mathbb{N}: \\ & \|H[2^m]\|_1 \geq 2^m p + s(\varepsilon, 2^m p), \\ \text{down} & \text{if } p \in \{\frac{1}{2}, 1 - \frac{1}{n}\} \wedge \exists m \in \mathbb{N}: \\ & \|H[2^m]\|_0 \geq 2^m p + s(\varepsilon, 2^m p), \\ \text{stay} & \text{else.} \end{cases} \quad (3)$$

We stop at the first (minimum) length 2^m that yields a significance. Thus, we check a history H in each iteration at most $\log_2 |H|$ times.

We now prove that the sig-cGA does not detect a significance at a position with no bias in selection (i.e., a *false significance*) w.h.p.

Lemma 2. *Consider the sig-cGA (Alg. 1) with $\varepsilon \geq 1$. Further, consider a position $i \in [n]$ and an iteration such that the distribution X of 1s of H_i follows a binomial law with k tries and success probability τ_i , i.e., there is no bias in selection at position i . Then τ_i changes in this iteration with a probability of at most $n^{-\varepsilon/3} \log_2 k$.*

Lemma 2 bounds the probability of detecting a false significance within a single iteration, assuming no bias in selection. The following corollary trivially bounds the probability of detecting a false significance within any number of iterations.

Corollary 3. *Consider the sig-cGA (Alg. 1) with $\varepsilon \geq 1$ running for k iterations such that, during each iteration, for each $i \in [n]$, a 1 is added to H_i with probability τ_i . Then*

at least one frequency changes during an interval of $k' \leq k$ iterations with a probability of at most $k'n^{1-\varepsilon/3} \log_2 k$.

Intuitively, this corollary states that ε should be chosen such that the term $k'n^{1-\varepsilon/3} \log_2 k$ represents the desired error probability, where k' is the length of an interval such that a frequency only drops with the error probability. Assuming that one chooses $k' = \Theta(n^r)$ for some constant $r > 0$ and desires an error probability of at most n^{-q} for some constant $q > 0$ (ignoring constant factors and the logarithm), it makes sense to choose $\varepsilon \geq 3(r + 1 + q)$.

B. Efficient Implementation of the sig-cGA

Recall that, in order to save on the computational cost of checking for a significance, we only do so in historic data in lengths of powers of 2. By precomputing the number of 1s in each such interval, checking a single history for a significance can be done in time logarithmically in its length. Note that the update to this precomputed data can also be done in logarithmic time, as each iteration only a single bit is added to the history and thus the number of 1s can only differ by at most one per interval from one iteration to the next. Consequently, the loop of the sig-cGA has a computational cost of $O(\sum_{i=1}^n \log |H_i|)$. Since a history can never be longer than the run time T of the sig-cGA, its total computational cost is $O(nT \log T)$. In comparison, many EAs have an extra cost of $O(n)$ per iteration. Thus, our significance-based approach is only more costly by a factor of $O(\log T)$.

One drawback of the approach above is that the full history needs to be stored. In the supplementary material, we describe a way of condensing a history to a size only logarithmic in the length of the full history. This approach does not allow anymore to access the exact number of 1s (or 0s) in all power-of-two length histories. However, for each $\ell \in [|H_i|]$, it yields the number of 1s in some interval of length ℓ' with $\ell \leq \ell' < 2\ell$. For reasons of readability, we nevertheless regard the original sig-cGA in the subsequent analyses, but it is quite evident that the mildly different accessibility of the history in our condensed implementation does not change our result.

IV. RUN TIME RESULTS FOR LO AND OM

We now prove our main results, that is, upper bounds of $O(n \log n)$ for the expected run time of the sig-cGA on LO and OM. We also note that the optimization process for the binary value function can be analyzed with arguments very similar to those for the LO process. Consequently, we here have an $O(n \log n)$ run time as well. Further, we consider the number of iterations T until the sig-cGA finds the optimal solution. Since it generates two offspring each iteration, the number of fitness function evaluations is at most $2T$.

Note that the sig-cGA treats 1s and 0s symmetrically, that is, it is unbiased in the sense of Lehre and Witt [31]. Hence, all results in this section hold for any type of function as defined in eqs. (1) or (2) where, for any position $i \in [n]$, a bit x_i can be flipped to $1 - x_i$ instead or swapped with another bit x_j .

The following lemma states a useful bound for convex combinations. We use it in order to bound the probability of an event that we decomposed into an event and its complement.

Lemma 4. *Let $\alpha, \beta, x, y \in \mathbb{R}$ such that $x \leq y$ and $\alpha \leq \beta$. Then $\alpha x + (1 - \alpha)y \geq \beta x + (1 - \beta)y$.*

A. Analysis of LO

We show that the frequencies are set to $1 - 1/n$ sequentially from the most significant bit position to the least significant, that is, from left to right. w.h.p., no frequency is decreased until the optimization process is finished. Thus, a frequency τ_i will stay at $1/2$ until all of the frequencies to its left are set to $1 - 1/n$. Then τ_i will become relevant for selection, as all of the frequencies left to it will only sample 1s w.h.p. This results in a significant surplus of 1s being saved at position i , and τ_i will be set to $1 - 1/n$ within $O(\log n)$ iterations and remain there. Then frequency τ_{i+1} becomes relevant for selection. As we need to set n frequencies to $1 - 1/n$, we get a run time of $O(n \log n)$.

Theorem 5. *Consider the sig-cGA (Alg. 1) with $\varepsilon > 12$ being a constant. Its run time on LO is $O(n \log n)$ w.h.p. and in expectation.*

Proof. We split this proof into two parts and start by showing that the run time is $O(n \log n)$ w.h.p. Then we prove the expected run time.

Run time w.h.p. For the first part of the proof, we consider the first $O(n \log n)$ iterations of the sig-cGA and condition on the event that no frequency decreases during this time, i.e., no (false) significance of 0s is detected. Since, for any position $i \in [n]$ in LO, having a 1 is always at least as good as having a 0, a 1 is saved in H_i with a probability of at least τ_i . Hence, by Corollary 3, no frequency decreases in the first $O(n \log n)$ iterations with a probability of at least $1 - O(n^{2-\varepsilon/3} \log^2 n)$. As $\varepsilon > 12$, for an $\varepsilon' > 2$, this probability is at least $1 - O(n^{-\varepsilon'})$, which is w.h.p. In the following, we condition on this event.

We now prove that the history of the leftmost position with a frequency at $1/2$ saves 1s significantly more often than 0s such that the frequency is set to $1 - 1/n$ after $O(\log n)$ iterations. For the second part of the proof, we use a similar argument, but the frequency is at $1/n$, and it takes $O(n \log n)$ steps to get to $1 - 1/n$. Since the calculations for both scenarios are very similar, we combine them.

Consider a position $i \in [n]$ and any of the first $O(n \log n)$ iterations such that $\tau_i \in \{1/n, 1/2\}$ and, for all positions $j < i$, $\tau_j = 1 - 1/n$. Let O denote the event that we save a 1 in H_i this iteration. We derive an upper bound on the probability to detect the significance of 1s in H_i within $O(\log n)$ iterations by calculating a lower bound on the probability of O . Note that the probability of O is the same for each iteration until τ_i is increased, since we condition on no frequency dropping within the first $O(n \log n)$ iterations.

In order to bound $\Pr[O]$, we consider the event A that the bit at position i of the winning individual is not relevant for selection. That is, A denotes the event that at least one of the two offspring during this iteration has a 0 at a position in $[i - 1]$. Thus, if A occurs, a 1 is saved with probability

$\tau_i \in \{1/n, 1/2\}$. Otherwise, a 1 is saved if not two 0s are sampled, which has a probability of $1 - \tau_i^2$. Hence,

$$\Pr[O] = \Pr[A] \cdot \tau_i + \Pr[\bar{A}] \cdot (1 - \tau_i^2),$$

which is a convex combination of τ_i and $1 - \tau_i^2$. By Lemma 4, decreasing $\Pr[\bar{A}]$ (as $1 - \tau_i^2 \geq \tau_i$) results in a lower bound of $\Pr[O]$. Since \bar{A} is equivalent to both offspring having only 1s at positions in $[i - 1]$, we see that

$$\Pr[\bar{A}] = \left(1 - \frac{1}{n}\right)^{2(i-1)},$$

due to our assumption that all frequencies left of position i are at $1 - 1/n$. As this term is minimal for $i = n$, using the well-known inequality $(1 - 1/n)^{n-1} \geq e^{-1}$, we bound $\Pr[\bar{A}] \geq e^{-2}$. Further, noting that $1 - \tau_i^2 \geq (3/2)\tau_i$ for $\tau_i \in \{1/n, 1/2\}$, we bound

$$\begin{aligned} \Pr[O] &\geq (1 - e^{-2}) \cdot \tau_i + e^{-2} \cdot (1 - \tau_i^2) \\ &\geq (1 - e^{-2}) \cdot \tau_i + \frac{3}{2}e^{-2}\tau_i = \left(1 + \frac{1}{2}e^{-2}\right) \cdot \tau_i. \end{aligned}$$

