

Black-Box Search by Unbiased Variation

Per Kristian Lehre* and Carsten Witt†

Technical University of Denmark

Kgs. Lyngby, Denmark

{pkle,cfw}@imm.dtu.dk

September 30, 2010

Abstract

The complexity theory for black-box algorithms, introduced by Droste et al. (2006), describes common limits on the efficiency of a broad class of randomised search heuristics. There is an obvious trade-off between the generality of the black-box model and the strength of the bounds that can be proven in such a model. In particular, the original black-box model allows polynomial complexity for certain NP-complete problems and provides for well-known benchmark problems relatively small lower bounds, which are typically not met by popular search heuristics.

In this paper, we introduce a more restricted black-box model which we claim captures the working principles of many randomised search heuristics including simulated annealing, evolutionary algorithms, randomised local search and others. The key concept worked out is an unbiased variation operator. Considering this class of algorithms, significantly better lower bounds on the black-box complexity are proved, amongst them an $\Omega(n \log n)$ bound for functions with unique optimum. Moreover, a simple unimodal function and gap functions are considered. We show that a simple (1+1) EA is able to match the runtime bounds in several cases.

*Supported by EPSRC under grant no. EP/D052785/1, and Deutsche Forschungsgemeinschaft (DFG) under grant no. WI 3552/1-1.

†Supported by Deutsche Forschungsgemeinschaft (DFG) under grant no. WI 3552/1-1.

1 Introduction

The theory of randomised search heuristics has advanced significantly over the last years. In particular, there exist now rigorous results giving the runtime of canonical search heuristics, like the (1+1) EA, on combinatorial optimisation problems [19]. Ideally, these theoretical advances will guide practitioners in their application of search heuristics. However, it is still unclear to what degree the theoretical results that have been obtained for the canonical search heuristics can inform the usage of the numerous, and often more complex, search heuristics that are applied in practice. While there is an ongoing effort in extending runtime analysis to more complex search heuristics, this often requires development of new mathematical techniques.

To advance the theoretical understanding of search heuristics, it would be desirable to develop a computational complexity theory of search heuristics. The basis of such a theory would be a computational model that captures the inherent limitations of search heuristics. The results would be a classification of problems according to the time required to solve the problems in the model. Such a theory has already been introduced for local search problems [11]. The goal in *local search* problems is to find any solution that is locally optimal with regards to the cost function and a neighbourhood structure over the solution set. Here, we are interested in *global search* problems, i.e. where the goal is to find any globally optimal solution.

Droste et al. introduced a model for global optimisation called the *black-box model* [5]. The framework considers search heuristics, called black-box algorithms, that optimise functions over some finite domain S . The black-box algorithms are limited by the amount of information that is given about the function to be optimised and have to make queries to an oracle in order to determine fitness values. In this framework, lower bounds on the number of queries required for optimisation can be obtained. An advantage of the black-box model is its generality. The model covers any realistic search heuristic. Despite this generality, the lower bounds in the model are in some cases close or equal to the corresponding upper bounds that hold for particular search heuristics. For example, it is shown that the needle-in-the-haystack and trap problems are hard problems, having black-box complexity of $(2^n + 1)/2$ [5].

However, the black-box model also has some disadvantages. For example, it is proved that the NP-hard MAX-CLIQUE problem has polynomial black box complexity [5]. One cannot expect any realistic search heuristic to solve all instances of this problem in polynomial time. Indeed, the black-box algorithm that achieves the polynomial runtime is contrived. The algorithm first queries for the function values of all possible pairs of nodes, thus uncovering all the edges in the graph. Given knowledge about the instance, the algorithm can find the optimal solution through offline brute force search, without making any further queries to the black-box. This algorithm requires in total no more than $\binom{n}{2} + 1$ function evaluations. Similar black-box algorithms can be envisaged for other NP-hard problems. This result exposes two weaknesses: the lower bounds in the model are often obtained by black-box algorithms that do not resemble

any randomised search heuristic. Secondly, the model is too unrestricted with respect to the amount of resources disposable to the algorithm. Black-box algorithms can spend unlimited time in-between queries to do computation.

In order to define a more realistic black-box model, one should consider which additional restrictions to consider. Droste et al. suggested to limit the available storage space available to the black box algorithm [5], but did not prove any lower bounds in the space restricted scenario.

The remaining of this paper is organised as follows. Section 2 introduces the new black-box model along with a description of the unbiased variation operators. Section 3 provides the first lower bound in the model for a simple, unimodal problem. Then, in Section 4, we consider a function class that contains a plateau. Finally, in Section 5, we prove lower bounds that hold for any function with a single, global optimum. The paper is concluded in Section 6.

2 A Refined Black-Box Model

We now present the refined black-box model that is obtained by imposing two additional restrictions in the old black-box model. We start with a preliminary informal description and motivation.

Firstly, we limit the degree of independence between the queries. The initial query is a bitstring chosen uniformly at random. Every subsequent query must be made for a search point that is obtained by applying an unbiased variation operator to one of the previously queried search points. Unbiased variation operators will be defined in Section 2.1.

Secondly, we put an additional restriction on the information that is available to the algorithm by preventing the algorithm from observing the bit values of the search points that are queried. Hence, the only information that can be exploited by the algorithm is the sequence of fitness values obtained when making queries to the black-box, and not the search points themselves. Note that without this restriction, any black-box algorithm could be simulated as follows: For each query x made by the unrestricted algorithm, the restricted algorithm would solve the problem corresponding of minimising the Hamming distance to search point x . This simulation would have an expected overhead factor of $O(n \log n)$ function evaluations.

