

# A Flexible Evolutionary Algorithm With Dynamic Mutation Rate Archive

Martin S. Krejca

Laboratoire d'Informatique (LIX), CNRS, École  
Polytechnique, Institut Polytechnique de Paris  
Palaiseau, France  
martin.krejca@polytechnique.edu

Carsten Witt

Technical University of Denmark  
Lyngby, Denmark  
cawi@dtu.dk

## ABSTRACT

We propose a new, flexible approach for dynamically maintaining successful mutation rates in evolutionary algorithms using  $k$ -bit flip mutations. The algorithm adds successful mutation rates to an archive of promising rates that are favored in subsequent steps. Rates expire when their number of unsuccessful trials has exceeded a threshold, while rates currently not present in the archive can enter it in two ways: (i) via user-defined minimum selection probabilities for rates combined with a successful step or (ii) via a stagnation detection mechanism increasing the value for a promising rate after the current bit-flip neighborhood has been explored with high probability. For the minimum selection probabilities, we suggest different options, including heavy-tailed distributions.

We conduct rigorous runtime analysis of the flexible evolutionary algorithm on the ONEMAX and JUMP functions, on general unimodal functions, on minimum spanning trees, and on a class of hurdle-like functions with varying hurdle width that benefit particularly from the archive of promising mutation rates. In all cases, the runtime bounds are close to or even outperform the best known results for both stagnation detection and heavy-tailed mutations.

## CCS CONCEPTS

• **Theory of computation** → **Theory of randomized search heuristics**; • **Mathematics of computing** → **Evolutionary algorithms**.

## KEYWORDS

Theory, runtime analysis, evolutionary algorithm, parameter adaptation, archive

### ACM Reference Format:

Martin S. Krejca and Carsten Witt. 2024. A Flexible Evolutionary Algorithm With Dynamic Mutation Rate Archive. In *Genetic and Evolutionary Computation Conference (GECCO '24)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3638529.3654076>

## 1 INTRODUCTION

The success of evolutionary algorithms (EAs) depends crucially on their parametrization [15]. At their core is the careful balance of exploration and exploitation. On the one hand, an EA should

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '24, July 14–18, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0494-9/24/07.

<https://doi.org/10.1145/3638529.3654076>

find a, potentially locally, optimal solution if it is close by. On the other hand, it should escape local optima and keep finding better solutions within reasonable time.

To achieve this behavior, an EA needs to be flexible enough to consider solutions that are at different distances from its current solutions. As there are many ways to make such a choice, this has led to various approaches that have been studied theoretically [15]. Some approaches consider a static distribution over the search radius (known as *mutation rate*). Others adjust the mutation rate dynamically during the run. Yet other approaches allow to accept worse solutions in order to search more locally for alternative, better solutions. And some approaches do not work with multi-sets of solutions but with distributions over solutions instead. We discuss some prominent ideas in more detail in Section 2.

To the best of our knowledge, almost all of these approaches base their choice for the next mutation rate on ad hoc knowledge, not exploiting the information from previous iterations extensively. This means that recurring structures in the search space are not necessarily well exploited. In contrast, approaches that do attempt to incorporate knowledge from previous iterations are hard to analyze and there are much fewer runtime results.

We propose a flexible EA (the *flex-EA*, Algorithm 1), which aims to satisfy both needs. The flex-EA maintains an archive of mutation rates that were successful in the past, called *active*. In each iteration, it picks a mutation rate based primarily on the active rates, but it still considers all other rates, based on user-defined probabilities. Successful rates are added to the archive while unsuccessful ones are discarded after a certain amount of time, defined by the user. If no rate is active, the flex-EA picks the successor of the last-active rate, reminiscent of stagnation detection [33]. If the archive becomes too full without seeing any progress, it is reset.

The flex-EA allows many of its parameters to be adjusted in order to incorporate problem-specific knowledge. However, in the absence of such knowledge, we recommend to choose the probabilities of inactive rates according to a heavy-tailed distribution – a common, general-purpose choice, first suggested by Doerr et al. [14], also known as a *power-law* – and to choose the times for discarding active rates in the same manner as done in stagnation detection [33]. This results in an essentially parameter-less algorithm.

We show the efficiency of the flex-EA by analyzing it on a diverse set of problems, namely, on unimodal and on jump functions, on the minimum spanning tree problem, and on hurdle-like functions with varying hurdle width. For all of these problems, the performance of the flex-EA is close to or even better than the best known results for stagnation detection and heavy-tailed mutation.

For unimodal functions (Section 5), the progress of the flex-EA is mainly determined by its probability to improve solutions locally

(Theorem 5.1). For ONEMAX and LEADINGONES of problem size  $n$  and with the recommended parametrization, this results in the common bounds of  $O(n \log(n))$  and  $O(n^2)$ , respectively.

For the JUMP benchmark of problem size  $n$  and gap size  $k$  (Section 6), the recommended parametrization results for  $k = o(n/\log n)$  in a bound of  $(2+o(1))\binom{n}{k}$  (Corollary 6.2), which is optimal for unbiased algorithms with unary mutation, up to a factor of  $2+o(1)$  [17]. For many other values of  $k$ , the result is a product of  $\binom{n}{k}$  and a factor that depends at most quadratically on  $k$ .

For the combinatorial problem of finding a minimum spanning tree on a graph with  $m$  edges (Section 7), we rely on the flexibility to incorporate problem-specific knowledge, and we adjust the parameters such that the algorithm mainly makes progress by flipping one or two edges into or out of solutions. With this choice, the flex-EA has an expected runtime of  $(1+o(1))(m^2 \log(m))$  (Theorem 7.1), which matches the best known bound [31].

Last, to showcase the benefit of an archive, we introduce a problem that features many local optima separated by gaps of alternating widths (Section 8). Once the flex-EA determines the necessary rates, it maintains them successfully until the global optimum is found. It does so in an expected time that is faster by a poly-logarithmic factor than heavy-tailed mutation and faster by a super-polynomial factor than stagnation detection (Theorem 8.1). Both, heavy-tailed mutation and stagnation detection are slowed down by requiring to find the correct mutation rate repeatedly, although heavy-tailed mutation does this far more efficiently than stagnation detection.

Overall, our results show that the flex-EA has a great performance with the recommended parameter setting while effectively not requiring any parameter choices. In settings with repeating local structures, this parametrization even outperforms the operators it is based on. In the case of problem-specific knowledge, the flex-EA can achieve results that match the best known bounds. For all of these settings, the state space of the algorithm remains simple enough that a mathematical analysis is still tractable.

## 2 RELATED WORK

There exists a plethora of different operators for choosing good mutation rates [15]. We provide an overview on established concepts that are more closely related to the flex-EA. In particular, we only concern unary operators. Furthermore, we focus on single-trajectory algorithms that only maintain a single search point at a time and update it iteratively. By using populations and diversity-maintaining operators, one opens up for a new level of complexity in the design and analysis of evolutionary algorithms. In addition, this excludes estimation-of-distribution algorithms (EDAs) [22], which maintain a probability distribution over the search space instead of a single search point.

*Non-elitism.* A non-elitist algorithm is one in which worse solutions can be selected over better solutions. Although this approach does not directly change the mutation rate, it can affect the success of the current rate by creating solutions in a distance that is more favorable [4, 12]. This requires a careful balance of the selection probability of solutions. In comparison, the flex-EA is an elitist algorithm, i. e., it aims at finding mutation rates that are good for the currently best solution without modifying it.

*Static rate distributions.* A large class of EAs, most prominently the  $(1+1)$  EA, use a static distribution for selecting mutation rates. Since this distribution is fixed, it needs to accommodate all situations that may arise during optimization. The  $(1+1)$  EA uses a single rate that creates solutions with high probability locally around the current search point. This leads to long waiting times for escaping local optima [21].

Artificial immune systems (AISes) use operators that change solutions more drastically [3]. However, since this makes it harder to find local improvements, AISes often use additional operators that influence how large the changes to a solution actually are [43].

Arguably the best of both worlds is achieved by choosing the mutation rate according to a heavy-tailed distribution, as introduced by Doerr et al. [14] in their seminal paper. This approach, known as fast mutation, finds local improvements quickly and also has a good probability to use larger mutation rates to escape local optima. Since its inception, this approach has seen successful applications in a variety of theoretical works [1–3, 5, 10, 16–19, 29, 42, 44]. A drawback of fast mutation is that it never adjusts its distribution. Thus, to find improving solutions at larger distances, it exclusively relies on the decent probability of choosing a sufficiently large rate. Depending on the distance, this can result in long waiting times.

*Dynamic rate distributions.* Some algorithms change the distribution for choosing the mutation rate during the optimization process [8]. While this technically allows for complex distributions, to the best of our knowledge, most algorithms that fall into this category distribute the probability mass around a single rate.