Given our bound on the probability of O , we now bound the probability to detect a significance of 1s in H_i within k iterations. To this end, let $X \sim \text{Bin}(k, (1 + e^{-2}/2)\tau_i)$, and note that X is stochastically dominated by the process of saving 1s at position i . We bound the probability that we do not detect a significance of 1s within k iterations:

$$\begin{aligned} \Pr[X < k\tau_i + s(\varepsilon, k\tau_i)] \\ \leq \Pr[X \leq E[X] - \left(\frac{k}{2}e^{-2}\tau_i - s(\varepsilon, k\tau_i)\right)]. \end{aligned}$$

Note that the minuend is positive for $k > (4/\tau_i)e^4\varepsilon^2 \ln n > \ln n$, which holds due to our assumption $\varepsilon > 12$. Let $c = (4/\tau_i)e^4\varepsilon^2$, and assume $k \geq 4c \ln n$. Thus, $(k/2)e^{-2}\tau_i - s(\varepsilon, k\tau_i) \geq (k/4)e^{-2}\tau_i =: \lambda$ and $\text{Var}[X] = k\tau_i(1 - \tau_i) \geq \lambda$. By Theorem 1, noting that $\lambda^2/\text{Var}[X] \leq \lambda$ and using $\text{Var}[X] \leq k\tau_i$, we bound

$$\begin{aligned} \Pr[X < k\tau_i + s(\varepsilon, k\tau_i)] &\leq \Pr[X \leq E[X] - \frac{k}{4}e^{-2}\tau_i] \\ &\leq e^{-\frac{1}{3} \cdot \frac{\lambda^2}{\text{Var}[X]}} \leq e^{-\frac{1}{3} \cdot \frac{k^2 e^{-4} \tau_i^2}{16k\tau_i}} = e^{-\frac{1}{3} \cdot \frac{ke^{-4}\tau_i}{16}} \\ &\leq n^{-\frac{1}{3} \cdot \frac{ce^{-4}\tau_i}{4}} = n^{-\frac{\varepsilon^2}{3}}. \end{aligned}$$

Hence, τ_i is set to $1 - 1/n$ after $(4/\tau_i)e^4\varepsilon^2 \ln n = O((1/\tau_i) \log n)$ iterations with a probability of at least $1 - n^{-\varepsilon^2/3}$. By applying a union bound over all n different possibilities for index i , we see that each frequency (once all frequencies at positions $[i - 1]$ are at $1 - 1/n$) is set to $1 - 1/n$ within $O((1/\tau_i) \log n)$ with a probability of at least $1 - n^{1-\varepsilon^2/3} \geq 1 - n^{-47}$, since $\varepsilon > 12$, which is w.h.p.

Overall, we assume that no frequency decreases during the first $O(n \log n)$ iterations w.h.p., and we showed that each frequency (at $1/2$) is set to $1 - 1/n$ within $O(\log n)$ iterations w.h.p. once all frequencies to its left are at $1 - 1/n$. Thus, since there are n frequencies, all frequencies are at $1 - 1/n$ after $O(n \log n)$ iterations w.h.p. The probability to sample the optimum is now $(1 - 1/n)^n \geq 1/(2e) = \Omega(1)$. Hence, waiting an additional $O(\log n)$ iterations, the optimum is sampled w.h.p. This concludes the first part of this proof.

Expected run time. Since we showed above that the sig-cGA optimizes LO in $O(n \log n)$ iterations w.h.p., we

are left to bound its run time in the event that at least one frequency decreases within the first $O(n \log n)$ iterations. As we discussed at the beginning of the first part of this proof, this only happens with a probability of $O(n^{-\varepsilon'})$, for $\varepsilon' > 2$.

Consider an interval of length t' . By Corollary 3, during the first t iterations, no frequency decreases for t' iterations with a probability of at least $1 - t'n^{1-\varepsilon'/3} \log_2 t$. Assume $t \leq n^{2n}$ and $t' = \Theta(n^2 \log n)$. Then no frequency decreases for t' iterations w.h.p., since $\varepsilon > 12$.

By using the result calculated in the first part, we see that a leftmost frequency τ_i at $1/n$ is increased during $O((1/\tau_i) \log n) = O(n \log n)$ iterations w.h.p. Thus, in overall, the sig-cGA finds the optimum during an interval of length $t' = \Theta(n^2 \log n)$ w.h.p., as n frequencies need to be increased to $1 - 1/n$. We pessimistically assume that the optimum is only found with a probability of at least $1/2$ during t' iterations. Hence, the expected run time in this case is $2t' = \Theta(t')$.

Last, we assume that we did not find the optimum during n^{2n} iterations, which only happens with a probability of at most $2^{-n^{2n}/t'}$. Then, the expected run time is at most n^n by pessimistically assuming that all frequencies are at $1/n$.

We conclude the proof by combining all of the three different regimes we just discussed, we see that we can bound the expected run time by

$$O(n \log n) + O(n^{-\varepsilon'}) \cdot O(t') + 2^{-n^{2n}/t'} \cdot n^n = O(n \log n). \quad \square$$

The proof of Theorem 5 shows that the sig-cGA rapidly makes progress when optimizing LO. In fact, after $O(i \log n)$ iterations, with $i \in [n]$, the sig-cGA finds a solution with fitness i w.h.p. (if i is large) and in expectation. Thus, in the fixed-budget perspective introduced by Jansen and Zarges [32], the sig-cGA performs very well on LO. For comparison, for the $(1 + 1)$ EA, it is known that the time to reach a fitness of i is $\Theta(in)$ in expectation and (again, when i is sufficiently large) w.h.p., see [33].

The reason that the sig-cGA optimizes LO so quickly is that the probability of saving a 1 at position i is increased by a constant factor once all frequencies at positions less than i are at $1 - 1/n$. This boost is a result of position i being the most relevant position for selection, assuming that all bits at positions less than i are 1.

Binary Value: A very similar boost in relevance occurs when considering the function BINVAL (BV), which returns the bit value of a bit string. Formally, BV is defined as

$$\text{BV}(x) = \sum_{i=1}^n 2^{n-i} x_i.$$

Note that the most significant bit is the leftmost.

BV imposes a lexicographic order from left to right on a bit string x , since a bit x_i has a greater weight than the sum of all weights at positions greater than i . This is similar to LO. The main difference is that, for BV, a position i can also be relevant for selection when bits at positions less than i are 0. More formally, for LO, position i is only relevant for selection when all of the bits at positions less than i are 1, whereas position i is relevant for selection for BV when all the bits at positions less than i are *the same*. With this insight, we adapt the proof of Theorem 5 for BV.

Corollary 6. Consider the sig-cGA (Alg. 1) with $\varepsilon > 12$ being a constant. Its run time on BV is $O(n \log n)$ w.h.p. and in expectation.

B. Analysis of OM

In order to analyze how likely it is that two individuals sampled from the sig-cGA have the same OM value, we use the following estimate, whose proof can be found, e.g., in [34].

Lemma 7. For $c \in \Theta(1)$, $\ell \in \mathbb{N}^+$, let $k \in [\ell/2 \pm c\sqrt{\ell}]$ and let $X \sim \text{Bin}(1/2, \ell)$. Then $\Pr[X = k] = \Omega\left(\frac{1}{\sqrt{\ell}}\right)$.

For the proof of the run time of the sig-cGA on OM, we show that, during *each* of the first $O(n \log n)$ iterations, each position can become relevant for selection with a decent probability of $\Omega(1/\sqrt{n})$. In contrast to LO, there is no sudden change in the probability that 1s are saved. Thus, it takes $O(n \log n)$ iterations to set a frequency to $1 - 1/n$. However, this is done for all frequencies in parallel. Thus, the overall run time remains $O(n \log n)$.

Theorem 8. Consider the sig-cGA (Alg. 1) with $\varepsilon > 12$ being a constant. Its run time on OM is $O(n \log n)$ w.h.p. and in expectation.

Proof. We first show that the run time holds w.h.p. Then we prove the expected run time.

Run time w.h.p. We consider the first $O(n \log n)$ iterations and condition on the event that no frequency decreases during that time. This can be argued in the same way as at the beginning in the proof of Theorem 5.

We now show that a single frequency (starting at $1/2$) is increased to $1 - 1/n$ within the first $O(n \log n)$ iterations w.h.p. as long as the other frequencies are at $1/2$ or at $1 - 1/n$. Hence, *all* frequencies are increased during that time w.h.p. when applying a union bound.

Similar to the proof of Theorem 5, when proving the expected run time, we use that, if all frequencies start at $1/n$, they are set to $1 - 1/n$ w.h.p. within $O(n^2 \log n)$ iterations in parallel. Thus, we combine both cases in the following argumentation.

Let $s \in \{1/2, 1/n\}$ denote the starting value of a frequency that we consider, and let $\ell \in [n]$ denote the number of frequencies *not* at $1 - 1/n$ during an arbitrary single iteration. Further, let $i \in [n]$ be a position in that iteration such that $\tau_i = s$. We prove that H_i saves 1s more likely by a factor of $1 + \Theta(1/\sqrt{\ell})$ when compared to the hypothesis. This results in τ_i being increased to $1 - 1/n$ within $O((\ell/s) \log n)$ iterations.

We determine the bias in saving a 1 by making the following observation: ignoring position i , if the absolute difference in the number of 1s of both offspring is greater than one, then bit i is not relevant for determining which offspring is selected. However, if the difference in the number of 1s (except position i) of both offspring is at most 1, having a 1 at position i makes it more likely for an individual to be selected. We now formalize this idea. To this end, let O denote the event that H_i saves a 1, and let A denote the event that the difference

of both offspring (except position i) is greater than 1. Note that in the case of A , the probability to save a 1 is τ_i .