We now turn our considerations into a formal definition of a refined black-box model, as stated in Algorithm 1. Let us first pick up the unrestricted black-box scenario [5]. The black-box algorithm A is given a class of pseudo-Boolean functions \mathcal{F} . An adversary selects a function f from this class and presents it to the algorithm as a black-box. At this point, the only information available to the algorithm about the function f is that it belongs to function class \mathcal{F} . The black-box algorithm can now start to query an oracle for the function values of any search points. The runtime $T_{A,f}$ of the algorithm on function f is the number of function queries on f until the algorithm queries the function value of an optimal search point for the first time. The runtime $T_{A,\mathcal{F}}$ on the class of functions is defined as the maximum runtime over the class of functions. In the

Algorithm 1 Unbiased Black-Box Algorithm

- 1: $t \leftarrow 0$.
 - 2: Choose $x(t)$ uniformly at random from $\{0, 1\}^n$.
 - 3: **repeat**
 - 4: $t \leftarrow t + 1$.
 - 5: Compute $f(x(t - 1))$.
 - 6: $I(t) \leftarrow (f(x(0)), \dots, f(x(t - 1)))$.
 - 7: Depending on $I(t)$, choose a probability distribution p_s on $\{0, \dots, t - 1\}$.
 - 8: Randomly choose an index j according to p_s .
 - 9: Depending on $I(t)$, choose an unbiased variation operator $p_v(\cdot \mid x(j))$.
 - 10: Randomly choose a bitstring $x(t)$ according to p_v .
 - 11: **until** termination condition met.
-

unbiased black-box model, queries of new search points must be made according to Algorithm 1. I.e., the initial search point is chosen uniformly at random by an oracle, and subsequent search points are obtained by asking the oracle to apply a given unbiased variation operator to a previously queried search point.

The *unbiased black-box complexity* of a function class \mathcal{F} is the minimum worst case runtime $T_{A, \mathcal{F}}$ among all unbiased black-box algorithms A satisfying the framework of Algorithm 1. Hence, any upper bound on the worst case of a particular unbiased black-box algorithm, also implies the same upper bound on the unbiased black-box complexity. To prove a lower bound on the unbiased black-box complexity, it is necessary to prove that the lower bound holds for any unbiased black-box algorithm. Note that since unbiased black-box algorithms are a special case of black-box algorithms, all the lower bounds that hold for the general black-box model also hold for the unbiased black-box model.

2.1 Unbiased Variation Operators

We formalise variation operators as conditional probability distributions over the search space. Given k search points x_1, \dots, x_k , a variation operator p produces a search point y with probability $p(y \mid x_1, \dots, x_k)$. We put some restrictions on the probability distribution to capture the essential characteristics of the variation operators that are used by the common randomised search heuristics. Firstly, one can limit the number k of search points that are used to produce the new search point. The number k determines the *arity* of the variation operator. Here, we will only consider *unary* variation operators, i.e. when $k = 1$. Furthermore, we will impose the following two unbiasedness-conditions on the operators:

- 1) $\forall x, y, z, p(y \mid x_1, \dots, x_k) = p(y \oplus z \mid x_1 \oplus z, \dots, x_k \oplus z)$,
- 2) $\forall x, y, \sigma, p(y \mid x_1, \dots, x_k) = p(\sigma_b(y) \mid \sigma_b(x_1), \dots, \sigma_b(x_k))$,

where \oplus denotes the exclusive or-operation on bitstrings, and for any permutation σ over $[n]$, σ_b is an associated permutation over the bitstrings defined

as

$$\sigma_b(x_1x_2 \cdots x_n) := x_{\sigma(1)}x_{\sigma(2)} \cdots x_{\sigma(n)}.$$

Variation operators that satisfy the first condition are called \oplus -invariant, while variation operators that satisfy the second condition are called σ -invariant. In this paper, unbiased variation operators is defined as variation operators that satisfy both conditions, i.e. \oplus - σ -invariant operators. Note that this is a special case of the framework by Rowe, Vose and Wright [21], who study invariance from a group-theoretical point of view.

We claim that the two conditions are natural. Firstly, the variation operators used by common randomised search heuristics are typically \oplus - σ -invariant, including flipping a randomly chosen bit position, bitwise mutation with any mutation probability, and uniform sampling of bitstrings. Furthermore, the black-box algorithm is allowed to select a different variation operator in each iteration, as long as the variation operators are statistically independent. This generality allows adaptive variation operators to be covered, including rank-based mutation [20]. Secondly, the two conditions on the variation operators can also be motivated by practical concerns. In applications, the set of bitstrings typically encode the variable settings of candidate solutions, and is rarely the optimisation domain *per se*. Hence, the encoding from variable setting to bit value, and from variable to bitstring position, can be arbitrary. E.g., whether the user encodes a “high temperature” variable setting as 0 or 1, or decides to encode the temperature variable by the first instead of the last variable in the bitstring, should not influence the behaviour of a search heuristic.

Droste and Wiesmann recommended that all search points that are within the same distance of the originating search point should have the same probability of being produced [6]. This unbiasedness criterion, which we call *Hamming-invariance*, can be formalised as follows:

$$3) \forall x, y, z, \quad H(x, y) = H(x, z) \implies p(y \mid x) = p(z \mid x),$$

where $H(x, y)$ denotes the Hamming distance between x and y . We now show that these criteria are related.