A common approach is the 1/5 success rule (e. g., [7, 34, 37]), which increases the current rate if more than one fifth of the last tries with the current rate are successful, and it decreases it otherwise. While this approach automatically adjusts the rate to one for finding an improving solution, it forgets the previously used rate and, instead, aims to rediscover it if it becomes useful again later.

A more deterministic approach to a similar idea is stagnation detection [32, 33]. This operator starts with the smallest rate and picks it repeatedly until the algorithm is certain with high probability that it cannot find an improving solution with the current rate. It then increases the rate and repeats the process. This approach has the benefit that it can escape local optima very well. However, similar to the one-fifth success rule, the algorithm focuses on a single rate, meaning that it needs to rediscover previously good rates. This problem was addressed partially by adding a radius memory to the algorithm [31], which saves the last successful rate in order to re-use it. However, if more than a single previous rate are useful, this approach falls back to the problem of the initial algorithm.

Another, very recent approach [17] proposes to combine stagnation detection with fast mutation in the following way: The algorithm follows classic stagnation detection but does not always create a solution with the current mutation rate; it also has a chance of picking a mutation rate from a heavy-tailed distribution around the current rate. This proves very effective in escaping certain local optima where there are many improving solutions at a further distance [17]. The resulting runtime is better than that of classic stagnation detection and thus also than that of the flex-EA. However, we believe that this variant still has problems of re-using previously good rates, as it does not save them.

*Learning-based algorithms.* Doerr et al. [9] propose an algorithm that chooses its mutation rate according to a distribution that is adjusted over time. This adjustment follows the idea of reinforcement learning and incorporates information from many past iterations. A related concept are *hyper-heuristics*, e.g., [25, 26], which are algorithms that aim at selecting in each iteration an algorithm that is best suited for the current state. This choice is based on the previously observed performance of the algorithms to select from.

The flex-EA is similar to these approaches, with the main difference being that it updates its distribution over the mutation rate in a simpler way. This simplification especially allows a rigorous analysis of the flex-EA on a wide variety of problems while still showcasing a strong performance.

### 3 PRELIMINARIES

The natural numbers  $\mathbb{N}$  include 0. For  $a, b \in \mathbb{R}$ , let  $[a..b] = [a, b] \cap \mathbb{N}$ , and let  $[a] = [1..a]$ . By  $\log(\cdot)$  we denote the binary logarithm.

We consider pseudo-Boolean optimization, that is, for a given  $n \in \mathbb{N}_{\geq 1}$ , the optimization of functions  $f: \{0, 1\}^n \rightarrow \mathbb{R}$ . We call  $f$  a *fitness function*, and we assume that  $n$  is given implicitly. When using big-O notation, we assume asymptotics in this implicit  $n$ . Furthermore, we call each  $x \in \{0, 1\}^n$  an *individual*, and  $f(x)$  its *fitness*. For each  $i \in \mathbb{N}$ , let  $x_i$  denote the value of  $x$  at position  $i$ . Last, let  $|x|_1$  be the number of 1s of  $x$ , and  $|x|_0$  its number of 0s.

For each  $x, y \in \{0, 1\}^n$ , let  $d_H(x, y) = |\{i \in \mathbb{N} \mid x_i \neq y_i\}|$  denote the Hamming distance of  $x$  and  $y$ . For each  $r \in [n]$ , we define the *r-bit flip neighborhood* of  $x$  to be the set  $\{y \in \{0, 1\}^n \mid d_H(x, y) = r\}$ .

Given a fitness function  $f$  and an algorithm  $A$  optimizing  $f$ , we define the *runtime* of  $A$  on  $f$  as the (random) number of evaluations of  $f$  until the first time that an optimum of  $f$  is evaluated. To this end, we assume that each individual that  $A$  creates is evaluated exactly once. For the flex-EA in Section 4, the runtime is essentially the number of iterations until an optimum is created for the first time (plus one, due to the initialization).

## 4 THE FLEXIBLE EVOLUTIONARY ALGORITHM

The *flexible evolutionary algorithm* (flex-EA, Algorithm 1) is an elitist  $(1 + 1)$ -type evolutionary algorithm choosing its mutation rate each iteration according to a probability distribution given by a vector (the *frequency vector*). The flex-EA updates its frequency vector in each iteration, aiming to give promising rates a high probability. This update is influenced by different factors. In the following, we first discuss the high-level idea of the flex-EA and its different parts. Afterward, in Section 4.1, we propose good default choices for the various parameters of the algorithm. The result is a basically parameter-less algorithm that can be thought of as a superposition of heavy-tailed mutation [14] and stagnation detection [33].

*Update of the frequency vector.* The update of the vector  $p \in [0, 1]^n$  aims to give (mutation) rates a high probability if choosing them in the past succeeded in improving the fitness of the current solution. If a rate repeatedly fails to be successful, the probability of choosing it (called its *frequency*) is reduced. To this end, the flex-EA groups all rates into those that receive a high frequency (called *active*) and those that receive a low frequency (called *inactive*).

The active frequencies follow a uniform distribution. Each active frequency has a counter of how often it failed to succeed. If this counter exceeds a user-defined limit, the respective frequency becomes inactive. If this results in no frequencies being active and the frequency at index  $r \in [n]$  was the last active one, then the frequency at index  $r + 1$  is set to active (identifying  $n + 1$  with 1).

The update process is superimposed by user-defined lower bounds per frequency. If a frequency would be set to a value less than its lower bound, it takes on its lower bound instead.

In addition, the flex-EA keeps a global counter that counts how many times in a row no success was achieved with any mutation rate. If this counter reaches a specific limit, all frequencies become inactive, and the frequency at position 1 becomes active. This effectively resets the frequency vector.

*The flex-EA in more detail.* Algorithm 1 shows the detailed flex-EA. The frequency vector is denoted by  $p \in [0, 1]^n$ , and the set of active rates (called the *archive*) by  $A \subseteq [n]$ . The lower bounds for each frequency are provided by the user to the flex-EA via the *lower-bound vector*  $\ell \in [0, 1]^n$ . We assume that  $\sum_{i \in [n]} \ell_i \leq 1$ .

Furthermore, each rate  $i \in [n]$  has a counter  $c_i \in \mathbb{N}$  for counting its number of failed attempts to improve the current solution. If  $c_i$  reaches a user-defined limit  $C_i \in \mathbb{R}_{\geq 0}$ , then frequency  $i$  becomes inactive. These limits are provided by the user as the *count bound vector*  $C \in \mathbb{R}_{\geq 0}^n$ .

The global counter is denoted by  $g \in \mathbb{N}$ . Its limit that, if reached, resets the archive of the flex-EA is denoted by  $G \in \mathbb{R}_{\geq 0}$ . The value  $G$  is chosen as the count bound of the smallest active rate  $m \in [n]$  divided by its current frequency. This resembles the normal bound  $C_m$  but accounts for rate  $m$  being only chosen once in  $p_m^{-1}$  iterations in expectation.

The update of the frequency vector is a bit involved, as frequencies may be capped by their respective lower bound, thus disallowing to distribute the probability mass uniformly among all active frequencies. Note that the flex-EA aims to give each active frequency the same probability. The total probability to distribute among these  $|A| = \varphi$  frequencies is  $1 - \sum_{i \in [n] \setminus A} \ell_i =: M$ , as all inactive frequencies are at their respective lower bound. Hence, each active frequency should get a probability of  $\frac{M}{\varphi}$ . In order to account for the lower bounds of the active frequencies, the active frequencies are adjusted sequentially, starting with the largest lower bound, going to the smallest. If the lower bound of the first frequency is less than the intended value  $\frac{M}{\varphi}$ , then all active frequencies get this value. Otherwise, the first frequency is set to its lower bound, and the remaining probability mass is adjusted. This process is repeated.

Algorithm 2 shows this approach algorithmically. We note that the algorithm is presented in a way that aims to easily understand the logic of the update. It is not optimized for efficiency. A more efficient implementation would make use of a data structure for the archive that quickly allows to add and delete indices while keeping them ordered with respect to their lower bound. This can be achieved, for example, with a balanced search tree.

### 4.1 Recommended Parameter Choices

The flex-EA requires the user to provide  $2n$  parameters, which can be daunting. Hence, we propose a recommendation that works well in many cases, as we prove in most of the following sections. For

---

**Algorithm 1:** The flexible evolutionary algorithm (flex-EA) with given lower-bound vector  $\ell$  and count bound vector  $C$ , maximizing a pseudo-Boolean function  $f: \{0, 1\}^n \rightarrow \mathbb{R}$ , stopping at a user-defined termination criterion.