We now consider the case of \bar{A} , that is, the absolute difference in the number of 1s of both offspring (excluding position i) is at most one. If it is zero, then H_i saves a 1 if none of the offspring has a 0 at position i . Thus, the respective probability is $1 - (1 - \tau_i)^2 = 2\tau_i - \tau_i^2$. In the case of the numbers of 1 differing by exactly one, a 1 is saved if the winner (with respect to all bits but bit i) has a 1 at position i (which it has with a probability of τ_i), or if the winner has a 0 at position i , the loser has a 1, and the loser wins the tie-breaking. The probability of this event is $(1/2)\tau_i(1 - \tau_i) \geq (1/4)\tau_i$. All in all, the probability to save a 1 conditional on \bar{A} is at least $\tau_i + (1/4)\tau_i = (5/4)\tau_i$, since $(5/4)\tau_i \leq 2\tau_i - \tau_i^2$ for $\tau_i \in \{1/2, 1/n\}$.

Taking both cases together, we bound

$$\Pr[O] \geq \Pr[A] \cdot \tau_i + \Pr[\bar{A}] \cdot \frac{5}{4}\tau_i.$$

By Lemma 4, we lower bound this term even further by calculating a lower bound for $\Pr[\bar{A}]$. We first show that the frequencies at $1 - 1/n$ and $1/n$ sample the same bits in both offspring with at least a constant probability. For the $n - \ell$ positions with frequencies at $1 - 1/n$, both offspring have a 1 at the respective positions with a probability of $(1 - 1/n)^{2(n-\ell)} \geq e^{-2}$, since $n - \ell \leq n - 1$. Analogously, for all positions with frequencies at $1/n$ (but τ_i), both offspring have a 0 at position i also with a probability of at least $(1 - 1/n)^{2(n-1)} \geq e^{-2}$. Hence, both offspring have the same bits at all positions with frequencies not at $1/2$ with a probability of at least e^{-4} .

We now consider the number of 1s of an offspring at the remaining $\ell' \leq \ell - 1$ (for $\ell \geq 2$) positions (except i) with frequencies at $1/2$. We call this number Y . Note that the expected value of Y is $\ell'/2$. By Lemma 7, for a $k \in [\ell'/2 \pm \sqrt{\ell'/2}]$, the probability that $Y = k$ is $\Omega(1/\sqrt{\ell'})$. Thus, the probability that both offspring have the same number of 1s at the ℓ' positions we consider is $d/\sqrt{\ell'}$, for a constant $d > 0$, since there are $\sqrt{\ell'}$ possible values of k and the probability that both offspring have k bits as 1 is $\Omega(1/\ell')$. Factoring in the probability of all remaining $n - \ell'$ positions to sample the same values in both offspring and for a sufficiently small constant $d' > 0$, we bound

$$\begin{aligned} \Pr[O] &\geq \left(1 - e^{-4} \frac{d}{\sqrt{\ell'}}\right) \cdot \tau_i + e^{-4} \frac{d}{\sqrt{\ell'}} \cdot \frac{5}{4}\tau_i \\ &\geq \left(1 + \frac{d'}{\sqrt{\ell'}}\right) \tau_i. \end{aligned}$$

Recall that we assumed $\ell \geq 2$ for this bound. For $\ell = 1$, i.e., $\ell' = 0$, we have $n - 1$ positions with frequencies at $1 - 1/n$ or $1/n$. Thus, $\Pr[\bar{A}] \geq e^{-4}$, as we discussed before. Consequently, we bound $\Pr[O] \geq (1 - e^{-2}) \cdot \tau_i + e^{-2} \cdot (5/4)\tau_i \geq (1 + d'/\sqrt{\ell})\tau_i$ if we choose d' sufficiently small. Overall, we use $(1 + d'/\sqrt{\ell})\tau_i$ as a lower bound for $\Pr[O]$.

Analogous to the proof of Theorem 5, we now consider the probability to detect a significance of 1s in H_i within k iterations. To this end, let $X \sim \text{Bin}(k, (1 + d'/\sqrt{\ell})\tau_i)$ and note that X is stochastically dominated by the process of

saving 1s at position i . We bound the probability to not detect a significance of 1s as follows:

$$\Pr[X < k\tau_i + s(\varepsilon, k\tau_i)] \\ \leq \Pr\left[X \leq E[X] - \left(\frac{kd'}{\sqrt{\ell}}\tau_i - s(\varepsilon, k\tau_i)\right)\right].$$

Let $k \geq 4(\varepsilon^2/d'^2)(\ell/\tau_i) \ln n$. Then $(kd'/\sqrt{\ell})\tau_i - s(\varepsilon, k\tau_i) \geq (kd'/(2\sqrt{\ell}))\tau_i =: \lambda$. Further note that $\text{Var}[X] = k\tau_i(1-\tau_i) \geq \lambda$ if d' is sufficiently small, which implies $\lambda^2/\text{Var}[X] \leq \lambda$. By Theorem 1 with $\text{Var}[X] \leq k\tau_i$, we see that

$$\Pr[X < k\tau_i + s(\varepsilon, k\tau_i)] \leq \Pr\left[X \leq E[X] - \frac{kd'}{2\sqrt{\ell}}\tau_i\right] \\ \leq e^{-\frac{1}{3} \cdot \frac{4k^2d'^2\tau_i^2}{4\ell k\tau_i}} = e^{-\frac{1}{3} \cdot \frac{kd'^2}{\ell}\tau_i} \leq e^{-\frac{4}{3}\varepsilon^2 \ln n} = n^{-\frac{4}{3}\varepsilon^2}.$$

Hence, τ_i is increased to $1-1/n$ within $4(\varepsilon^2/d'^2)(\ell/\tau_i) \ln n = O((\ell/\tau_i) \log n)$ iterations with a probability of at least $1 - n^{-4\varepsilon^2/3}$. By applying a union bound over all n frequencies, each frequency reaches $1-1/n$ within $O((\ell/\tau_i) \log n)$ iterations with a probability of at least $1 - n^{1-4\varepsilon^2/3} \geq 1 - n^{-191}$, as $\varepsilon > 12$, which is w.h.p.

Since we assume that no frequency drops within the first $O(n \log n)$ iterations w.h.p. and since all frequencies start at $1/2$, all of them reach $1-1/n$ within that time w.h.p. Then, the optimum is sampled during a single iteration with a probability of at least $(1-1/n)^n \geq 1/(2e) = \Omega(1)$. Thus, the optimum is sampled after $O(\log n)$ additional iterations w.h.p.

Expected run time. This part follows the same arguments as outlined in the respective part in the proof of Theorem 5. Different from there, assuming that a frequency is at $1/n$, it now takes $O(n^2 \log n)$ iterations to be increased to $1-1/n$ w.h.p., as we proved above. However, since $\varepsilon > 12$, a union bound over all n frequency again results in all frequencies being increased during $O(n^2 \log n)$ iterations w.h.p. The rest remains the same, which concludes this proof. \square

OM vs. LO. While the sig-cGA has the same asymptotic run time on LO and OM (w.h.p. and in expectation), the reasons differ. For LO, the frequencies are increased consecutively to $1-1/n$, where each frequency only needs $O(\log n)$ iterations, which is also the asymptotic minimum number of iterations to do so. This speed results from the sudden boost in probability once a position i becomes relevant, that is, its preceding frequencies are all at $1-1/n$ and thus sample only 1s with at least a constant probability. Given that both offspring have only 1s at positions in $[i-1]$, it suffices that position i has at least one 1, which is quite likely. In contrast, the probability that the bias in selection is also detected at positions after i declines exponentially in the distance to i , making the bias negligible. This fact is also exploited by Friedrich et al. [7] in the analysis (and design) of the scGA, which is why it has the same run time on LO.

For OM, the impact of the bias in selection depends on the number ℓ of other frequencies that are not at $1-1/n$. In order for a position i to detect the bias, the number of 1s in these positions has to almost be identical in both offspring, i.e., it can differ by at most one. This then adds a bias of

roughly $1/\sqrt{\ell}$ for saving a 1 in H_i . Since ℓ is large for a long time (for example, $\ell = n$ at the beginning), this bias remains small during that period. However, this bias is there constantly for each position. Thus, all frequencies can be optimized in parallel, whereas for LO this is done sequentially.

OM vs. BV. BV is often considered one extremal case of the class of linear functions, as its weights impose a lexicographic order on the bit positions. The other extreme is OM, where all weights are identical and basically no order among the positions exists. Our results show that the sig-cGA optimizes both functions in $O(n \log n)$. It remains an open question whether the sig-cGA is capable of optimizing any linear function in that time, a feat that the $(1+1)$ EA, a classical EA, is known to be capable of [9]. Contrary to that, it was proven for the cGA (an EDA) that it performs worse on BV than on OM [25]. Thus, a uniform performance on the class of linear functions would be a great feat for an EDA.