Proposition 1. *Every unary variation operator that is \oplus - σ -invariant is also Hamming-invariant.*

Proof. Assume that p is an \oplus - σ -invariant variation operator. Let x, y, u and v be any search points such that $d(x, y) = d(u, v) =: d$. Given these assumptions, we will prove that $p(y \mid x) = p(v \mid u)$ holds. There must exist a permutation σ such that $\sigma(y \oplus x) = v \oplus u$. We then have

$$\begin{aligned} p(y \mid x) &= p(y \oplus x \mid 0^n) = p(\sigma(y \oplus x) \mid 0^n) \\ &= p(v \oplus u \mid 0^n) = p(v \mid u). \end{aligned}$$

By setting $u = x$, it is clear that p is Hamming-invariant. □

When referring to unbiased black-box algorithms in the following, we mean any algorithm that follows the framework of Algorithm 1. In particular, by *unary, unbiased black-box algorithms*, we mean such algorithms that only use unary, unbiased variation operators. The class of unary, unbiased black-box algorithms is general, and includes many well known algorithms, including simulated annealing [13], random local search (RLS) [17], $(\mu+\lambda)$ EA [4, 10, 23], and many other population-based EAs that do not use crossover. Note that the restriction to unary, unbiased variation operators excludes some randomised search heuristics. In particular, the model does not cover EAs that use crossover. Many of the commonly used diversity mechanisms are excluded. Also, estimation of distribution algorithms [14], ant colony optimisation [3] and particle swarm optimisation [12] are not covered by the model.

3 Simple Unimodal Functions

As an initial example, we consider the simple unimodal function LEADINGONES $(x) = \sum_{i=1}^n \prod_{j=1}^i x_j$. (The function ONEMAX is covered by the results in Section 5.) The expected runtime of the (1+1) EA on this function is $\Theta(n^2)$ [4], which seems optimal among commonly analysed EAs. Increasing either the offspring or parent population sizes does not reduce the runtime. For $(\mu+1)$ EA, the runtime is $\Theta(n^2 + \mu n \log n)$ [23], and for $(1+\lambda)$ EA, the runtime is $\Theta(n^2 + \lambda n)$ [10].

We now show that the runtime of the (1+1) EA on LEADINGONES is asymptotically optimal in the unbiased black-box model. We define the potential of the algorithm at time step t as the largest number of leading 1- or 0-bits obtained so far, i.e. $k := \max_{0 \leq j \leq t} \{\text{LO}(x(j)), \text{LZ}(x(j))\}$. The number of 0-bits must be considered because flipping every bit in a bitstring with i leading 0-bits will produce a bitstring with i leading 1-bits. The increase of the potential will be studied using drift analysis [7, 8]. To lower bound the drift, it is helpful to proceed as in the analysis of (1+1) EA on LEADINGONES [4], i.e. to first prove that the substring after the first 0-bit in a given time step is uniformly distributed. We will prove a more general statement that will also be used in Section 4. For notational convenience, subsets of $[n]$ and bitstrings of length n will be used interchangeably, i.e., the bitstring $x \in \{0, 1\}^n$ is associated with the subset $\{i \in [n] \mid x_i = 1\}$.

Lemma 1. *For any $t \geq 0$, let $X(t) = \{x(0), x(1), \dots, x(t)\}$ be the search points visited by any unary, unbiased black-box algorithm until iteration t when optimising function f . If there exists a subset of indices $z \subseteq [n]$ such that*

$$\forall y \subseteq z, \forall x \in X(t) \quad f(x \oplus y) = f(x),$$

then the bits $\{x_i(t) \mid i \in z\}$ are independent and uniformly distributed.

Proof. We first prove the claim that for all $t \geq 0$, $b \in \{0, 1\}$, and $i \in z$ it holds that $\Pr[x_i(t) = b] = 1/2$. The proof is by induction over the time step

t . The initial search point is sampled uniformly at random, so the claim holds for $t = 0$. Assume that the claim holds for $t = t_0 \geq 0$. The fitness profile $I(t_0)$ does not depend on any of the bits in position i in the previously visited search points. The choice of search point $x(j)$ the algorithm makes is therefore independent of these bits. By the induction hypothesis, search point $x(j)$ has a bit value of b in position i with probability $p := 1/2$. Letting r be the probability that the variation operator flips bit position i , we have $\Pr[x_i(t_0 + 1) = b] = p(1 - r) + (1 - p)r = 1/2$. By induction, the claim now holds for all t .

We now prove the lemma by induction over the time t . The lemma holds for time step $t = 0$. Assume that the lemma holds for $t = t_0 \geq 0$. Let $U \in \{0, 1\}^n$ be a random vector where the elements $U_i, 1 \leq i \leq n$, take the value $U_i = 1$ if bit position i flipped in step t_0 , and $U_i = 0$ otherwise. Then for all bitstrings $y \in \{0, 1\}^n$,

$$\Pr \left[\bigwedge_{i \in z} x_i(t_0 + 1) = y_i \right] = \sum_{u \in \{0, 1\}^n} \Pr[U = u] \Pr \left[\bigwedge_{i \in z} x_i(t_0) = y_i \oplus u_i \right],$$

which by the induction hypothesis (more precisely, the statement about independence) equals

$$\sum_{u \in \{0, 1\}^n} \Pr[U = u] \prod_{i \in z} \Pr[x_i(t_0) = y_i \oplus u_i] = 2^{-|z|} = \prod_{i \in z} \Pr[x_i(t_0 + 1) = y_i],$$

where the last equality follows from the claim above. This proves the independence at time $t_0 + 1$, and, therefore, the induction step. \square