---

```

1  $t \leftarrow 0$ ;
2  $g^{(t)} \leftarrow 0$ ;
3  $A^{(t)} \leftarrow \{1\}$ ;
4 for  $i \in [n]$  do  $c_i^{(t)} \leftarrow 0$ ;
5  $\mathbf{x}^{(t)} \leftarrow$  solution from  $\{0, 1\}^n$  chosen uniformly at random;
6 while termination criterion not met do
7    $\mathbf{p}^{(t)} \leftarrow$  distribute probability mass over  $A^{(t)}$ ,
   respecting  $\ell$  (Algorithm 2);
8    $m \leftarrow \min A^{(t)}$ ;
9    $G^{(t)} \leftarrow \frac{C_m}{p_m^{(t)}}$ ;
10  choose  $r \in [n]$  according to  $\mathbf{p}^{(t)}$ ;
11   $\mathbf{y}^{(t)} \leftarrow$  copy  $\mathbf{x}^{(t)}$  and flip  $r$  bits uniformly at random in
   this copy;
12  if  $f(\mathbf{y}^{(t)}) > f(\mathbf{x}^{(t)})$  then
13     $\mathbf{x}^{(t+1)} \leftarrow \mathbf{y}^{(t)}$ ;
14     $A^{(t+1)} \leftarrow A^{(t)} \cup \{r\}$ ;
15     $g^{(t+1)} \leftarrow 0$ ;
16     $c_r^{(t+1)} \leftarrow 0$ ;
17  else
18    if  $f(\mathbf{y}^{(t)}) = f(\mathbf{x}^{(t)})$  then  $\mathbf{x}^{(t+1)} \leftarrow \mathbf{y}^{(t)}$ ;
19     $g^{(t+1)} \leftarrow g^{(t)} + 1$ ;
20     $c_r^{(t+1)} \leftarrow c_r^{(t)} + 1$ ;
21    if  $g^{(t+1)} \geq G^{(t)}$  then
22       $g^{(t+1)} \leftarrow 0$ ;
23       $A^{(t+1)} \leftarrow \{1\}$ ;
24       $c_1^{(t+1)} \leftarrow 0$ ;
25    else if  $c_r^{(t+1)} \geq C_r$  then
26       $A^{(t+1)} \leftarrow A^{(t)} \setminus \{r\}$ ;
27      if  $A^{(t+1)} = \emptyset$  then
28         $r' \leftarrow$  if  $r + 1 \leq n$  then  $r + 1$  else  $1$ ;
29         $A^{(t+1)} \leftarrow A^{(t+1)} \cup \{r'\}$ ;
30         $c_{r'}^{(t+1)} \leftarrow 0$ ;
31   $t \leftarrow t + 1$ ;

```

---

the lower-bounds vector  $\ell$ , we recommend to choose a heavy-tailed distribution, inspired by the fast-mutation operator [14]. For the count bound vector  $C$ , we recommend to choose the same values as for stagnation detection [33]. We explain both recommendations in more detail below and then briefly discuss the resulting algorithm.

*Heavy-tailed lower bounds.* Let  $\beta \in (1, 2)$  (called the *power-law exponent*), and let  $N: (1, 2) \rightarrow \mathbb{R}_{\geq 1}$  denote a normalization with  $N: \alpha \mapsto \sum_{i \in [n]} i^{-\alpha}$ . We recommend for all  $i \in [n]$  to choose

$$\ell_i = \frac{1}{2N(\beta)} i^{-\beta}.$$

---

**Algorithm 2:** The algorithm that, given an index set  $A \subseteq [n]$  and a lower-bound vector  $\ell$  of size  $n$ , returns a frequency vector  $\mathbf{p}$  of size  $n$  where all possible probability mass is distributed as evenly as possible over the frequencies with indices in  $A$ , respecting  $\ell$ .

---

```

1  $M \leftarrow 1 - \sum_{i \in [n] \setminus A} \ell_i$ ;
2 for  $i \in [n] \setminus A$  do  $p_i \leftarrow \ell_i$ ;
3  $\varphi \leftarrow |A|$ ;
4  $(S_i)_{i \in [\varphi]} \leftarrow A$  sorted in descending order by  $\ell$ ;
5 for  $i \in [\varphi]$  do
6   if  $\ell_{S_i} \leq \frac{M}{\varphi - i + 1}$  then
7     for  $j \in [i.. \varphi]$  do  $p_{S_j} \leftarrow \frac{M}{\varphi - i + 1}$ ;
8     return  $\mathbf{p}$ ;
9   else
10     $p_{S_i} \leftarrow \ell_{S_i}$ ;
11     $M \leftarrow M - \ell_{S_i}$ ;
12 return  $\mathbf{p}$ ;

```

---

Note that  $\sum_{i \in [n]} \ell_i = \frac{1}{2}$ . Hence, the flex-EA still has a probability mass of  $\frac{1}{2}$  remaining to distribute among the active frequencies.

This recommendation follows the ideas of *fast mutation*, introduced in [14]. This mutation operator first chooses a rate  $r \in [n]$  according to a heavy-tailed distribution and performs standard bit mutation afterward. That is, it flips each bit of a given individual independently with probability  $\frac{r}{n}$ . Doerr et al. [14] name the variant of the  $(1+1)$  EA that uses fast mutation the Fast  $(1+1)$  EA.

Doerr et al. [14] choose the heavy-tailed distribution because even large rates have a reasonable probability of at least  $\frac{1}{n^2}$  of being chosen. The authors show that the Fast  $(1+1)$  EA has a super-exponential runtime speed-up in the parameter  $k$  compared to the classic  $(1+1)$  EA when optimizing  $\text{JUMP}_k$  and has seen great success in other settings (see also Section 2). Hence, we propose to also use it for the flex-EA.

*Stagnation detection bounds.* Let  $R \in \mathbb{R}_{>1}$ . We recommend

$$C_i = \binom{n}{i} \ln(R)$$

for all  $i \in [n]$  and call this the *standard-SD choice*.

This recommendation follows the ideas of *stagnation detection* (SD), introduced by Rajabi and Witt [33]. In their paper, the authors introduce the algorithm  $\text{SD-RLS}^r$ , which is also a  $(1+1)$ -type elitist algorithm that performs the same mutation as the flex-EA but chooses its mutation rate differently in the following deterministic manner. The  $\text{SD-RLS}^r$  starts with rate 1 and counts how many times it used this rate. It then uses rate 1 for at most  $C_1$  iterations. If it is successful in finding a strictly improving solution during this time, the algorithm resets its counter. If rate 1 is not successful, the  $\text{SD-RLS}^r$  chooses rate 2 from now on, for a total of at most  $C_2$  iterations. If it is successful during this time, it resets its counter and reverts back to rate 1. Otherwise, it picks the next-higher rate and adjusts the iteration limit according to the formula above. We note that the  $\text{SD-RLS}^r$  only tries out rates up to  $\frac{n}{2}$ , but the formula for the bounds generalizes to all values in  $[n]$ .

Rajabi and Witt [33] choose the bounds above because this guarantees that if rate  $r \in [n]$  is currently used and there is at least one solution in Hamming distance  $r$  of the currently best solution, the SD-RLS<sup>f</sup> find this improvement with probability at least  $1 - R^{-1}$ , as the probability of not doing so is at most  $(1 - \binom{n}{r}^{-1})^{C_r} \leq R^{-1}$ . Hence, the value  $R$  determines how unlikely it is to fail in such a case. Since the flex-EA is in some aspects similar to the SD-RLS<sup>f</sup>, we recommend to copy its bounds.

*Discussion.* When considering the flex-EA with heavy-tailed lower bounds and with the standard-SD choice, it has two parameter values left to choose, namely, the power-law exponent  $\beta \in (1, 2)$  as well as the parameter  $R \in \mathbb{R}_{>1}$ . However, the exact choice of these parameters is typically not very relevant, as we show in this article. Doerr et al. [14] recommend to choose  $\beta = \frac{3}{2}$ . Since the flex-EA can act similarly to the algorithm discussed by Doerr et al. [14], we recommend this choice too. For the parameter  $R$ , Rajabi and Witt [33] show that a choice of  $R \geq n^{4+\epsilon}$  for a small constant  $\epsilon > 0$ , with  $R$  being a polynomial, is usually sufficient. We show in our results below that this is also the case for the flex-EA.

Since we provide recommendations for the only two parameter choices left, we consider the flex-EA to be an essentially parameter-less algorithm. However, one needs to keep in mind that the flex-EA allows to choose its parameters  $\ell$  and  $C$  more flexibly, which can be a good choice if given extra information (see also Section 7).

## 5 RUNTIME ON UNIMODAL FUNCTIONS

We analyze the flex-EA on unimodal pseudo-Boolean functions, that is, functions  $f: \{0, 1\}^n \rightarrow \mathbb{R}$  with a unique global optimum  $z^* \in \{0, 1\}^n$  such that, for all  $x \in \{0, 1\}^n \setminus \{z^*\}$ , there is a  $y \in \{0, 1\}^n$  with  $d_H(x, y) = 1$  such that  $f(x) < f(y)$ . For each unimodal function  $f$ , we call the cardinality of the range of  $f$  the number of *fitness levels* of  $f$ , and each set of all individuals of the same fitness a *fitness level*. Each individual (besides the global optimum) has at least one neighbor in a fitness level of better fitness.