We would like to note that the result of Droste [25] considered the cGA without frequency borders, that is, the frequencies could reach values of 0. Once this is the case, the algorithm is stuck (as it only samples 0 at this position) and the optimization fails. It is unknown up to date whether the cGA still performs worse on BV when the frequencies are bound to the interval $[1/n, 1-1/n]$. However, the main idea of Droste's proof that frequencies drop very low remains. Thus, if sufficiently many frequencies were to drop to $1/n$, the cGA would still perform badly on BV. Note that this is exactly the problem that the sig-cGA circumvents with its update rule, resulting in its run time of $O(n \log n)$.

The only other known EDA run time result for BV was recently proven by Lehre and Nguyen [20]. They show that the PBIL optimizes BV with $O(n^2)$ fitness function evaluations in expectation (considering best parameter choices).

V. RUN TIME ANALYSIS FOR THE SCGA

Another variant of the cGA [30] that is able to optimize LO in $O(n \log n)$ w.h.p. is the *stable compact genetic algorithm* (scGA; found in the supplementary material) introduced by Friedrich et al. [7] with the intent to provide an EDA that optimizes LO in $o(n^2)$. The update procedure of the scGA is very similar to that of the cGA, that is, a frequency at position i is changed by a value of ρ with respect to the difference of the bits at i of the winner and the loser. However, an update toward $1/2$ is stronger by an additive term of a , where $a \in O(\rho)$ is an additional parameter.

Different from many other EDAs, the scGA does not have a margin and explicitly makes use of the frequency values 0 and 1. In fact, the scGA has another parameter $d \in (1/2, 1)$, which indicates a value that is sufficient in order to set a frequency to 1. The value $1-d$ is used symmetrically in order to set a frequency to 0. Thus, the scGA fixes frequencies once they leave the interval $(1-d, d)$.

Theorem 9. *Let $\alpha \in (0, 1]$ be a constant. Consider the scGA with $\rho = O(1/\log n)$, $a = \alpha\rho$, and $1/2 < d \leq 5/6$ with $d = \Theta(1)$. Its run time on OM is $\Omega(\min\{2^{\Theta(n)}, 2^{c/\rho}\})$ in expectation and w.h.p. for a constant $c > 0$.*

VI. RUN TIME ANALYSIS FOR THE CONVEX SEARCH ALGORITHM

The following *convex search algorithm* was proposed by Moraglio [18]. Its sole parameter is a population size $\mu \in \mathbb{N}$. The algorithm starts with a first population of μ random individuals $x^{(1,1)}, \dots, x^{(1,\mu)} \in \{0,1\}^n$. In each iteration $t = 1, 2, \dots$, the algorithm generates from the current “parent” population $x^{(t,1)}, \dots, x^{(t,\mu)}$ a new “offspring” population $x^{(t+1,1)}, \dots, x^{(t+1,\mu)}$ as follows.

- If the parent population contains only copies of a single individual, the algorithm stops and outputs this solution.
- If all individuals of the parent population have the same fitness, the offspring population is the parent population.
- Otherwise, the individuals with lowest fitness value are removed from the parent population (giving the “reduced parent population”) and the offspring population is obtained by μ times independently sampling from the convex hull of the reduced parent population. In other words, for all $i \in [n]$ and $j \in [\mu]$ independently, $x_i^{(t+1,j)}$ is chosen randomly from $\{0,1\}$ if the reduced parent population contains both an individual having a 0 at the i -th position and an individual having a 1 at this position. If all individuals of the reduced parent population have the same value $b \in \{0,1\}$ in the i -th position, then $x_i^{(t+1,j)} := b$.

The convex search algorithm with $\mu \geq 8 \log_2(4n^2 + n)$ and a suitable restart strategy was shown to optimize the LO problem in expected time $O(n \log n)$ [18]. We now show that its performance on the OM problem is not very attractive, namely it is asymptotically larger than any polynomial even when employing a suitable restart strategy. We suspect that much stronger lower bounds hold, but given the only moderate general interest in this algorithm so far, we restrict ourselves to this super-polynomial lower bound.

Theorem 10. *Let $c > 0$. Regardless of the population size, a run of the convex search algorithm on OM with probability at least $1 - O(n^{-c})$*

- *either reaches a state from which the optimum cannot be found,*
- *or within n^c iterations does not fix any bit-position.*

Consequently, at least $\Omega(n^c)$ iterations are necessary to find the optimum.

VII. CONCLUSIONS

We introduced the novel EDA sig-cGA, which optimizes both OM and LO in $O(n \log n)$ w.h.p. and in expectation. This is the first result of this kind for an EDA or even an EA. These run times are a result of the update process of the sig-cGA: it only updates its probabilistic model if it finds a significance in the history of its samples. In contrast, common EDAs or EAs that are analyzed theoretically do not store the entire history of their samples; EAs keep some samples as their population, and EDAs learn from samples iteratively and store the gained information implicitly in their model.

Since storing the entire history of samples demands a lot of memory if the sig-cGA runs longer, we proposed a

method that stores the history compactly while maintaining its important information. We want to note that this method can be improved even further. Currently, the sig-cGA saves new data in each iteration, even if no information is gained. In order to further reduce the memory demands of the sig-cGA, it should save a bit only if it is different from that of the competing offspring, that is, if there actually was a bias in both offspring at that position. Note that this is more similar to how the cGA updates its frequencies. However, if a frequency of the sig-cGA is at $1/2$, the number of samples that can contain important information (that is the pairs $(0,1)$ and $(1,0)$) is, in expectation, only half the number of all samples. Thus, the memory is only reduced by roughly a factor of 2.

All in all, the approach of the sig-cGA to reduce run times for a slight increase in memory appears to pay off very well. In this first work, as often in the theory of evolutionary algorithms, we only regarded the two unimodal benchmark functions OM and LO. Since it has been observed, e.g., recently in [35], that insights derived from such analyses can lead to wrong conclusions for more difficult functions, an interesting next step would be to analyze the performance of the sig-cGA on objective functions that have true local optima or that have larger plateaus of equal fitness. Two benchmark functions have been suggested in this context, namely jump functions [9] having an easy to reach local optimum with a scalable basin of attraction and plateau functions [36] having a plateau of scalable diameter around the optimum. We are vaguely optimistic that our sig-cGA has a good performance on these as well. We expect that the sig-cGA, as when optimizing OM, quickly fixes a large number of bits to the correct value and then, different from classic EAs, profits from the fact that the missing bits are sampled with uniform distribution, leading to a much more efficient exploration of the small subhypercube formed by these undecided bits. Needless to say, transforming this speculation into a formal proof would be a significant step forward to understanding the sig-cGA.

From a broader perspective, our work shows that by taking into account a longer history and only updating the model when the history justifies it, the performance of a classic EDA can be improved and its usability can be increased (since the difficult choice of the model update strength is now obsolete). An interesting question from this viewpoint would be to what extent similar ideas can be applied to other well-known EDAs.

From a very broad perspective, our work suggests that generally EC could profit from enriching the iterative evolutionary process with mechanisms that collect and exploit information over several iterations. So far, such learning-based concepts are rarely used in EC. The only theoretical works in this direction propose a history-based choice of the mutation strength [37] and analyze hyperheuristics that stick to a chosen subheuristic until its performance over the last τ iterations, τ a parameter of the algorithms, appears insufficient (see, e.g., [38] and the references therein).

REFERENCES

- [1] B. Doerr and M. S. Krejca, “Significance-based estimation-of-distribution algorithms,” in *Proc. of GECCO’18*. ACM, 2018, pp. 1483–1490.