Theorem 1. *The expected runtime of any unary, unbiased black-box algorithm on LEADINGONES is $\Omega(n^2)$.*

Proof. We first prove the claim that w.o.p., the potential of the algorithm will at some time be in the interval between $n/2$ and $3n/4$, before the optimum has been found. With probability $1 - 2^{-n/2}$, the initial search point will have potential less than $n/2$. Let integer i be the number of 0-bits in the interval from $n/2$ and $3n/4$ in the bitstring that the algorithm selects next. By Lemma 1 with the index set $z := [n/2 + 1, n]$, and a Chernoff bound, there is a constant $\delta > 0$ such that with overwhelming probability, this integer satisfies $(1 - \delta)n/8 < i < (1 + \delta)n/8$. In order to increase the potential from less than $n/2$ to at least $3n/4$, it is necessary to flip every 0-bit in the interval from $n/2$ to $3n/4$ and no other bits. We optimistically assume that exactly i bits are flipped in this interval. However due to the unbiasedness condition, every choice of i among $n/4$ bits to flip is equally likely. The probability that only the 0-bits are flipped is therefore at most $\binom{n/4}{i}^{-1} \leq (4i/n)^i \leq ((1 + \delta)/2)^{n/4(1 - \delta)}$. The claim therefore holds.

We now apply drift analysis according to the potential k of the algorithm, only counting the steps starting from a potential in the interval $n/2 \leq k < 3n/4$. In order to find the optimum, the potential must be increased by at least $n/4$. Assume that the selected search point has r 0-bits in the first $k + 1$ bit positions.

In order to increase the potential, it is necessary to flip all these 0-bits, and none of the 1-bits within this interval. This corresponds to consecutively selecting all the r red balls from an urn containing $k + 1$ balls. The probability of this event is less than

$$\frac{r}{k+1} \cdot \frac{r-1}{k} \cdots \frac{2}{k-r+2} \cdot \frac{1}{k-r+1} \leq \frac{1}{k+1}.$$

The drift in each step is bounded from above by $\Delta_i(t) \leq (1 + \mathbf{E}[Y_t])/(k + 1)$, where random variable Y_t is the number of free-riders [4] in step t . Applying Lemma 1 with the index set $z := [k + 2, n]$ gives $\mathbf{E}[Y_t] \leq \sum_{i=1}^{\infty} 2^{-i} \leq 1$. The polynomial drift theorem [8] now implies that $\mathbf{E}[T] \geq (n/4)/\Delta_t(i) = (k + 1)n/4 = \Omega(n^2)$. \square

Note that the complexity of LEADINGONES in the unrestricted black-box model¹ is bounded above by $n/2 - o(n)$ [5]. This illustrates that the complexity of a function class can be significantly higher in the unbiased black-box model than in the unrestricted black-box model.

4 Enforcing Expected Runtimes

We are interested in problem classes where the unbiased black-box complexity depends on some parameter of the problem. More specifically, is there a class of functions $\mathcal{F} = \{f_i \mid i \in \mathbb{Z}\}$, such that the unbiased black-box complexity of the functions f_m in the class \mathcal{F} increases with the problem parameter m , e. g. as $\Theta(2^m)$?

In the case of the (1+1) EA, it is known that the runtime can depend on the size of so-called Hamming cliffs. A Hamming cliff of size m is a search point where all other search points within Hamming distance m are of inferior quality. Hamming cliffs are common in combinatorial optimisation problems. For example, on the minimum spanning tree problem, the (1+1) EA must flip two edges to make an improvement [17]. This is accounted for by an $\Omega(n^2)$ factor in the expected runtime. More generally, there exists combinatorial optimisation problems that for any m have instances with Hamming cliffs of size m [16]. The (1+1) EA needs $\Omega(n^m)$ steps to overcome a Hamming cliff of size m , and there is an example called JUMP _{m} [4] where also an upper bound $O(n^m)$ holds, i. e., an expected runtime can be enforced. This corresponds to a hierarchy consisting of a class of functions with increasing difficulty.

We are aiming at generalising this result to black-box algorithms with unary unbiased variation. We pick up the general idea of the JUMP _{m} function in [4],

¹Unlike unbiased black-box algorithms, unrestricted black-box algorithms (the general model studied in [5]) can optimise function classes containing a single function in constant time by querying the optimum in the first iteration. To obtain meaningful results in the unrestricted model, it is therefore necessary to consider the generalised class of functions containing $\text{LEADINGONES}_z(x) := \text{LEADINGONES}(x \oplus z)$ for every bitstring z .

but modify the “gap”. For any m , where $0 \leq m \leq n$, define

$$\text{JUMP}_m(x) := \sum_{i=1}^{n-m} x_i + \prod_{i=1}^m x_i.$$

Hence, the gap corresponds to the last m bit positions which do not influence the function value, except at the optimal search point 1^n , i.e., a *plateau*. For $m = 0$, we obtain the easy ONEMAX function, and for $m = n$, we obtain the hard NEEDLE function. The following result shows that the hardness of JUMP_m depends on the problem parameter m .