The flex-EA optimizes unimodal functions efficiently in expectation if its parameters are chosen well. To this end, we rely exclusively on the lower bound of rate 1.

**THEOREM 5.1.** *Let  $f$  be a unimodal fitness function with  $L$  fitness levels. For all  $j \in [L-1]$ , let  $q_j$  denote a lower bound on the probability that, given an individual  $z$  in fitness level  $j$ , flipping a single bit in  $z$  creates an offspring with strictly better fitness. Furthermore, consider the flex-EA with  $\ell_1 > 0$ .*

*Then the expected runtime of the flex-EA on  $f$  is at most*

$$1 + \ell_1^{-1} \sum_{j \in [L-1]} q_j^{-1}.$$

*Especially, it is at most  $1 + n(L-1)\ell_1^{-1}$ .*

**PROOF.** We start with the first claim. The plus 1 is for initialization. For the remaining bound, we apply the fitness level method (Theorem A.1). Let  $j \in [L-1]$  be the fitness level of the current solution of the flex-EA. In order to leave level  $j$ , it is sufficient to choose rate 1 and then create an offspring with strictly better fitness. The probability of the former is  $\ell_1$ , and the one of the latter is at least  $q_j$ . Since these events are independent, the flex-EA leaves level  $j$  with a probability of at least  $\ell_1 q_j$ . This concludes this case.

For the second claim, note that since  $f$  is unimodal, there is always at least one improving solution that only requires a single bit to be flipped. Hence, it holds that  $q_j \geq \frac{1}{n}$ . Applying the first claim concludes the proof.  $\square$

**Theorem 5.1** leads to useful bounds on various functions. For example, the coarser bound at the end results already in the common expected runtime of  $O(n^2)$  for the LEADINGONES benchmark [36] if we assume a constant lower bound for rate 1, as is the case with heavy-tailed lower bounds. For functions where the probability of leaving a fitness level depends on the current state, we need to estimate these transition probabilities more carefully. Doing so leads to an expected runtime bound of  $O(n \log n)$  for the ONEMAX benchmark, as we show below (Corollary 5.2). Since we require bounds for similar functions in the following sections, we introduce a slight generalization of ONEMAX and analyze it instead.

For all  $k \in [0..n]$ , let *trimmed ONEMAX* ( $OM_k$ ) be defined as

$$x \mapsto \begin{cases} k + OM(x) & \text{if } |x|_1 \leq n - k, \\ n - OM(x) & \text{else.} \end{cases}$$

Note that  $OM_k$  is a unimodal function with  $n + 1$  fitness levels. Its global optimum is achieved for all individuals with exactly  $n - k$  1s. The case  $k = 0$  results in ONEMAX. We get the following bound.

**COROLLARY 5.2.** *Let  $k \in [0..n]$ . Consider the flex-EA with  $\ell_1 > 0$  optimizing  $OM_k$ . Then the expected runtime is  $O(\ell_1^{-1} n \log(n))$ .*

**PROOF.** We apply Theorem 5.1 and use its notation. For all  $j \in [n + 1]$ , denote the fitness level that represents an  $OM_k$ -value of  $j - 1$ . We say that individuals with at most  $n - k$  1s are on the *left branch* of the function, and the other ones on the *right branch*.

Let  $j \in [n]$ . We bound  $q_j$  with respect to whether the current individual is on the left or the right branch. If it is on the left branch, it has  $j - 1 - k$  1s. If mutation flips any of the  $n - j + 1 + k$  0s, then the result strictly improves the fitness. Hence,  $q_j = \frac{n-j+1+k}{n}$ .

If the current individual is on the right branch, it has  $n - j + 1$  1s. If mutation flips any of these 1s, then the result strictly improves the fitness. Hence,  $q_j = \frac{n-j+1}{n}$ .

Combining both cases, by Theorem 5.1 and disregarding the plus 1, we obtain an upper bound on the expected runtime of

$$\begin{aligned} \ell_1^{-1} \sum_{j \in [n]} \left( \frac{n}{n-j+1+k} \cdot \mathbb{1}\{j \geq k+1\} + \frac{n}{n-j+1} \cdot \mathbb{1}\{j < k+1\} \right) \\ \leq \ell_1^{-1} \cdot 2n \sum_{j \in [n]} j^{-1} = O(\ell_1^{-1} n \log(n)). \end{aligned} \quad \square$$

## 6 RUNTIME ON JUMP FUNCTIONS

We consider the well known JUMP benchmark, which, for all  $k \in [n]$  is defined as

$$x \mapsto \begin{cases} k + OM(x) & \text{if } |x|_1 \in [n-k] \cup \{n\}, \\ n - OM(x) & \text{else.} \end{cases}$$

This function is identical to  $OM_k$  except that its unique global optimum is at the all-1s bit string. Consequently, the individuals with exactly  $n - k$  1s are all local optima. The set of all local optima is often called the *plateau* of the function. For elitist algorithms, in order to leave the plateau, exactly all  $k$  0s of an individual need

to be changed into a 1, which requires, in expectation, at least  $\binom{n}{k}$  tries for unbiased algorithms with unary mutation [17].

With [Corollary 6.2](#), we show that the flex-EA with the recommended parameter choices optimizes JUMP in this time. Beforehand, we prove a runtime bound for a more general parametrization.

**THEOREM 6.1.** *Let  $R \in \mathbb{R}_{>1}$  and  $k \in [2..n]$ . Consider the flex-EA with the standard-SD choice  $C$  with parameter  $R$  and with lower bounds  $\ell$  such that  $\ell_1, \ell_k > 0$  and  $1 - \sum_{i \in [n] \setminus \{k\}} \ell_i = \Theta(1)$ . Furthermore, consider that the flex-EA optimizes JUMP $_k$ . Let  $L = (1 - \sum_{i \in [n] \setminus \{k\}} \ell_i)^{-1}$ . Then the expected runtime is*

$$\begin{aligned} & O\left(n\ell_1^{-1} \log(nR)\right) + \min \left\{ \binom{n}{k} \cdot \ell_k^{-1}, \right. \\ & \binom{n}{k} \cdot \left( L + O\left(\frac{k}{n-2k+3} \log(R)\right) + \ell_k^{-1} R^{-1/L} \right) \\ & \left. + 2^n \cdot \mathbb{1}\left\{k \geq \frac{n}{2}\right\} \ln(R) \right\}. \end{aligned}$$

**PROOF.** Let  $T$  denote the runtime of the algorithm on JUMP $_k$ . We split  $T$  into the following two phases: Phase 1 considers the time until the current solution of the algorithm has a fitness of at least  $n$ , that is, the current solution is on the plateau or optimal. Phase 2 starts in the iteration that phase 1 stops and considers the time until the current solution is optimal. Let  $T_1$  and  $T_2$  denote the runtimes of the respective phases. By the linearity of expectation and by  $T = T_1 + T_2$ , it holds that  $E[T] = E[T_1] + E[T_2]$ . We consider the expectations of both phases separately.

**Phase 1.** Note that the time until the algorithm finds an optimum of OM $_k$  is an upper bound for  $T_1$ . Hence, the same is true for their expectations. By [Corollary 5.2](#), we obtain  $E[T_1] = O(\ell_1^{-1} n \log(n))$ .

**Phase 2.** Note that before the global optimum is found, no rate enters the archive due to being successful, since the global optimum is the only remaining strictly better solution. We bound the expected time of this phase in two different ways. The first way relies exclusively on the lower bound  $\ell_k$  for sampling the global optimum, and the second way considers a course of events similar to stagnation detection. It assumes that rate 1 is active in iteration  $T_1$ , that is, at the beginning of phase 2. Calling the respective bounds  $B_1$  and  $B_2$ , we then conclude that  $E[T_2] \leq \min\{B_1, B_2\}$ , as either bound is valid.

**1. Relying only on the lower bound.** In each iteration, the probability for choosing rate  $k$  is at least  $\ell_k$ . Then, since the current individual has exactly  $k$  0s, the probability to mutate it into the global optimum is  $\binom{n}{k}^{-1}$ . Hence, in each iteration, the probability to create the global optimum is at least  $\ell_k \binom{n}{k}^{-1} =: q$ . Since each iteration has an independent chance to create the global optimum,  $T_2$  follows a geometric distribution with a success probability of at least  $q$ . Hence,  $E[T_2] \leq q^{-1} = \ell_k^{-1} \binom{n}{k}$ .

**2. Rate 1 is active at the start.** Let  $D$  denote the event that  $1 \in A^{(T_1)}$ , and assume that  $D$  occurs. Then  $G^{(T_1)} = C_1/p_1^{(T_1)} \leq C_1/\ell_1$ . Thus, if no strict improvement is found within at most  $C_1/\ell_1$  iterations, the algorithm reverts to classic stagnation detection. Since the current individual is on the plateau, finding a strict improvement means to find the global optimum (and thus ending the phase).