- [2] M. Pelikan, M. Hauschild, and F. G. Lobo, "Estimation of distribution algorithms," in *Springer Handbook of Computational Intelligence*, 2015, pp. 899–928.
- [3] M. S. Krejca and C. Witt, "Theory of estimation-of-distribution algorithms," in [39], 2020, pp. 405–442, also available at <http://arxiv.org/abs/1806.05392>.
- [4] D. Sudholt and C. Witt, "On the choice of the update strength in estimation-of-distribution algorithms and ant colony optimization," *Algorithmica*, vol. 81, no. 4, pp. 1450–1489, 2019.
- [5] M. S. Krejca and C. Witt, "Lower bounds on the run time of the univariate marginal distribution algorithm on onemax," in *Proc. of FOGA'17*. ACM, 2017, pp. 65–79.
- [6] J. Lengler, D. Sudholt, and C. Witt, "Medium step sizes are harmful for the compact genetic algorithm," in *Proc. of GECCO'18*. ACM, 2018, pp. 1499–1506.
- [7] T. Friedrich, T. Kötzing, and M. S. Krejca, "EDAs cannot be balanced and stable," in *Proc. of GECCO'16*. ACM, 2016, pp. 1139–1146.
- [8] W. Zheng, G. Yang, and B. Doerr, "Working principles of binary differential evolution," in *Proc. of GECCO'18*. ACM, 2018, pp. 1103–1110.
- [9] S. Droste, T. Jansen, and I. Wegener, "On the analysis of the $(1+1)$ evolutionary algorithm," *Theoretical Computer Science*, vol. 276, no. 1–2, pp. 51–81, 2002.
- [10] C. Witt, "Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions," *Evolutionary Computation*, vol. 14, pp. 65–86, 2006.
- [11] T. Jansen, K. A. De Jong, and I. Wegener, "On the choice of the offspring population size in evolutionary algorithms," *Evolutionary Computation*, vol. 13, pp. 413–440, 2005.
- [12] B. Doerr and M. Künnemann, "Optimizing linear functions with the $(1+\lambda)$ evolutionary algorithm – different asymptotic runtimes for different instances," *Theoretical Computer Science*, vol. 561, pp. 3–23, 2015.
- [13] G. Badkobeh, P. K. Lehre, and D. Sudholt, "Unbiased black-box complexity of parallel search," in *Proc. of PPSN'14*. Springer, 2014, pp. 892–901.
- [14] B. Doerr, C. Gießen, C. Witt, and J. Yang, "The $(1 + \lambda)$ evolutionary algorithm with self-adjusting mutation rate," *Algorithmica*, vol. 81, no. 2, pp. 593–631, 2019.
- [15] B. Doerr, C. Witt, and J. Yang, "Runtime analysis for self-adaptive mutation rates," in *Proc. of GECCO'18*. ACM, 2018, pp. 1475–1482.
- [16] D. Antipov, B. Doerr, J. Fang, and T. Hetet, "A tight runtime analysis for the $(\mu + \lambda)$ EA," in *Proc. of GECCO'18*. ACM, 2018, pp. 1459–1466.
- [17] B. Doerr and C. Doerr, "Optimal static and self-adjusting parameter choices for the $(1+(\lambda, \lambda))$ genetic algorithm," *Algorithmica*, vol. 80, no. 5, pp. 1658–1709, 2018.
- [18] A. Moraglio and D. Sudholt, "Principled design and runtime analysis of abstract convex evolutionary search," *Evolutionary Computation*, vol. 25, no. 2, pp. 205–236, 2017.
- [19] D. Dang, P. K. Lehre, and P. T. H. Nguyen, "Level-based analysis of the univariate marginal distribution algorithm," *Algorithmica*, vol. 81, no. 2, pp. 668–702, 2019.
- [20] P. K. Lehre and P. T. H. Nguyen, "Level-based analysis of the population-based incremental learning algorithm," in *Proc. of PPSN'18*. Springer, 2018, pp. 105–116.
- [21] C. Witt, "Upper bounds on the running time of the univariate marginal distribution algorithm on onemax," *Algorithmica*, vol. 81, no. 2, pp. 632–667, 2019.
- [22] F. Neumann and C. Witt, "Runtime analysis of a simple ant colony optimization algorithm," *Algorithmica*, vol. 54, no. 2, pp. 243–255, 2009.
- [23] B. Doerr, F. Neumann, D. Sudholt, and C. Witt, "Runtime analysis of the 1-ANT ant colony optimizer," *Theoretical Computer Science*, vol. 412, no. 17, pp. 1629–1644, 2011.
- [24] F. Neumann, D. Sudholt, and C. Witt, "Analysis of different MMAS ACO algorithms on unimodal functions and plateaus," *Swarm Intelligence*, vol. 3, no. 1, pp. 35–68, 2009.
- [25] S. Droste, "A rigorous analysis of the compact genetic algorithm for linear functions," *Natural Computing*, vol. 5, no. 3, pp. 257–283, 2006.
- [26] S. Droste, T. Jansen, and I. Wegener, "Upper and lower bounds for randomized search heuristics in black-box optimization," *Theory of Computing Systems*, vol. 39, pp. 525–544, 2006.
- [27] G. Anil and R. P. Wiegand, "Black-box search by elimination of fitness functions," in *Proc. of FOGA'09*. ACM, 2009, pp. 67–78.
- [28] P. Afshani, M. Agrawal, B. Doerr, C. Doerr, K. G. Larsen, and K. Mehlhorn, "The query complexity of finding a hidden permutation," *Discrete Applied Mathematics*, pp. 28–50, 2019.
- [29] B. Doerr, "Probabilistic tools for the analysis of randomized optimization heuristics," in [39], 2020, pp. 1–87, also available at <https://arxiv.org/abs/1801.06733>.
- [30] G. R. Harik, F. G. Lobo, and D. E. Goldberg, "The compact genetic algorithm," *IEEE Transactions on Evolutionary Computation*, vol. 3, no. 4, pp. 287–297, 1999.
- [31] P. K. Lehre and C. Witt, "Black-box search by unbiased variation," *Algorithmica*, vol. 64, no. 4, pp. 623–642, 2012.
- [32] T. Jansen and C. Zarges, "Performance analysis of randomised search heuristics operating with a fixed budget," *Theoretical Computer Science*, vol. 545, pp. 39–58, 2014.
- [33] B. Doerr, T. Jansen, C. Witt, and C. Zarges, "A method to derive fixed budget results from expected optimisation times," in *Proc. of GECCO'13*. ACM, 2013, pp. 1581–1588.
- [34] B. Doerr and C. Winzen, "Ranking-based black-box complexity," *Algorithmica*, vol. 68, pp. 571–609, 2014.
- [35] B. Doerr, H. P. Le, R. Makhmara, and T. D. Nguyen, "Fast genetic algorithms," in *Proc. of GECCO'17*. ACM, 2017, pp. 777–784.
- [36] D. Antipov and B. Doerr, "Precise runtime analysis for plateaus," in *Proc. of PPSN'18*. Springer, 2018, pp. 117–128.
- [37] B. Doerr, C. Doerr, and J. Yang, " k -bit mutation with self-adjusting k outperforms standard bit mutation," in *Proc. of PPSN'16*. Springer, 2016, pp. 824–834.
- [38] B. Doerr, A. Lissovoi, P. S. Oliveto, and J. A. Warwicker, "On the runtime analysis of selection hyper-heuristics with adaptive learning periods," in *Proc. of GECCO'18*. ACM, 2018, pp. 1015–1022.
- [39] B. Doerr and F. Neumann, *Theory of Evolutionary Computation—Recent Developments in Discrete Optimization*. Springer, 2020.
- [40] P. S. Oliveto and C. Witt, "Simplified drift analysis for proving lower bounds in evolutionary computation," *Algorithmica*, vol. 59, no. 3, pp. 369–386, 2011.
- [41] —, "Erratum: simplified drift analysis for proving lower bounds in evolutionary computation," *CoRR*, vol. abs/1211.7184, 2012. [Online]. Available: <http://arxiv.org/abs/1211.7184>
- [42] B. Doerr, "Analyzing randomized search heuristics via stochastic domination," *Theoretical Computer Science*, vol. 773, pp. 115–137, 2019.
- [43] C. Gießen and C. Witt, "The interplay of population size and mutation probability in the $(1 + \lambda)$ EA on OneMax," *Algorithmica*, vol. 78, pp. 587–609, 2017.
- [44] B. Doerr and D. Johannsen, "Edge-based representation beats vertex-based representation in shortest path problems," in *Proc. of GECCO'10*. ACM, 2010, pp. 759–766.



Benjamin Doerr received the Diploma, Ph.D., and the Habilitation degrees in mathematics from the Christian-Albrechts-Universität zu Kiel, Kiel, Germany, in 1998, 2000, and 2005, respectively.

He is a Full Professor with the École Polytechnique, Palaiseau, France. From 2005 to 2013, he was a Senior Researcher (tenured) at the Max Planck Institute for Informatics, Saarbrücken, Germany, and an Adjunct Professor at Saarland University, Saarbrücken. He is a member of the editorial boards of several scientific journals, among them, *Evolutionary Computation*, *Natural Computing*, and the *Artificial Intelligence Journal*.

His research interests include mathematical analyses of problem-specific algorithms and randomized search heuristics as well as the development of new algorithms based on such analyses.



Martin S. Krejca received the B.S. and M.S. degrees in computer science from the University of Jena, Jena, Germany, in 2012 and 2014, respectively, and the Ph.D. degree in computer science from the Hasso Plattner Institute, University of Potsdam, Potsdam, Germany, in 2019. He is currently a Post-Doctoral Researcher with the Algorithm Engineering Group, Hasso Plattner Institute, University of Potsdam, Potsdam, Germany.

His research interests include the theoretical analysis of randomized processes, especially the analysis of estimation-of-distribution algorithms, and the development of efficient tools for this goal.

APPENDIX

Here we provide all of the material that we omitted in the main paper.

PRELIMINARIES

A. Detailed Description of the sig-cGA

Proof of Lemma 2. In order for τ_i to change, the number of 0s or 1s in X needs to deviate significantly from the hypothesis, which follows the same distribution as X by assumption. We are going to use Theorem 1 in order to show that, in such a scenario, X will deviate significantly from its expected value only with a probability of at most $n^{-\varepsilon/3} \log_2 k$ for any number of trials at most k .

Let $\tau'_i = \min\{\tau_i, 1 - \tau_i\}$. Note that, in order for τ_i to change, a significance of values sampled with probability τ'_i needs to be sampled. That is, for $\tau_i = 1/2$, either a significant amount of 1s or 0s needs to occur; for $\tau_i = 1 - 1/n$, a significant amount of 0s needs to occur; and, for $\tau_i = 1/n$, a significant amount of 1s needs to occur. Further, let X' denote the number of values we are looking for a significance within $k' \leq k$ trials. That is, if $\tau_i = 1/2$, X' is either the number of 1s or 0s; if $\tau_i = 1 - 1/n$, X' is the number of 0s; and if $\tau_i = 1/n$, X' is the number of 1s.