Theorem 2. *For $0 \leq m \leq n$, the expected optimisation time of any unary, unbiased black-box algorithm on JUMP_m is at least 2^{m-2} . Furthermore, the probability that the optimum is found within $2^{m(1-\varepsilon)}$ iterations for any $\varepsilon, 0 < \varepsilon < 1$, is no more than $2^{-\varepsilon m}$.*

Proof. The time T to find the optimum is bounded from below by the time T' until all the last m bits are 1-bits. The time is analysed as if the algorithm was presented with the function $f(x) = \sum_{i=1}^{n-m} x_i$ instead of JUMP_m . This function differs from JUMP_m only on the optimal search point 1^n . The distribution of T' will therefore be the same for f and JUMP_m .

Lemma 1 now applies for the function f and the set of indices $z = [n - m + 1, n]$. The probability that the search point visited in any given iteration has 1^m as suffix is 2^{-m} . By union bound, the probability that this suffix is obtained before iteration $t \geq 0$ is no more than $t \cdot 2^{-m}$. In particular, the probability that the optimisation time is shorter than 2^{m-1} is less than $1/2$, so the expected optimisation time is at least 2^{m-2} . Furthermore, the probability that the optimum has been found within $2^{m(1-\varepsilon)}$ iterations is no more than $2^{-\varepsilon m}$. \square

We supplement an upper bound.

Theorem 3. *There exists a unary, unbiased black-box algorithm whose expected runtime on JUMP_m is $ne^m + O(n \log n)$.*

Proof. We show that the well-known algorithm random local search (RLS) [17] has the stated expected runtime.

The *first phase* lasts as long as the algorithm can increase the number of 1-bits in the first $n - m$ positions. By a coupon collector argument, the expected time of this phase is $O(n \log n)$.

In the *second phase*, the function value of the current search point is $n - m$, and the algorithm will only accept a new search point if it was obtained by flipping one of the m bits in the suffix. We call such steps *relevant*, and we ignore the other steps. Assume that the current search point has $i, 1 \leq i \leq m$, 0-bits in the suffix. To reach the optimum, it is sufficient to have i relevant steps, each flipping a 0-bit. The probability of such a sequence can be bounded

by Stirling's approximation as

$$\frac{i}{m} \cdot \frac{(i-1)}{m} \cdots \frac{1}{m} = \frac{i!}{m^i} \geq \left(\frac{i}{em}\right)^i.$$

The right hand side of this inequality is monotonically decreasing in $i \leq m$, and hence at least e^{-m} . The expected waiting time for one relevant step is n/m . By linearity of expectation, the expected waiting time for i relevant steps is $in/m < n$. The expected time until the suffix of the current search point only contains 1-bits is no more than ne^m .

The theorem now follows by adding the expected durations of the two phases. \square

5 General Functions

In the previous sections, we provided bounds on particular pseudo-Boolean functions that are commonly considered in the runtime analysis of randomised search heuristics. In this section, we focus on finding lower bounds that hold for any function. Such bounds are only interesting when we consider functions that correspond to realistic optimisation problems, as trivial functions like constant functions can be optimised with a single function evaluation. We therefore focus on functions that have a unique global optimum.

It is of interest to compare the lower bounds in the black-box models with those bounds that have been obtained for specific EAs. Wegener proved a lower bound of $\Omega(n \log n)$ for the (1+1) EA on any function with a unique optimum [22]. This bound is significantly larger than the $\Omega(n/\log n)$ bound that holds for a generalisation of the ONEMAX problem in the black-box model [5]. Given this discrepancy, one can ask whether there is room to design better EAs which overcome the $n \log n$ barrier, or whether the black-box bound is too loose. Jansen et al. provided evidence that there is little room for improvement by showing that any EA that uses uniform initialisation, selection and bitwise mutation with probability $1/n$ needs $\Omega(n \log n)$ function evaluations to optimise functions with a unique optimum [10].

In the following, we will generalise this result further, showing that the $n \log n$ -barrier for functions of a unique optimum even holds for the wider class of unary, unbiased black-box algorithms. The idea behind the proof is to show that the probability of making an improving step reduces as the algorithm approaches the optimum. To implement this idea, we will apply Theorem 5. This is a lower-bound analogue to a technique which is called expected multiplicative weight decrease in the evolutionary computation literature [18]. Theorem 5 will be proved using the following polynomial drift-theorem.

Theorem 4 ([9]). *Let X_1, X_2, \dots be random variables with bounded support and let T be the stopping time defined by $T := \min\{t \mid X_1 + \dots + X_t \geq g\}$ for a given $g > 0$. If $\mathbf{E}[T]$ exists and $\mathbf{E}[X_i \mid T \geq i] \leq u$ for $i \in \mathbb{N}$, then $\mathbf{E}[T] \geq g/u$.*

Theorem 5. Let $\beta(n), \gamma(n)$, and $b(n)$ be positive reals where $b(n) > 2\beta(n)$. Let Z_1, Z_2, \dots be a stochastic process with bounded support on the set of non-negative integers, and define T to be the smallest t such that $Z_t = 0$. If for all $t > 1$, and z , where $0 < z \leq b(n)$, it holds that

1. $\mathbf{E}[z - Z_t \mid Z_{t-1} = z, T > t] \leq z\gamma(n)$
2. $\mathbf{Pr}[z - Z_t \geq \beta(n) \mid Z_{t-1} = z, T > t] \leq \gamma(n)/z$,

then

$$\mathbf{E}[T \mid Z_1 \geq b(n)] \geq \frac{1}{3\gamma(n)} \cdot \ln\left(\frac{b(n) + 1}{2\beta(n) + 1}\right).$$