Let  $T_{2,1}$  denote the first time in phase 2 such that  $g^{(T_{2,1}+T_1)} \geq C_1/\ell_1$ , and let  $T_{2,2} = T_2 - T_{2,1}$ . Since no mutation rate but  $k$  can lead to creating the global optimum, the global counter increases at most  $C_1/\ell_1$  times. Hence,  $T_{2,1} \leq C_1/\ell_1$ , and thus  $E[T_{2,1} | D] \leq C_1/\ell_1$ .

After  $T_{2,1}$ , if the algorithm did not find the optimum yet, it performs classic stagnation detection by maintaining a single active rate, based on  $C$ . Since all rates in  $[k-1]$  cannot lead to an improvement, a total of  $\sum_{i \in [k-1]} \binom{n}{i} \ln(R)$  iterations is spent until rate  $k$  becomes active. For  $k \leq \frac{n}{2}$ , we obtain by [Lemma A.2](#)

$$\begin{aligned} \sum_{i \in [k-1]} \binom{n}{i} \ln(R) & \leq \frac{n - (k-2)}{n - (2k-3)} \binom{n}{k-1} \ln(R) \\ & = \frac{n - (k-2)}{n - (2k-3)} \frac{k}{n-k+1} \binom{n}{k} \ln(R) \\ & = O\left(\frac{k}{n-2k+3}\right) \binom{n}{k} \ln(R). \end{aligned}$$

For  $k > \frac{n}{2}$ , we obtain  $\sum_{i \in [k-1]} \binom{n}{i} \ln(R) \leq 2^n \ln(R)$ .

Overall, for  $\gamma := O\left(\frac{k}{n-2k+3}\right) \binom{n}{k} + 2^n \cdot \mathbb{1}\{k > \frac{n}{2}\}$ , it takes at most  $\gamma \ln(R)$  iterations until rate  $k$  is active or the optimum is found.

Assume that the optimum is not found until rate  $k$  is active. Let  $S$  denote the event that the algorithm samples the optimum within the next  $C_k \ln(R)$  iterations. The probability (unconditional on  $S$ ) to create the optimum in a single iteration is  $(1 - \sum_{i \in [n] \setminus \{k\}} \ell_i) \cdot \binom{n}{k}^{-1} = L^{-1} \binom{n}{k}^{-1}$ . The process of creating the optimum, conditional on  $S$ , follows a truncated geometric distribution [33] and is dominated by a geometric distribution with success probability  $L^{-1} \binom{n}{k}^{-1}$ , since the actual process stops after at most  $C_k \ln(R)$  attempts. Hence,  $E[T_{2,2} | D, S] \leq \gamma \ln(R) + L \binom{n}{k}$ .

In the case that  $S$  does not hold, we use the same arguments as when relying only on  $\ell_k$  and obtain  $E[T_{2,2} | D, \bar{S}] \leq \gamma \ln(R) + \ell_k^{-1} \binom{n}{k}$ .

Next, we trivially bound  $\Pr[D, S] \leq 1$  as well as  $\Pr[D, \bar{S}] \leq \Pr[\bar{S} | D]$ . By the definition of  $S$ , the estimates above of finding the global optimum in a single iteration, as well as that  $\binom{n}{k} = C_k$  due to the standard-SD choice,  $\Pr[\bar{S} | D] \leq (1 - L^{-1} C_k^{-1})^{C_k \ln(R)} \leq R^{-1/L}$ . Thus, we obtain

$$E[T_2 | D] \leq \frac{C_1}{\ell_1} + \gamma \ln(R) + L \binom{n}{k} + \ell_k^{-1} \binom{n}{k} R^{-1/L}.$$

We conclude by trivially bounding  $\Pr[D] \leq 1$  and by adding the term  $C_1/\ell_1 = n\ell_1^{-1} \ln(R)$  also to our first bound for phase 2.  $\square$

[Theorem 6.1](#) shows the intricate interplay of the parameters of the flex-EA. For  $k \geq \frac{n}{2}$ , the value of the lower bound for rate  $k$  has a huge impact on the expected runtime, as it is better to choose rate  $k$  due to its lower bound than to wait for it to become the solely active rate. The latter requires to wait in the order of at least the central binomial coefficient, which is exponential in  $n$ .

In contrast, if  $k < \frac{n}{2}$  is sufficiently small, the value  $L$  of all probability mass put on rate  $k$ , except for the lower bounds, plays an important role, as it can speed up the process of choosing rate  $k$ . If the optimum is not found when rate  $k$  is active, the algorithm uses the lower bound of rate  $k$  again. However, choosing the parameter  $R$  sufficiently large makes it unlikely for this case to occur.

The following corollary summarizes how the expected runtime improves with the recommended parameter choices. Up to a factor

of  $2 + o(1)$ , the result is asymptotically optimal if  $k$  is sufficiently far away from  $\frac{n}{2}$ . The factor of 2 is a result of the lower bounds, which take up a total probability mass of  $\frac{1}{2}$  in the recommended setting. Hence, even if rate  $k$  is the only active rate, the probability of choosing it is only about  $\frac{1}{2}$ .

**COROLLARY 6.2.** *Let  $\beta \in (1, 2)$  and  $\varepsilon \in \mathbb{R}_{>0}$  be constants, let  $R = n^{2(\beta+\varepsilon)}$ , and let  $k \in [2..n]$ . Consider the flex-EA with the standard-SD choice  $C$  with parameter  $R$  and with heavy-tailed lower bounds  $\ell$ . Furthermore, consider that the flex-EA optimizes  $\mathcal{J}_{\text{MST}_k}$ . Then the expected runtime is*

$$(1 + o(1)) \min \left\{ \binom{n}{k} \cdot 2N(\beta)k^\beta, \binom{n}{k} \cdot \left( 2 + O\left(\frac{k}{n-2k+3} \log(n) + n^{-\varepsilon}\right) \right) + 2^n \cdot \mathbb{1}\left\{k \geq \frac{n}{2}\right\} \ln(R) \right\}.$$

*Especially, for  $k = o\left(\frac{n}{\log(n)}\right)$ , this simplifies to  $(2 + o(1))\binom{n}{k}$ .*

**PROOF.** The first bound follows from [Theorem 6.1](#). Due to the assumptions on  $\ell$ , it holds that  $\ell_1^{-1} = \Theta(1)$ , that  $\ell_k = \frac{1}{2N(\beta)}k^{-\beta}$ , and that  $1 - \sum_{i \in [n] \setminus \{k\}} \ell_i \geq \frac{1}{2}$ . Furthermore, due to the choice of  $R$ , it follows that  $O(n\ell_1^{-1} \log(nR)) = O(n \log(n))$ . Since  $k \geq 2$ , the other summand is  $\Omega(n^2)$ , so the summand  $O(n\ell_1^{-1} \log(nR))$  is of lower order. Last, since  $L^{-1} \geq \frac{1}{2}$ , it follows that  $R^{-1/L} \leq R^{-1/2} \leq n^{-(\beta+\varepsilon)}$ . Substituting all of the values yields the desired result.

For the case  $k = o\left(\frac{n}{\log(n)}\right)$ , note that  $\frac{k}{n-2k+3} \log(n) = o(1)$ . Hence, the second term in the minimum is  $(2 + o(1))\binom{n}{k}$ . This term is smaller than the first term in the minimum, since  $N(\beta) \geq 1$  as well as  $k^\beta \geq 2$ , concluding the proof.  $\square$

## 7 BENEFITS OF THE RATE ARCHIVE ON MINIMUM SPANNING TREES

The prior works on SD-RLS<sup>t</sup> (i. e., RLS with a classical stagnation detection mechanism) in [\[33\]](#) and on SD-RLS<sup>m</sup>, a variant with a radius memory, in [\[31\]](#) include analyses of the algorithms on the classical minimum spanning tree problem. Interestingly, both stagnation detection algorithms find optimal solutions to the problem in efficient polynomial expected time (see detailed expressions below), which stands in contrast to the bound  $O(m^2(\log n + \log w_{\max}))$ , where  $n$  is the number of vertices,  $m$  the number of edges and  $w_{\max}$  is the largest edge weight, from the seminal work [\[27\]](#) analyzing the (1+1) EA on the MST problem. It is well known [\[35\]](#) that the dependency of the bound on  $w_{\max}$  can be removed and that the runtime bound  $O(m^3 \ln n)$  is obtained by studying the (1+1) EA on a rank-equivalent fitness function, however, even that bound is probably far from tight. Before that, polynomial bounds were only known for randomized local search algorithms with 1- and 2-bit flip neighborhoods, but not for globally searching algorithms.