Given the definition of τ'_i , we see that $E[X'] = k'\tau'_i$ and $\text{Var}[X'] = k'\tau_i(1 - \tau_i) \leq k'\tau'_i$. Since we want to apply Theorem 1, let $\lambda = s(\varepsilon, E[X']) = s(\varepsilon, k'\tau'_i)$ and $\sigma^2 = \text{Var}[X']$.

First, consider the case that $\lambda = s(\varepsilon, k'\tau'_i) = \varepsilon \ln n$, i.e., that $(k'\tau'_i \ln n)^{1/2} \leq \ln n$, which is equivalent to $k' \leq (1/\tau'_i) \ln n$. Note that $\lambda^2/\sigma^2 \geq \varepsilon^2 \ln n \geq \ln n$, as $\varepsilon \geq 1$. Thus, $\min\{\lambda^2/\sigma^2, \lambda\} \geq \varepsilon \ln n$.

Now consider the case $\lambda = s(\varepsilon, k'\tau'_i) = \varepsilon(k'\tau'_i \ln n)^{1/2}$, i.e., that $(k'\tau'_i \ln n)^{1/2} \geq \ln n$, which is equivalent to $k' \geq (1/\tau'_i) \ln n$. We see that $\lambda \geq \varepsilon \ln n$ and $\lambda^2/\sigma^2 \geq \varepsilon^2 \ln n$. Hence, as before, we get $\min\{\lambda^2/\sigma^2, \lambda\} \geq \varepsilon \ln n$.

Combining both cases and applying Theorem 1, we get

$$\begin{aligned} \Pr[X' \geq k'\tau'_i + s(\varepsilon, k'\tau'_i)] &= \Pr[X' \geq E[X'] + \lambda] \\ &\leq e^{-\frac{1}{3} \min\{\frac{\lambda^2}{\sigma^2}, \lambda\}} \leq e^{-\frac{\varepsilon}{3} \ln n} = n^{-\frac{\varepsilon}{3}}. \end{aligned}$$

That is, the probability of detecting a (false) significance during k' trials is at most $n^{-\varepsilon/3}$. Since we look for a significance a total of at most $\log_2 k$ times during an iteration, we get by a union bound that the probability of detecting a significance within a history of length k is at most $n^{-\varepsilon/3} \log_2 k$. \square

Proof of Corollary 3. For any $i \in [n]$ during any of the k iterations, by Lemma 2, the probability that τ_i changes is at most $n^{-\varepsilon/3} \log_2 k$. Via a union bound over all k' relevant iterations and all n frequencies, the statement follows. \square

B. Efficient Implementation of the sig-cGA

We now elaborate on how to implement the histories of the sig-cGA efficiently. First, we provide an intuitive description, then we provide some pseudocode.

For our condensed storage of the history, we have a list of blocks, each storing the number of 1s in some discrete interval

$[t_1..t_2]$ of length equal to a power of two (including 1). When a new item has to be stored, we append a block of size 1 to the list. Then, traversing the list starting with the newest element, we check if there are three consecutive blocks of the same size, and if so, we merge the two oldest ones into a new block of twice the size. By this, we always maintain a list of blocks such that, for a certain power 2^k , there are between one and two blocks of length 2^j for all $j \in [0..k - 1]$. This structural property implies both that we only have a logarithmic number of blocks (as we have $k = O(\log |H_i|)$) and that we can (in amortized constant time) access all historic intervals consisting of full blocks, which in particular implies that we can access an interval with length in $[2^j, 2^{j+1} - 1]$ for all $j \in [0..k]$.

For the pseudocode of this approach, assume that a list element has a pointer to the next list element (next), its previous element (prev), and stores an integer value (load).

Algorithm 2: The algorithm used by the sig-cGA in order to condense a history H , given a new bit value x . Let L be a list with elements of non-decreasing load.

```

1  $X \leftarrow$  a new list element with load  $x$ ;
2 Append  $X$  to the head of  $L$ ;
3  $C \leftarrow$  head of  $L$ ;
4  $N \leftarrow C.\text{next}$ ;
5  $r \leftarrow 0$ ;
6 while  $N$  is not null do
7   if  $C.\text{load} = N.\text{load}$  then  $r \leftarrow r + 1$ ;
8   if  $r = 2$  then
9      $r \leftarrow 0$ ;
10     $X \leftarrow$  a new list element with load  $C.\text{load} +$ 
11       $N.\text{load}$ ;
12     $C.\text{prev}.\text{next} \leftarrow X$ ;
13     $X.\text{next} \leftarrow N.\text{next}$ ;
14     $N \leftarrow X$ ;
15   $C \leftarrow N$ ;
16   $N \leftarrow N.\text{next}$ ;

```

RUN TIME ANALYSIS FOR THE SIG-CGA ON BV

Proof of Corollary 6. We can use the same arguments as in the proof of Theorem 5. The only difference is in how we define the event A that position i is *not* relevant for selection. As described above, let A denote the event that there exists a position $j \in [i - 1]$ such that, for the two offspring x and y sampled during the iteration we consider, $x_i \neq y_i$. Further, let O denote again the event that we save a 1 this iteration. We see that the following equation also holds for BV:

$$\Pr[O] = \Pr[A] \cdot \tau_i + \Pr[\bar{A}] \cdot (1 - \tau_i^2).$$

Again, according to Lemma 4, we get a lower bound for $\Pr[O]$ if we get a lower-bound for $\Pr[\bar{A}]$. The event \bar{A} occurs if all the bits sampled at positions in $[i - 1]$ have the same value. Thus, a subset of this event is that all of the bits sampled at positions in $[i - 1]$ are 1, which yields that

$$\Pr[\bar{A}] \geq \left(1 - \frac{1}{n}\right)^{2(i-1)}.$$

This probability is minimized by choosing $i = n$. Hence, we get $\Pr[\bar{A}] \geq e^{-2}$, which is the same lower bound for \bar{A} as in the proof of Theorem 5.

Since this is the only part of the proof where the fitness function comes into play, the remaining proof is identical to the proof of Theorem 5. \square

RUN TIME ANALYSIS FOR THE SCGA

Algorithm 3: The scGA [7] with parameters ρ , a , and d optimizing f

```

1  $t \leftarrow 0$ ;
2 for  $i \in [n]$  do  $\tau_i^{(t)} \leftarrow \frac{1}{2}$ ;
3 repeat
4    $x, y \leftarrow$  offspring sampled with respect to  $\tau^{(t)}$ ;
5    $(x, y) \leftarrow$  winner/loser of  $x$  and  $y$  with respect to  $f$ ;
6   for  $i \in [n]$  do
7     if  $x_i > y_i$  then
8       if  $\tau_i^{(t)} \leq \frac{1}{2}$  then  $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)} + \rho + a$ ;
9       else if  $\frac{1}{2} < \tau_i^{(t)} < d$  then
10         $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)} + \rho$ ;
11        else  $\tau_i^{(t+1)} \leftarrow 1$ ;
12     else if  $x_i < y_i$  then
13       if  $\tau_i^{(t)} \geq \frac{1}{2}$  then  $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)} - \rho - a$ ;
14       else if  $1 - d < \tau_i^{(t)} < \frac{1}{2}$  then
15         $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)} - \rho$ ;
16        else  $\tau_i^{(t+1)} \leftarrow 0$ ;
17     else  $\tau_i^{(t+1)} \leftarrow \tau_i^{(t)}$ ;
18    $t \leftarrow t + 1$ ;
19 until termination criterion met;
```

Before we prove the theorem, we mention two other theorems that we are going to use in the proof. The first bounds the probability of a randomly sampled bit string having $s \in \{0\} \cup [n]$ 1s. We use it in order to bound the probability of both offspring having the same number of 1s. Note that the values $1/6$ and $5/6$ in the lemma are somewhat arbitrary and can be exchanged for any constant in $(0, 1/2)$ and $(1/2, 1)$, respectively.

Lemma 11 ([4]). *Let S denote the sum of n independent Poisson trials with probabilities τ_1, \dots, τ_n such that, for all $i \in [n]$, $1/6 \leq \tau_i \leq 5/6$. Then, for all $s \in \{0\} \cup [n]$,*

$$\Pr[S = s] = O\left(\frac{1}{\sqrt{n}}\right).$$

The next theorem provides an upper bound on the probability of a random process stopping after a certain time. We use it in order to show that it is unlikely for a frequency of the scGA when optimizing OM to get to 1 within a certain number of iterations.

Theorem 12 (Negative Drift; [40], [41]). *Let $(X_t)_{t \in \mathbb{N}}$ be real-valued random variables describing a stochastic process over some state space, with $X_0 \geq b$. Suppose there exist an interval*

$[a, b] \subseteq \mathbb{R}$, two constants $\delta, \varepsilon > 0$, and, possibly depending on $\ell := b - a$, a function $r(\ell)$ satisfying $1 \leq r(\ell) = o(\ell / \log(\ell))$ such that, for all $t \in \mathbb{N}$, the following two conditions hold:

1) $E[X_{t+1} - X_t \mid X_t \wedge a < X_t < b] \geq \varepsilon$ and,

2) for all $j \in \mathbb{N}$, $\Pr[|X_{t+1} - X_t| \geq j \mid X_t \wedge X_t > a] \leq \frac{r(\ell)}{(1+\delta)^j}$.