Proof. The proof generalises the proof of Theorem 1 in [2]. The random variable T is non-negative, so if the expectation of T does not exist, then it is positive infinite and the theorem holds. We condition on the events $T > t$ and $Z_{t-1} = z$, but we omit stating these events in the expectations for notational convenience. We define the stochastic process $Y_t := \ln(Z_t + 1)$, and apply Theorem 4 with respect to the random variables $\Delta_t(z) := (Y_{t-1} - Y_t \mid Z_{t-1} = z)$. For technical reasons, we only consider the time until $Z_t \leq 2\beta(n)$, and therefore use the parameter $g := \ln(b(n) + 1) - \ln(2\beta(n) + 1) > 0$. We then have

$$\begin{aligned} \mathbf{E}[\Delta_t(z)] &= \mathbf{E}[\ln(z + 1) - \ln(Z_t + 1)] \\ &= \mathbf{E}[\ln((z + 1)/(Z_t + 1))]. \end{aligned}$$

The logarithmic function is concave, so by Jensen's inequality, this is less than

$$\ln\left(\mathbf{E}\left[\frac{z + 1}{Z_t + 1}\right]\right) = \ln\left(1 + \mathbf{E}\left[\frac{z - Z_t}{Z_t + 1}\right]\right).$$

By using the inequality $\ln(1 + x) \leq x$, and the law of total probability, this expectation can be bounded by the two terms

$$\begin{aligned} \mathbf{E}\left[\frac{z - Z_t}{Z_t + 1}\right] &\leq \mathbf{Pr}[z - Z_t \leq \beta(n)] \mathbf{E}\left[\frac{z - Z_t}{Z_t + 1}\right] + \\ &\quad \mathbf{Pr}[z - Z_t > \beta(n)] \mathbf{E}\left[\frac{z - Z_t}{Z_t + 1} \mid z - Z_t > \beta(n)\right]. \end{aligned} \quad (1)$$

In the first term, we have omitted the condition $z - Z_t \leq \beta(n)$, which will only increase the expectation. By condition 1, and the assumption that $z \geq 2\beta(n)$, the first term in Eq. (1) is bounded from above by

$$\frac{z\gamma(n)}{z - \beta(n)} = \frac{\gamma(n)}{(1 - \beta(n)/z)} \leq 2\gamma(n).$$

It follows from the assumption $T > t$ that $Z_t \geq 1$ and hence $(z - Z_t)/(Z_t + 1) \leq z$. The second term in Eq. (1) can therefore be bounded by applying the second condition, giving $\mathbf{Pr}[z - Z_t > \beta(n)] z \leq \gamma(n)$.

From the above, one can conclude that $\mathbf{E}[\Delta_t(z)] \leq 3\gamma(n)$. From Theorem 4, it now follows that $\mathbf{E}[T \mid Z_1 \geq b(n)] \geq \ln((b(n) + 1)/(2\beta(n) + 1))/3\gamma(n)$. \square

When analysing unbiased black-box algorithms, it is helpful to model the application of an unbiased variation operator as a classical urn experiment. Assume that the black-box algorithm chooses a search point that has m 0-bits and that the variation operator flips r bits. This corresponds to drawing r balls without replacement from an urn containing m red balls and $n - m$ white balls. The number of red balls Z in the sample, i.e., the number of flipped 0-bits, is a *hypergeometrically* distributed random variable with expectation rm/n .

If the optimal search point is 1^n , the improvement made by the algorithm in one step can be expressed as $Z - (r - Z) = 2Z - r$, i.e., as the reduction in the number of 0-bits. Clearly, the algorithm only makes an improvement if this number is positive. Hence, in order to apply Theorem 5, it will be helpful to have an estimate of the expectation of a hypergeometric random variable, conditional on the event that this variable takes at least a certain value.

Lemma 2. *Let Z be a hypergeometrically distributed random variable with parameters n (number of balls), r (number of samples) and m (number of red balls), then for all $k, 0 \leq k \leq r$, $\mathbf{E}[Z \mid Z \geq k] \leq k + (r - k)(m - k)/(n - k)$.*

Proof. The remaining number of trials where red balls can be obtained is maximised if already all of the first k sampled balls were red. Then the number of additionally sampled red balls is denoted by Y and hypergeometrically distributed with parameters $n - k$, $r - k$ and $m - k$. Hence,

$$\begin{aligned} \mathbf{E}[Z \mid Z \geq k] &= k + \sum_{i=k}^r (\Pr[Z = i \mid Z \geq k] \cdot \mathbf{E}[Z - k \mid Z = i \geq k]) \\ &= k + \sum_{i=k}^r (\Pr[Z = i \mid Z \geq k] \cdot (i - k)) \\ &\leq k + \sum_{i=k}^r (\Pr[Y = i - k] \cdot (i - k)) \\ &= k + \mathbf{E}[Y] = k + \frac{(r - k) \cdot (m - k)}{n - k}. \end{aligned}$$

□

We also need upper bounds on the tail of the hypergeometric distribution. The following result due to Chvátal [1] is an analogue to the Chernoff bounds for the binomial distribution.

Lemma 3 (Chvátal, [1]). *If X is a hypergeometrically distributed random variable with parameters n (number of balls), m (number of red balls) and r (number of samples), then $\Pr[X \geq \mathbf{E}[X] + r\delta] \leq \exp(-2\delta^2 r)$, where $\mathbf{E}[X] = \frac{rm}{n}$.*

We now state the main result of this section.