In this section, we will analyze the flex-EA on the MST problem and prove results matching the best classical stagnation detection algorithm. The setting is as follows. Given is an undirected, weighted graph  $G = (V, E)$ , where  $n = |V|$ ,  $m = |E|$  and the weight of edge  $e_i$ , where  $i \in \{1, \dots, m\}$ , is a positive integer  $w_i$ . Let  $c(x)$  denote the number of connected components in the subgraph described by

the search point  $x \in \{0, 1\}^m$ . The fitness function  $f: \{0, 1\}^m \rightarrow \mathbb{R}$  considered in [\[27\]](#), to be minimized, is defined by

$$f(x) := M^2(c(x) - 1) + M\left(\sum_{i \in [m]} x_i - (n - 1)\right) + \sum_{i \in [m]} w_i x_i$$

for an integer  $M \geq n^2 w_{\max}$ , where  $w_{\max}$  denotes the largest edge weight. Hence,  $f$  returns the total weight of a given spanning tree and penalizes unconnected graphs as well as graphs containing cycles so that such graphs are always inferior than spanning trees. The authors of [\[33\]](#) showed that SD-RLS<sup>t</sup> starting with an arbitrary spanning tree finds an MST in  $(1 + o(1))(m^2 \ln m + (4m \ln m)E[S])$  fitness calls where  $E[S]$  is the expected number of strict improvements. The variant SD-RLS<sup>m</sup> with radius memory analyzed in [\[31\]](#) has an expected runtime of at most

$$(1 + o(1))((m^2/2)(1 + \ln(\sum_{i \in [m]} r_i))) \leq (1 + o(1))(m^2 \ln m),$$

where  $r_i$  is the rank of the  $i$ th edge in the sequence sorted by increasing edge weights. This holds since SD-RLS<sup>m</sup> for a sufficiently long time only flips two bits uniformly at random; then its stochastic search trajectory on the given MST instance is indistinguishable from an instance, where the  $i$ -largest weight from the original set  $w_1, \dots, w_m$  is replaced by the rank number  $r_i$  [\[30\]](#). Essentially, the bound for SD-RLS<sup>m</sup> is better by a factor of 2 than for the classical RLS<sup>1,2</sup> algorithm choosing uniformly at random a rate  $r \in \{1, 2\}$  and flipping  $r$  bits uniformly. In fact, RLS<sup>1,2</sup> wastes every second step since 1-bit flips are never accepted on spanning trees.

We obtain [Theorem 7.1](#) below. Choosing  $\ell_1 = o(1)$ , it matches the bound from [\[31\]](#). While it does not improve the bound, which seems to be the best obtainable with the present methods, it shows the robustness of the flex-EA and features a simpler analysis than in [\[31\]](#). Finally, it assumes a uniform starting point instead of an initialization with a spanning tree as in [\[31\]](#).

We note that the parameter settings for the  $C_i$  and particularly the  $\ell_i$  in [Theorem 7.1](#) deviate from the recommended settings discussed in [Section 4.1](#) and used earlier in this paper. In fact, we take a gray-box perspective here and use a minor amount of problem-specific knowledge to favor the small rates 1 and 2. If the expected value was to hold conditioned on an event of probability  $1 - o(1)$  only, then the conditions on  $\ell_i$  for  $i \geq 3$  could be relaxed.

**THEOREM 7.1.** *Consider an instance to the MST problem, modeled with the classical fitness function from [\[27\]](#). Let flex-EA with the following parameter settings run on this function:*

- $C_i = m^i \ln(m^c)$  for a sufficiently large constant  $c > 0$  and  $i = 1, 2$  (the other  $C_i$  may be chosen arbitrarily),
- $\ell_1 = \omega(1/\ln m)$  and  $\ell_1 \leq 1/2$
- $\ell_2 = \Omega(1/(m^2 \ln m))$ , and  $\ell_i = o(1/(m^5 \ln^2 m))$  for  $i = 3..n$ .

*Then with probability  $1 - o(1)$ , the runtime is at most*

$$\begin{aligned} & ((1 - \ell_1)^{-1} + o(1))((m^2/2)(1 + \ln(\sum_{i \in [m]} r_i))) \\ & \leq ((1 - \ell_1)^{-1} + o(1))(m^2 \ln m), \end{aligned}$$

*where  $r_i$  is the rank of the  $i$ th edge in the sequence sorted by increasing edge weights. The expected runtime is bounded in the same way.*

The proof had to be omitted from this paper. Its main idea is to consider a phase of length  $\Theta((1 - \ell_1)^{-1} m^2 \ln m)$  after finding the first spanning tree and to show that within the phase, the archive

only contains rates 1 and 2 w. h. p. Then the above-mentioned stochastic equivalence to RLS with a 2-bit flip neighborhood is applied.

## 8 A FUNCTION WHERE THE ARCHIVE STORES SEVERAL RATES SIMULTANEOUSLY

Most of the above results do not exploit the full power of the archive  $A$ ; most of the time, it was sufficient that the archive contained only one rate and larger archive sizes were even detrimental for the analysis. A major difference of the flex-EA to classical stagnation detection is that a successful rate is kept for the next iteration; however, to some extent this idea was already present in stagnation detection with radius memory [31], which we will compare against.

The purpose of this section is to present an example where flex-EA needs to focus on two rates simultaneously. This can be beneficial on a multimodal function where fitness gaps of different sizes occur frequently. Our example resembles a HURDLE function [28] and features two different gap sizes, so that the archive should include the two corresponding rates. It seems possible to extend the example to a larger number of gap sizes.

Assume the numbers  $s := (3/4)n - \sqrt{n}$ ,  $(3/4)n$ ,  $g := \log n$  and  $\sqrt{n}/g$  are integers. We define

$$\text{TWO RATES}(x) := \begin{cases} \text{OM}(x) & \text{if } |x| < s \text{ or } |x| \geq (3/4)n \\ & \text{or } |x| = s + ig \text{ for } i \in \{0, 2, 3, 5, 6, \dots\}, \\ -1 & \text{otherwise,} \end{cases}$$

where the values of  $i$  alternately increase by 2 and 1.

The hurdles of width  $g$ , which are located in the interval  $[s..3n/4]$  for the number of one-bits, are overcome at the corresponding rate with probability at least

$$\frac{\binom{n/4}{g}}{\binom{n}{g}} \geq \left(\frac{n/4 - g}{n}\right)^g = \left(1 - \frac{O(\log n)}{n}\right)^{\log n} 4^{-\log n} = \frac{(1 - o(1))}{n^2}$$

(using the inequalities  $\frac{\binom{n-k}{k}}{k!} \leq \frac{\binom{n}{k}}{k!} \leq \frac{n^k}{k!}$ ) and at most

$$\frac{\binom{n/4 + \sqrt{n}}{g}}{\binom{n}{g}} \leq \left(\frac{n/4 + \sqrt{n}}{n - g}\right)^g = 2^{\log(1/4 + O(1/\sqrt{n}))(\log n)} \leq \frac{(1 + o(1))}{n^2} \quad (1)$$

and accordingly those of width  $2g$  with the probability bounds  $(1 - o(1))/n^4$  and  $(1 + o(1))/n^4$ , respectively.

We now prove that the flex-EA optimizes TWO RATES efficiently, while classical stagnation detection algorithms with and without radius memory need superpolynomial time. Moreover, the Fast (1+1) EA from [14] is by a at least a polylogarithmic factor of  $\Omega((\log n)^{\beta-1/2}) = \omega(\sqrt{\log n})$  slower than our approach.

**THEOREM 8.1.** *The expected runtime of the flex-EA with the recommended parameter set on TWO RATES is bounded by  $O(n^{4.5})$ . For the stagnation detection algorithms of the type SD-RLS in [31, 33] with typical parameter choice  $R = n^d$  for a constant  $d > 0$ , the expected runtime is  $n^{\Omega(\log n)}$ . Moreover, the expected runtime of the Fast (1+1) EA is by a factor  $\Omega((\log n)^{\beta-1/2})$  larger than for the flex-EA.*

In the proof, we shall need the following useful helper result. It relates to the fact that flipping exactly  $g$  zero-bits and no one-bits may not be optimal to cross a hurdle of size  $g$ , and it bounds the increase in probability when flipping both zero- and one-bits.

**LEMMA 8.2.** *Let  $x \in \{0, 1\}^n$  be a search point with  $a \geq 0$  one-bits and let  $p(n, a, d, i)$ , where  $d, i \geq 0$ , be the probability of mutating  $x$  into a point containing exactly  $a + d$  one-bits when flipping  $d + 2i$  bits uniformly at random. If  $d + 2i \leq n^{1/3}$ , then  $p(n, a, d, i)/p(n, a, d, 0) \leq (1 + o(1))(e(2 + d/i)(1 - a/n))^i \leq (1 + o(1))e^{d+2i}(1 - a/n)^i$ .*

The main idea for the upper bound for the flex-EA is as follows: the expected time to cross a hurdle is  $O(n^4)$  and there are  $\Theta(\sqrt{n}/\log n)$  hurdles. The probability of choosing a beneficial rate of  $g$  or  $2g$  is  $\Omega(1/\log n)$  since the cardinality of the archive is bounded by  $O(\log n)$  with high probability. The lower bound for the Fast (1+1) EA stems from the fact that rates of  $\Omega(\log n)$  must be chosen with high probability to cross a hurdle, which has probability  $O((\log n)^{-\beta})$  in the heavy-tailed mutation. Another factor of  $\Theta(\sqrt{\log n})$  is lost since the Fast (1+1) EA uses standard-bit mutation instead of a flipping a deterministic number of bits.