Then there is a constant $c > 0$ such that, for $T := \min\{t \in \mathbb{N} \mid X_t \leq a\}$, it holds that

$$\Pr\left[T \leq 2^{\frac{c\ell}{r(\ell)}}\right] = 2^{-\Omega\left(\frac{\ell}{r(\ell)}\right)}.$$

We can now prove our result.

Proof of Theorem 9. We only show that the run time is in $\Omega(\min\{2^{\Theta(n)}, 2^{c/\rho}\})$ w.h.p. The statement for the expected run time follows by lower-bounding the terms that occur with a probability of $o(1)$ with 0.

We first prove the bound of $\Omega(2^{c/\rho})$. We do so by showing that each frequency will stay in the non-empty interval $(1 - d, d) \subset [1/6, 5/6]$ w.h.p. Although OM introduces a bias into updating a frequency, it is too tiny in order to compensate the strong drift toward $1/2$ in the update.

We lower-bound the expected time it takes the scGA to optimize OM by upper-bounding the probability it takes a single frequency to leave the interval $(1 - d, d)$. Thus, we condition during the entire proof implicitly on the event that all frequencies are in the interval $(1 - d, d)$. Note that, in this scenario, the probability to sample the optimum during an iteration is at most $(5/6)^n$, which is exponentially small, even for a polynomial number of iterations.

Consider an index $i \in [n]$ with $\tau_i \in (1 - d, d)$. We only upper-bound the probability it takes τ_i to reach d . Note that the probability of τ_i reaching $1 - d$ is at most that large, as OM introduces a bias for 1s into the selection process. Hence, we could argue optimistically for τ_i reaching $1 - d$ as we do for τ_i reaching d by swapping 1s for 0s and considering $1 - \tau_i$ instead.

Let T denote the first point in time t such that $\tau_i^{(t)} \geq d$. We want to apply Theorem 12 and show that it is unlikely for τ_i to reach d within $2^{c/\rho}$ iterations. Hence, we define the following potential function $g: [0, 1] \rightarrow \mathbb{R}$:

$$g(\tau_i) = \frac{1}{\rho}(1 - \tau_i),$$

which we will use for our frequencies. Note that, at the beginning, τ_i is at $1/2$, i.e., $g(1/2) = 1/(2\rho)$. We stop once $\tau_i \geq d$, i.e., $g(\tau_i) \leq (1 - d)/\rho$. Thus, we consider an interval of length $\ell := 1/(2\rho) - (1 - d)/\rho = (2d - 1)/\rho = \Theta(1/\rho)$, as $d > 1/2$ is a constant.

We now argue how an update to τ_i is performed in order to estimate its expected value after an update, which is necessary in order to apply Theorem 12. Consider, similar to the proof of Theorem 8, that the bits of both offspring x and y for all positions but position i have been determined. If the difference of the number of 1s of both offspring is at least 2, i.e., $\|x - y\|_1 - x_i + y_i \geq 2$, then the outcome of neither x_i nor y_i can change the outcome of the selection process. Thus, τ_i increases with probability $\tau_i(1 - \tau_i)$, as the winner offspring needs to

sample a 1 and the loser a 0. Analogously, in this case, the probability that τ_i decreases is $\tau_i(1 - \tau_i)$, too.

If the difference of the number of 1s of both offspring is one, then, in order to increase τ_i , the winner (with respect to all bits but bit i) needs to sample a 1 and the loser a 0, or the winner needs to sample a 0, the loser a 1, and the loser wins. The first case has a probability of $\tau_i(1 - \tau_i)$, the second of $(1/2)\tau_i(1 - \tau_i)$, due to the uniform selection when the offspring have equal fitness. In order to decrease τ_i , the winner needs to sample a 0, the loser a 1, and the winner has to win, which has a probability of $(1/2)\tau_i(1 - \tau_i)$.

If the difference of the number of 1s of both offspring is zero, then τ_i is increased if any offspring samples a 1 and the other samples a 0. This has probability $2\tau_i(1 - \tau_i)$. In this case, it is not possible that τ_i is decreased.

In order to estimate the probabilities of when τ_i increases or decreases, we need to estimate the probabilities that the number of 1s of both offspring differ by at least two, differ by exactly one, and differ by exactly zero. Let p_1 denote the probability that this difference is one, and let p_0 denote the probability that the difference is zero. We now bound these probabilities.

Assume that offspring x has k 1s, where $k \in \{0\} \cup [n-1]$, since we assume that bit i has not been sampled yet. For p_0 , y needs to sample k 1s as well, and for p_1 , y needs to sample $k-1$ or $k+1$ 1s (such that the result is still in $\{0\} \cup [n-1]$). Due to Lemma 11, the probability for y to have this many 1s is $O(1/\sqrt{n})$, as we assume that all frequencies are in the interval $(1-d, d) \subset [1/6, 5/6]$. Hence, by the law of total probability, we get

$$p_0 = O\left(\frac{1}{\sqrt{n}}\right) \text{ and } p_1 = O\left(\frac{1}{\sqrt{n}}\right).$$

In the following, let $\gamma > 0$ be a constant such that $p_0 \leq \gamma/\sqrt{n}$ and $p_1 \leq \gamma/\sqrt{n}$.

We now consider the drift of $g(\tau_i)$ in any iteration t such that $1/2 < \tau_i^{(t)} < d$, i.e., we show that condition (1) of Theorem 12 holds. Let $\tau = \tau_i^{(t)}$ and $\tau' = \tau_i^{(t+1)}$. Note that conditioning on $g(\tau)$ is the same as conditioning on τ , as g is injective. If τ increases, it changes by ρ , and if it decreases, it changes by $\rho + a$.

$$\begin{aligned} & E\left[g(\tau') - g(\tau) \mid g(\tau) \wedge \frac{1-d}{\rho} < g(\tau) < \frac{1}{2\rho}\right] \\ &= \frac{1}{\rho} E\left[\tau - \tau' \mid \tau \wedge \frac{1}{2} < \tau < d\right] \\ &= \frac{1}{\rho} \left((\rho + a)((1 - p_0 - p_1)\tau(1 - \tau) + p_1 \cdot \frac{1}{2}\tau(1 - \tau)) \right. \\ &\quad \left. - \rho((1 - p_0 - p_1)\tau(1 - \tau) + p_1 \cdot \frac{3}{2}\tau(1 - \tau) \right. \\ &\quad \left. + p_0 \cdot 2\tau(1 - \tau)) \right) \\ &= \frac{1}{\rho} \tau(1 - \tau) \left((\rho + a)(1 - p_0 - \frac{1}{2}p_1) - \rho(1 + p_0 + \frac{1}{2}p_1) \right) \\ &= \frac{1}{\rho} \tau(1 - \tau) \left(\rho(-2p_0 - p_1) + a(1 - p_0 - \frac{1}{2}p_1) \right). \end{aligned}$$

For the negative terms with factor a , by using that $a \leq \rho$ and by applying the bounds on p_0 and p_1 , we get

$$\begin{aligned} & E\left[g(\tau') - g(\tau) \mid g(\tau) \wedge \frac{1-d}{\rho} < g(\tau) < \frac{1}{2\rho}\right] \\ &\geq \frac{1}{\rho} \tau(1 - \tau) \left(a - 3p_0\rho - \frac{3}{2}p_1\rho \right) \\ &\geq \frac{1}{\rho} \tau(1 - \tau) \left(a - 5\frac{\gamma}{\sqrt{n}}\rho \right). \end{aligned}$$

Due to $a = \alpha\rho$, there is a sufficiently small constant $\beta > 0$ such that $a - 5\gamma\rho/\sqrt{n} \geq \beta\rho$. Thus, we get

$$\begin{aligned} E\left[g(\tau') - g(\tau) \mid g(\tau) \wedge \frac{1-d}{\rho} < g(\tau) < \frac{1}{2\rho}\right] &\geq \beta\tau(1 - \tau) \\ &\geq \beta\frac{1}{6} \cdot \frac{5}{6}, \end{aligned}$$

which is constant.

We now show that condition (2) of Theorem 12 holds. For this, we define $r(\ell) = 2$ and $\delta = \sqrt{2} - 1 > 0$. Note that $1 \leq r(\ell) = o(\ell/\log(\ell)) = o(1/(\rho\log(1/\rho)))$ holds, as $\rho = o(1)$. Since τ can change by at most $\rho + a \leq 2\rho$ during a single update, $g(\tau)$ can change by at most 2. Thus, we only need to bound $\Pr[|g(\tau') - g(\tau)| \geq j \mid g(\tau) \wedge g(\tau) > (1-d)\rho]$ for $j \in \{0, 1, 2\}$. For all of these three cases, $r(\ell)/((1+\delta)^j) \geq 1$. Thus, condition (2) trivially holds for all $j \in \mathbb{N}$.