Theorem 6. *The expected runtime of any unary, unbiased black-box algorithm on any pseudo-Boolean function with a single global optimum is $\Omega(n \log n)$.*

Proof. Without loss of generality, assume that the optimum is the search point 1^n . Since this optimum can easily be obtained from search point 0^n by flipping all bits, the runtime will be bounded by the number of steps until either 1^n or 0^n is sampled for the first time. The potential P_t of the algorithm in a given iteration $t \geq 1$ is defined as the shortest Hamming distance from any previously sampled search point to either 1^n or 0^n . To find the optimum, it is necessary to reduce the potential to 0.

We consider it a failure if the initial potential is less than $b(n) := n/3$. By a Chernoff bound, the probability of this failure event is $e^{-\Omega(n)}$.

Assuming no failure occurs, we estimate the time to reduce the potential to 0 by applying Theorem 5 to the stochastic process $P_{t \geq 0}$. Clearly, this process has bounded support. If we can prove that the conditions of the theorem hold for the parameters $b(n) = n/3$, $\beta(n) = \ln^2(n)$, and $\gamma(n) = O(1/n)$, then the expected runtime conditional on no failure is $\Omega(n \log n)$. Since the failure probability is only $e^{-\Omega(n)}$, this also implies that the unconditional, expected runtime is $\Omega(n \log n)$. It therefore remains to show that the conditions of Theorem 5 holds.

We first verify the second condition. Let i be the number of 0-bits in the search point selected in iteration t , and $r \geq 1$ be the number of bits that was flipped by the variation operator. The number of 0-bits that is flipped by the variation operator is a hypergeometrically distributed random variable Z_t with parameters n (number of balls), $i < n/3$ (number of red balls) and r (number of samples). The expectation of Z_t is $ri/n < r/3$. In order to reduce the potential by $\beta(n)$, it is necessary to flip at least $r \geq \beta(n)$ bits. Furthermore, the distance to the optimum can only be reduced if the variation operator flips more 0-bits than 1-bits, i.e. $Z_t \geq r - Z_t$. The probability of this event is

$$\Pr[Z_t \geq r/2] \leq \Pr[Z_t \geq ri/n + r/6] = \Pr[Z_t \geq \mathbf{E}[Z_t] + r/6].$$

By Lemma 3, the probability of reducing the potential when flipping at least $\beta(n)$ bits is therefore no more than $\exp(-r/18) = \exp(-\Omega(\beta(n))) = n^{-\Omega(\ln n)}$. So the second condition holds, because $\gamma(n)/z = \Omega(n^{-2})$ is asymptotically larger than $n^{-\Omega(\ln n)}$.

We then verify the first condition, i.e., we must prove that if the current potential is $P_t = \alpha n$ for any $\alpha, 0 \leq \alpha < 1/3$, then the expected reduction in potential in one iteration is bounded from above by $O(\alpha)$, independently of r and c . The expected reduction in potential after varying a search point x is first estimated conditional on the event that the chosen variation operator p_v flips exactly $r \geq 1$ bits. Conditional on this event, the varied search point x' will be uniformly sampled among all bitstrings that have Hamming distance r to search point x . Assume that the chosen search point x has $\alpha n + cr$ number of 0-bits. In order to obtain a search point with less than αn 0-bits, it is necessary that $0 \leq c < 1$. Let X be the random variable such that the number of 0-bits in the new search point x' is $\alpha n - X$.

We first consider the case where $1 \leq r < n/2$. Let random variable Z denote the number of 0-bits that are flipped. The number of 0-bits in search point x' is $\alpha n - X = \alpha n + cr - Z + (r - Z)$, hence $X = 2Z - r(1 + c)$. Random

variable Z corresponds to the number of red balls obtained after sampling r balls without replacement from an urn containing n balls, where $\alpha n + cr$ of the balls are red. Random variable Z is therefore hypergeometrically distributed with expectation $r \cdot (\alpha + cr/n)$. The potential will only decrease when the new search point x' has at least cr fewer 0-bits than x . The probability of this event is $p_z := \Pr[X \geq 0] = \Pr[Z \geq r(1+c)/2]$. By Lemma 2, when $r < n/2$, the expected reduction in potential equals

$$\begin{aligned} p_z \cdot \mathbf{E}[X \mid X \geq 0] &= p_z \cdot \mathbf{E}\left[2Z - r(1+c) \mid Z \geq \frac{r(1+c)}{2}\right] \\ &\leq 2p_z \cdot \frac{(r - r(1+c)/2) \cdot (\alpha n + cr - r(1+c)/2)}{n - r(1+c)/2} \\ &\leq p_z \frac{r\alpha n}{n-r} = \alpha p_z r \cdot 1/(1-r/n) \leq 2\alpha p_z r. \end{aligned}$$

For $r \leq n/2$ and $\alpha < 1/3$, Lemma 3 gives

$$p_z = \Pr[Z > r(1+c)/2] \leq \exp(-2t^2r)$$

where $t = (1+c)/2 - \alpha - cr/n \geq 1/6$. So, expected reduction in potential when $r < n/2$ is $2\alpha r/e^{r/18} = O(\alpha)$.

For the case where $n/2 \leq r < n$, we exploit a symmetry in the hypergeometric distribution. Instead of selecting r bit positions to flip, one can select $q := n-r$ bit positions to keep and flip the other bit positions. Assume that the selected search point contains $\alpha n + cq$ 1-bits. Clearly, the constant c is bounded by $c < 1$, otherwise more than αn 1-bits will be flipped into 0-bits, and the potential would not decrease. Let random variable O denote the number of 1-bits selected not to be flipped. The number of 0-bits in search point x' is $\alpha n + cq - O + (q - O)$, hence $X = 2O - q(1+c)$. Random variable O is hypergeometrically distributed, corresponding to the number of red balls obtained after sampling q balls without replacement from an urn containing n balls, where $\alpha n + cq$ of the balls are red. This corresponds exactly to the case when $r < n/2$, where the roles of variables q and O are replaced with r and Z . It therefore follows that the expected decrease in potential is bounded from above by $O(\alpha)$.