We conjecture that the runtime of the SD-FEA from [17] on TWO RATES is asymptotically at least as large as the one of the Fast (1+1) EA. Essentially, rates in the order  $\Omega(\log n)$  must be used to obtain a sufficient probability of crossing a hurdle. Then in the same way as described above, the selection probability  $O((\log n)^{-\beta})$  in the heavy-tailed component impacts the runtime bound.

## 9 CONCLUSION

We introduced the flex-EA – a flexible evolutionary algorithm that maintains an archive of previously successful mutation rates. The algorithm is highly configurable and thus allows to exploit problem-specific knowledge. However, we also recommended a default configuration, based on heavy-tailed distributions and stagnation detection, which results in an essentially parameter-less algorithm. Using this recommended configuration, we proved that the flex-EA achieves standard runtimes on ONEMAX and on LEADINGONES, and that it achieves a near best-possible runtime on JUMP for unbiased mutation-only algorithms. Furthermore, for a hurdle-like problem that features repetitive sub-structures, this configuration outperforms the Fast (1+1) EA (which uses heavy-tailed mutations) and stagnation detection. When provided with some mild problem-specific knowledge, the flex-EA matches the best known runtime result for the minimum spanning tree problem.

Although the performance of the flex-EA is very good on the problems we considered, the algorithm still has some shortcomings that can be improved. For one, on shifted jump benchmarks [41], the combination of heavy-tailed mutation and stagnation detection by Doerr and Rajabi [17] should prove superior, as the flex-EA basically defaults to classic stagnation detection. However, this shortcoming can be overcome by adjusting the mutation of the flex-EA in a similar way as done by Doerr and Rajabi [17] for their algorithm. Another potential shortcoming is that the archive of the flex-EA can become very crowded when many different mutation rates prove useful. If the mutation rates are removed in an incorrect order, this can result in long waiting times. It is unclear under what circumstances this effect actually occurs.

Overall, the flex-EA proposes an arguably new paradigm for evolutionary algorithms that performs very well and can be studied in multiple domains that exceed the content of this paper.



## REFERENCES

- [1] Denis Antipov, Maxim Buzdalov, and Benjamin Doerr. 2022. Fast mutation in crossover-based algorithms. *Algorithmica* 84 (2022), 1724–1761. <https://doi.org/10.1007/s00453-022-00957-5>
- [2] Denis Antipov and Benjamin Doerr. 2020. Runtime analysis of a heavy-tailed  $(1 + (\lambda, \lambda))$  genetic algorithm on jump functions. In *Prof. of PPSN '20*. Springer, 545–559. [https://doi.org/10.1007/978-3-030-58115-2\\_38](https://doi.org/10.1007/978-3-030-58115-2_38)
- [3] Dogan Corus, Pietro S. Oliveto, and Donya Yazdani. 2021. Fast immune system-inspired hypermutation operators for combinatorial optimization. *IEEE Transactions on Evolutionary Computation* 25, 5 (2021), 956–970. <https://doi.org/10.1109/TEVC.2021.3068574>
- [4] Duc-Cuong Dang, Anton V. Eremeev, and Per Kristian Lehre. 2021. Escaping local optima with non-elitist evolutionary algorithms. In *Proc. of AAAI '21*. AAAI Press, 12275–12283. <https://doi.org/10.1609/AAAI.V35I14.17457>
- [5] Duc-Cuong Dang, Anton V. Eremeev, Per Kristian Lehre, and Xiaoyu Qin. 2022. Fast non-elitist evolutionary algorithms with power-law ranking selection. In *Proc. of GECCO '22*. ACM, 1372–1380. <https://doi.org/10.1145/3512290.3528873>
- [6] Benjamin Doerr. 2020. Probabilistic tools for the analysis of randomized optimization heuristics. In *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 1–87. [https://doi.org/10.1007/978-3-030-29414-4\\_1](https://doi.org/10.1007/978-3-030-29414-4_1)
- [7] Benjamin Doerr and Carola Doerr. 2015. Optimal parameter choices through self-adjustment: Applying the 1/5-th rule in discrete settings. In *Proc. of GECCO '15*. ACM Press, 1335–1342. <https://doi.org/10.1145/2739480.2754684>
- [8] Benjamin Doerr and Carola Doerr. 2020. Theory of parameter control for discrete black-box optimization: Provable performance gains through dynamic parameter choices. In *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 271–321. [https://doi.org/10.1007/978-3-030-29414-4\\_1](https://doi.org/10.1007/978-3-030-29414-4_1)
- [9] Benjamin Doerr, Carola Doerr, and Jing Yang. 2016.  $k$ -bit mutation with self-adjusting  $k$  outperforms standard bit mutation. In *Proc. of PPSN '16*. Springer, 824–834. [https://doi.org/10.1007/978-3-319-45823-6\\_77](https://doi.org/10.1007/978-3-319-45823-6_77)
- [10] Benjamin Doerr, Yassine Ghannane, and Marouane Ibn Brahim. 2022. Towards a stronger theory for permutation-based evolutionary algorithms. In *Proc. of GECCO '22*. ACM, 1390–1398. <https://doi.org/10.1145/3512290.3528720>
- [11] Benjamin Doerr and Leslie Ann Goldberg. 2013. Adaptive Drift Analysis. *Algorithmica* 65, 1 (2013), 224–250. <https://doi.org/10.1007/S00453-011-9585-3>
- [12] Benjamin Doerr, Taha El Ghazi El Houssaini, Amirhossein Rajabi, and Carsten Witt. 2023. How well does the metropolis algorithm cope with local optima?. In *Proc. of GECCO '23*. ACM, 1000–1008. <https://doi.org/10.1145/3583131.3590390>
- [13] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. 2012. Multiplicative drift analysis. *Algorithmica* 64 (2012), 673–697. <https://doi.org/10.1007/s00453-012-9622-x>
- [14] Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. 2017. Fast genetic algorithms. In *Proc. of GECCO '17*. 777–784. <https://doi.org/10.1145/3071178.3071301>
- [15] Benjamin Doerr and Frank Neumann (Eds.). 2020. *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*. Springer. <https://doi.org/10.1007/978-3-030-29414-4>
- [16] Benjamin Doerr and Zhongdi Qu. 2023. A first runtime analysis of the NSGA-II on a multimodal problem. *IEEE Transactions on Evolutionary Computation* 27, 5 (2023), 1288–1297. <https://doi.org/10.1109/TEVC.2023.3250552>
- [17] Benjamin Doerr and Amirhossein Rajabi. 2023. Stagnation detection meets fast mutation. *Theoretical Computer Science* 946 (2023), 113670. <https://doi.org/10.1016/j.tcs.2022.12.020>
- [18] Tobias Friedrich, Andreas Göbel, Francesco Quinzan, and Markus Wagner. 2018. Heavy-tailed mutation operators in single-objective combinatorial optimization. In *Proc. of PPSN '18*. Springer, 134–145. [https://doi.org/10.1007/978-3-319-99253-2\\_11](https://doi.org/10.1007/978-3-319-99253-2_11)
- [19] Tobias Friedrich, Francesco Quinzan, and Markus Wagner. 2018. Escaping large deceptive basins of attraction with heavy-tailed mutation operators. In *Proc. of GECCO '18*. ACM, 293–300. <https://doi.org/10.1145/3205455.3205515>
- [20] Jun He and Xin Yao. 2001. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence* 127, 1 (2001), 51–81. [https://doi.org/10.1016/S0004-3702\(01\)00058-3](https://doi.org/10.1016/S0004-3702(01)00058-3)
- [21] Thomas Jansen and Ingo Wegener. 1999. On the analysis of evolutionary algorithms – A proof that crossover really can help. In *Proc. of ESA '99*. Springer, 184–193. [https://doi.org/10.1007/3-540-48481-7\\_17](https://doi.org/10.1007/3-540-48481-7_17)
- [22] Martin Krejca and Carsten Witt. 2020. Theory of estimation-of-distribution algorithms. In *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 405–442. <https://doi.org/10.1007/978-3-030-29414-4>
- [23] Timo Kötzing and Martin S. Krejca. 2019. First-hitting times under drift. *Theoretical Computer Science* 796 (2019), 51–69. <https://doi.org/10.1016/j.tcs.2019.08.021>
- [24] Johannes Lengler. 2020. Drift Analysis. In *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 89–131. [https://doi.org/10.1007/978-3-030-29414-4\\_1](https://doi.org/10.1007/978-3-030-29414-4_1)
- [25] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2020. Simple Hyper-Heuristics Control the Neighbourhood Size of Randomised Local Search Optimally for LeadingOnes. *Evolutionary Computation* 28, 3 (2020), 437–461. [https://doi.org/10.1162/EVCO\\_A\\_00258](https://doi.org/10.1162/EVCO_A_00258)
- [26] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2023. When move acceptance selection hyper-heuristics outperform Metropolis and elitist evolutionary algorithms and when not. *Artificial Intelligence* 314 (2023), 103804:1–103804:23. <https://doi.org/10.1016/j.artint.2022.103804>
- [27] Frank Neumann and Ingo Wegener. 2007. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Theoretical Computer Science* 378 (2007), 32–40. <https://doi.org/10.1016/j.tcs.2006.11.002>
- [28] Phan Trung Hai Nguyen and Dirk Sudholt. 2018. Memetic algorithms beat evolutionary algorithms on the class of hurdle problems. In *Proc. of GECCO '18*. ACM, 1071–1078. <https://doi.org/10.1145/3205455.3205456>
- [29] Francesco Quinzan, Andreas Göbel, Markus Wagner, and Tobias Friedrich. 2021. Evolutionary algorithms and submodular functions: Benefits of heavy-tailed mutations. *Natural Computing* 20, 3 (2021), 561–575. <https://doi.org/10.1007/s11047-021-09841-7>
- [30] Günther R. Raidl, Gabriele Koller, and Bryant A. Julstrom. 2006. Biased mutation operators for subgraph-selection problems. *IEEE Transaction on Evolutionary Computation* 10, 2 (2006), 145–156. <https://doi.org/10.1109/TEVC.2006.871251>
- [31] Amirhossein Rajabi and Carsten Witt. 2021. Stagnation detection in highly multimodal fitness landscapes. In *Proc. of GECCO '21*. ACM Press, 1178–1186. <https://doi.org/10.1145/3449639.3459336>
- [32] Amirhossein Rajabi and Carsten Witt. 2022. Self-adjusting evolutionary algorithms for multimodal optimization. *Algorithmica* 84, 6 (2022), 1694–1723. <https://doi.org/10.1007/s00453-022-00933-z>
- [33] Amirhossein Rajabi and Carsten Witt. 2023. Stagnation detection with randomized local search. *Evolutionary Computation* 31, 1 (2023), 1–29. [https://doi.org/10.1162/evco\\_a\\_00313](https://doi.org/10.1162/evco_a_00313)
- [34] I. Rechenberg. 1973. *Evolutionstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboog Verlag.
- [35] Joachim Reichel and Martin Skutella. 2009. On the size of weights in randomized search heuristics. In *Proc. of FOGA 2009*. 21–28. <https://doi.org/10.1145/1527125.1527130>
- [36] Günter Rudolph. 1997. *Convergence properties of evolutionary algorithms*. Verlag Dr. Kovač.
- [37] M. Schumer and Kenneth Steiglitz. 1968. Adaptive step size random search. *IEEE Transactions on Automatic Control* 13, 3 (1968), 270–276. <https://doi.org/10.1109/TAC.1968.1098903>
- [38] Dirk Sudholt. 2013. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation* 17, 3 (2013), 418–435. <https://doi.org/10.1109/TEVC.2012.2202241>
- [39] Ingo Wegener. 2001. Theoretical aspects of evolutionary algorithms. In *Proc. of ICALP '01*. Springer, 64–78. [https://doi.org/10.1007/3-540-48224-5\\_6](https://doi.org/10.1007/3-540-48224-5_6)
- [40] Carsten Witt. 2014. Fitness levels with tail bounds for the analysis of randomized search heuristics. *Information Processing Letters* 114, 1-2 (2014), 38–41. <https://doi.org/10.1016/j.ipl.2013.09.013>
- [41] Carsten Witt. 2023. How majority-vote crossover and estimation-of-distribution algorithms cope with fitness valleys. *Theoretical Computer Science* 940 (2023), 18–42. <https://doi.org/10.1016/j.tcs.2022.08.014>
- [42] Mengxi Wu, Chao Qian, and Ke Tang. 2018. Dynamic mutation based Pareto optimization for subset selection. In *Proc. of ICIC '18*. Springer, 25–35. [https://doi.org/10.1007/978-3-319-95957-3\\_4](https://doi.org/10.1007/978-3-319-95957-3_4)
- [43] Christine Zarges. 2020. Theoretical foundations of immune-inspired randomized search heuristics for optimization. In *Theory of Evolutionary Computation – Recent Developments in Discrete Optimization*, Benjamin Doerr and Frank Neumann (Eds.). Springer, 443–474. [https://doi.org/10.1007/978-3-030-29414-4\\_1](https://doi.org/10.1007/978-3-030-29414-4_1)
- [44] Weijie Zheng and Benjamin Doerr. 2023. Theoretical analyses of multiobjective evolutionary algorithms on multimodal objectives. *Evolutionary Computation* 31, 4 (2023), 337–373. [https://doi.org/10.1162/EVCO\\_A\\_00328](https://doi.org/10.1162/EVCO_A_00328)