Overall, by applying Theorem 12 and recalling that $\ell = \Theta(1/\rho)$, there are constant $c, c', c'' > 0$ such that

$$\Pr\left[T \leq 2^{\frac{c'\ell}{r(\ell)}}\right] = \Pr\left[T \leq 2^{\frac{c}{\rho}}\right] \leq 2^{-\Omega\left(\frac{1}{\rho}\right)} \leq n^{-c''}.$$

Thus, w.h.p., τ_i does not reach b within $2^{c/\rho}$ iterations, given that all frequencies are in $(1-d, d)$. As discussed before, the probability of τ_i reaching $1-d$ has at most the same probability. Note that conditioning on never sampling the optimum during any of these t iterations increases these probabilities only by a factor of $1 - t(5/6)^n$, which is constant if $t = o(2^{\Theta(n)})$. Otherwise, we choose $2^{\Theta(n)}$ as run time bound. This concludes the proof. \square

RUN TIME ANALYSIS FOR THE CONVEX SEARCH ALGORITHM

While the result may seem natural, proving it is made difficult by the dependencies inflicted from restricting the parent population to all but the lower fitness level. We remove these dependencies by suitable pessimistic estimates (e.g., estimating that a position does not become fixed to 1 when at least Δ zeros are sampled), by suitable domination arguments (cf. [42]), and by first regarding an artificial process in which bits can only be fixed to 1.

Proof of Theorem 10. Since we are aiming at an asymptotic statement, we assume in the following that n is sufficiently large.

To ease the following proof, we first argue that only the case $\mu = \Theta(\log n)$ is interesting. If $\mu \leq \frac{1}{2} \log_2 n$, then with probability

$$1 - (1 - 2^{-\mu})^n \geq 1 - \exp(-2^{-\mu}n) \geq 1 - \exp(-n^{1/2}),$$

at least one of the bit-positions of the initial population is already converged to zero. Let now $\mu \geq K \log_2 n$ for a sufficiently large constant K (which may depend on the constant c). We show that a random population with probability n^{-2c} fixes no bit (that is, the next population is again fully random). By elementary properties of the binomial distribution, with probability at least $1 - (2/3)^\mu$, the lowest fitness value of the

random population is below $n/2$. Since the probability of having a fitness of at least $n/2$ is at least $\frac{1}{2}$, the additive Chernoff bound ([29, Theorem 10.7]) gives that with probability at least $1 - \exp(-\mu/8)$, the number μ^+ of individuals having a fitness of at least $n/2$ is at least $\mu/4$.

We now condition on $\mu^+ \geq \mu/4$ and that there is an individual with fitness less than $n/2$ (and recall that this event happens with probability at least $1 - (2/3)^\mu - \exp(-\mu/8)$). We first note that in this case the μ^+ individuals with fitness $n/2$ or more surely belong to the reduced population, which defines the next population. For each of these μ^+ individuals and for each of their bit-positions, the probability to have a one is between $1/2$ and $3/4$, since these individuals are random individuals conditional on having at least $n/2$ ones. Since these individuals are stochastically independent, the probability that a bit-position in all these μ^+ individuals has the value 0 is at most $(1/2)^{\mu^+} \leq (1/2)^{\mu/4}$ and the probability is at most $(3/4)^{\mu^+} \leq (3/4)^{\mu/4}$ for the event that they are all one.

In summary, we obtain that the probability that a bit-position becomes fixed is at most $(2/3)^\mu + \exp(-\mu/8) + (1 - (2/3)^\mu - \exp(-\mu/8))n((1/2)^{\mu/4} + (3/4)^{\mu/4})$. By taking the constant K in the lower bound for μ sufficiently large, this probability is at most n^{-2c} . Hence a union bound over n^c iterations shows that within this time frame, with probability $1 - n^{-c}$ no bit-position becomes fixed.

In the remainder we thus assume that $\mu = \Theta(\log n)$ with implicit constants depending on the constant c only.

We first regard the artificial random process which equals a true run of the CSA on OM except that for all bits $i \in [n]$ where the reduced parent population contains only individuals with bit-value 0 we still sample the offspring bits randomly from $\{0, 1\}$. In other words, we prevent the algorithm from letting a bit-value converge to the wrong value of 0.

Let $\Delta = \lceil 2 + 2c + 2\frac{\mu}{\log_2 n} \rceil$ and $t = \lceil \frac{2^\mu}{4(2\mu)^\Delta} \rceil$. We call a bit-position $i \in [n]$ at time $s \in [t]$ *unsafe* if the population at time s contains at least $\mu - \Delta$ ones in this bit-position, that is, if $\sum_{j=1}^{\mu} x_i^{(s,j)} \geq \mu - \Delta$, and if this bit-position was determined by sampling random bit-values (that is, not by setting all bit-values to one because the previous reduced population was converged to 1 in this position). Let X_{is} be the indicator random variable for this event.

We easily see that

$$\begin{aligned} \Pr[X_{is} = 1] &= \Pr[\text{Bin}(\mu, \frac{1}{2}) \geq \mu - \Delta] \leq 2^{-(\mu - \Delta)} \binom{\mu}{\mu - \Delta} \\ &\leq 2^{-\mu} (2\mu)^\Delta, \end{aligned}$$

where the estimate for the binomial distribution is well-known (see [43, Lemma 3] or [29, Lemma 10.37]).

Regarding the correlation of the X_{is} , we see that either X_{is} is a fresh random sample independent from all $X_{i's'}$ with $s' < s$ and $i' \in [n]$ or, namely if the reduced parent population in iteration s has the i -th bit converged, $X_{is} = 0$ with probability one. Consequently, the number $X := \sum_{s=1}^t \sum_{i=1}^n X_{is}$ of unsafe bit-positions in the time frame $[t]$ is dominated by a sum of nt independent Bernoulli random variables with

success probability $2^{-\mu} (2\mu)^\Delta$, see [44, Lemma 11] or [29, Lemma 10.22].

For these reasons, we have $E[X] \leq \frac{1}{4}n$ and $\Pr[X \geq \frac{1}{2}n] \leq \exp(-\frac{1}{8}n)$ by the additive Chernoff bound.

We now argue that having at least Δ zeros in some bit-position is often enough sufficient for the position not being fixed to 1. For each $s \in [t]$, let B_s be the following event.

- If in the sampling process of the s -th population at least $\frac{n}{2}$ bit-positions are not already fixed (“fat s -th population”), then B_s is the event that there is a fitness value $z \in [0..n]$ such that at least Δ individuals of the s -th population have fitness exactly z .
- Otherwise (“thin s -th population”) let the event B_s be true with probability $p := n^{1-\Delta/2} (2\mu)^\Delta$ independent of all other random decisions of the algorithm.

Since a random variable with binomial distribution with parameters n and $\frac{1}{2}$ attains each value in $[0..n]$ with probability at most $2/\sqrt{n}$, this follows from elementary estimates of binomial coefficient, see, e.g., [29, Lemma 4.9], a union bound over the $n+1$ possible values of z and a similar estimate as above shows

$$\Pr[B_s] \leq (n+1) \binom{\mu}{\Delta} 2^\Delta n^{-\Delta/2} \leq n^{1-\Delta/2} (2\mu)^\Delta = p$$

also in the first case, where the last estimate exploits that $\Delta \geq 2$.

Denote by $\bar{B} = \bigwedge_{s=1}^t \bar{B}_s$ the event that none of the B_s comes true. By a simple union bound,

$$\Pr[(X \leq \frac{n}{2}) \wedge \bar{B}] \geq 1 - \exp(-\frac{n}{8}) - tp \geq 1 - \exp(-\frac{n}{8}) - \frac{1}{4}n^{-c}.$$

A simple induction shows that the event “ $(X \leq \frac{n}{2}) \wedge \bar{B}$ ” implies that in each iteration $s \in [2..t]$ at most those bit-positions which have been unsafe before can be converged to one. For $s=2$ this follows from the fact that all positions of the initial population are sampled randomly; consequently, the event \bar{B}_2 means that all fitness values occurred less than Δ times. This, however, implies that only a bit-position i which was unsafe in the first iteration (that is, $X_{i1} = 1$) can be converged in the second population. Since the total number of unsafe positions is at most $n/2$, also the second population is fat, that is, contains at least $n/2$ random bits. Repeating the previous arguments, we obtain that at most bit-positions which were unsafe at least once can be converged to one, and further, that all populations up to time t are fat.

Conditioning on the event “ $(X \leq \frac{n}{2}) \wedge \bar{B}$ ”, we now regard the difference between the artificial process and a true run of the CSA. We have just seen that during the run of the artificial process, at least $tn/2$ times a bit-position was sampled randomly (without becoming unsafe). The number of ones in such a bit-position is described by a random variable $(Z \mid Z \leq \mu - \Delta)$, where Z follows a binomial law with parameters μ and $\frac{1}{2}$. In particular, with probability at least $2^{-\mu}$, this number is zero. Note that when a bit-position is sampled with zero ones, then the true process differs from the artificial process and the run of the CSA reaches a state from which it cannot generate the optimum of OM. The probability

that none of the at least $tn/2$ safe samplings of variables leads to this negative event is at most

$$\begin{aligned} (1 - 2^{-\mu})^{tn/2} &\leq \exp\left(-\frac{2^{-\mu}tn}{2}\right) \\ &\leq \exp\left(-\frac{2^{-\mu}n}{2} \frac{2^\mu}{4(2\mu)^\Delta}\right) = \exp\left(-\frac{n}{16(2\mu)^\Delta}\right). \end{aligned}$$

In summary, we see that with probability at least

$$\begin{aligned} (1 - \exp(-\frac{n}{8}) - \frac{1}{4}n^{-c}) \left(1 - \exp\left(-\frac{n}{16(2\mu)^\Delta}\right)\right) \\ = 1 - O(n^{-c}), \end{aligned}$$

the run of the true CAS fixes a position to zero. \square