Both conditions of Theorem 5 hold, and the theorem follows. \square

6 Conclusions

This paper takes a step forward in building a unified theory of randomised search heuristics. We have defined a new black-box model that captures essential aspects of randomised search heuristics. The new model covers many of the common search heuristics, including simulated annealing and EAs commonly considered in theoretical studies. We have proved upper and lower bounds on the runtime of several commonly considered pseudo-Boolean functions. For some functions, the lower bounds coincide with the upper bounds for the (1+1) EA, implying that this simple EA is asymptotically optimal on the function class.

It is shown that any search heuristic in the model needs $\Omega(n \log n)$ function evaluations to optimise functions with a unique optimum. Also, it is shown that a function with a plateau can pose a difficulty for any black-box search heuristic in the model.

This work can be extended in several ways. Firstly, it is interesting to consider more problem classes than those considered here. Secondly, the analysis should be extended to variation operators with greater arity than one. Finally, alternative black-box models could be defined that cover ant colony optimisation, particle swarm optimisation and estimation of distribution algorithms.

Acknowledgements

We thank Benjamin Doerr and Timo Kötzing for pointing out some errors in an earlier version of this report, and in the publication at GECCO 2010 [15]. In particular, Section 4 has been subject to a thorough correction. We also thank Daniel Johannsen for reading the proof of Theorem 5.

References

- [1] V. Chvátal. The tail of the hypergeometric distribution. *Discrete Mathematics*, 25(3):285–287, 1979.
- [2] B. Doerr, M. Fouz, and C. Witt. Quasirandom evolutionary algorithms. In *GECCO 10: Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 1457–1464, New York, NY, USA, 2010. ACM.
- [3] M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, 2004.
- [4] S. Droste, T. Jansen, and I. Wegener. On the analysis of the (1+1) Evolutionary Algorithm. *Theoretical Computer Science*, 276:51–81, 2002.
- [5] S. Droste, T. Jansen, and I. Wegener. Upper and lower bounds for randomized search heuristics in black-box optimization. *Theory of Computing Systems*, 39(4):525–544, 2006.
- [6] S. Droste and D. Wiesmann. Metric based evolutionary algorithms. In *Proceedings of Genetic Programming, European Conference*, volume 1802 of *LNCS*, pages 29–43. Springer, 2000.
- [7] B. Hajek. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 13(3):502–525, 1982.
- [8] J. He and X. Yao. A study of drift analysis for estimating computation time of evolutionary algorithms. *Natural Computing*, 3(1), 2004.
- [9] J. Jägerskupper. Algorithmic analysis of a basic evolutionary algorithm for continuous optimization. *Theoretical Computer Science*, 39(3):329–347, 2007.

- [10] T. Jansen, K. A. D. Jong, and I. Wegener. On the choice of the offspring population size in evolutionary algorithms. *Evolutionary Computation*, 13(4):413–440, 2005.
- [11] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.
- [12] J. Kennedy and R. C. Eberhart. *Swarm intelligence*. Morgan Kaufmann Publishers Inc., 2001.
- [13] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [14] P. Larrañaga and J. A. Lozano. *Estimation of distribution algorithms: a new tool for evolutionary computation*. Kluwer Academic Publishers, 2002.
- [15] P. K. Lehre and C. Witt. Black-box search by unbiased variation. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation (GECCO'10)*, pages 1441–1448, New York, NY, USA, 2010. ACM.
- [16] P. K. Lehre and X. Yao. Runtime analysis of (1+1) EA on computing unique input output sequences. In *Proceedings of 2007 IEEE Congress on Evolutionary Computation (CEC'07)*, pages 1882–1889. IEEE Press, 2007.
- [17] F. Neumann and I. Wegener. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Theoretical Computer Science*, 378(1):32–40, 2007.
- [18] F. Neumann and C. Witt. Ant colony optimization and the minimum spanning tree problem. In *Proceedings of Learning and Intelligent Optimization (LION'2008)*, pages 153–166, 2008.
- [19] P. S. Oliveto, J. He, and X. Yao. Time complexity of evolutionary algorithms for combinatorial optimization: A decade of results. *International Journal of Automation and Computing*, 4(1):100–106, 2007.
- [20] P. S. Oliveto, P. K. Lehre, and F. Neumann. Theoretical analysis of rank-based mutation - combining exploration and exploitation. In *Proceedings of the 10th IEEE Congress on Evolutionary Computation (CEC '09)*, pages 1455–1462. IEEE, 2009.
- [21] J. E. Rowe, M. D. Vose, and A. H. Wright. Neighborhood graphs and symmetric genetic operators. In *Proceedings of Foundations of Genetic Algorithms 9*, number 4436 in LNCS, pages 110–122, 2007.
- [22] I. Wegener. Methods for the analysis of evolutionary algorithms on pseudo-Boolean functions. In R. Sarker, M. Mohammadian, and X. Yao, editors, *Evolutionary Optimization*, pages 349–369. Kluwer, 2002.
- [23] C. Witt. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation*, 14(1):65–86, 2006.