## A APPENDIX

This appendix contains material that does not fit the main paper, due to space restrictions. It is meant to be read at the reviewer's discretion only. It contains mathematical tools that we use in some of our proofs and proofs that had to be omitted from the main paper.

### UNIMODAL FUNCTIONS

The fitness level method was introduced by Wegener [39] for the analysis of evolutionary algorithms on a given fitness function  $f$ . It requires a partition  $(L_i)_{i \in [m]}$  of the search space  $\{0, 1\}^n$  (with  $m, n \in \mathbb{N}_{\geq 1}$ ) such that the fitness increases in each level. That is, for all  $i \in [m-1]$  and all  $\mathbf{x} \in L_i$  and  $\mathbf{y} \in L_{i+1}$ , it holds that  $f(\mathbf{x}) < f(\mathbf{y})$ . We call such a partition a *fitness level partition* and its elements *fitness levels*.

The following theorem by Sudholt [38] is used in the proof of [Theorem 5.1](#).

**THEOREM A.1** ([38, THEOREM 2]). *Let  $f$  be a pseudo-Boolean function, let  $m \in \mathbb{N}_{\geq 1}$ , and let  $(L_i)_{i \in [m]}$  be a fitness level partition of  $f$ . Let  $A$  be an elitist search algorithm, and, for all  $i \in [m-1]$ , let  $s_i$  denote a lower bound on the probability of  $A$  creating a new search point in  $\bigcup_{j \in [i+1..m]} L_j$ , provided that the best individual created so far by  $A$  is in  $L_i$ . Then the expected optimization time of  $A$  on  $f$  (without the initialization) is at most  $\sum_{i \in [m-1]} \frac{1}{s_i}$ .*

### JUMP FUNCTIONS

The following lemma by Rajabi and Witt [31] is used in the proof of [Theorem 6.1](#).

**LEMMA A.2** ([31, LEMMA 1 (A)]). *Let  $n, k \in \mathbb{N}$  with  $k \leq \frac{n}{2}$ . Then  $\sum_{i \in [k]} \binom{n}{i} \leq \frac{n - (k-1)}{n - (2k-1)} \binom{n}{k}$ .*

### BENEFITS OF THE RATE ARCHIVE ON MINIMUM SPANNING TREES (SECTION 7)

Drift analysis is the most prominent mathematical tool for analyzing evolutionary algorithms. It was introduced by He and Yao [20]. We use the following theorem by Doerr et al. [13] in the proof of [Theorem 7.1](#), presented in a version close to the one by Kötzing and Krejca [23]. The tail bound can be found in [11] and [24].

**THEOREM A.3** ([13, THEOREM 3]). *Let  $(X_t)_{t \in \mathbb{N}}$  be a Markov chain over  $\mathbb{R}_{\geq 0}$ , and let  $T = \inf\{t \in \mathbb{N} \mid X_t < 1\}$ . Suppose that there is a  $\delta > 0$  such that it holds that  $\mathbb{E}[X_t - X_{t+1} \mid X_t] \cdot \mathbb{1}\{t < T\} \geq \delta X_t \cdot \mathbb{1}\{t < T\}$ . Then  $\mathbb{E}[T] \leq (1 + \mathbb{E}[\ln(X_0)]) / \delta$ . Moreover,  $\Pr[T \geq (\mathbb{E}[\ln(X_0)] + r)] \leq e^{-r}$ .*

We now give the full proof of [Theorem 7.1](#) that was omitted from the main paper.

**PROOF OF THEOREM 7.1.** Similarly to the classical analysis of the (1+1) EA on the problem [27], we divide the run into three phases. The first two phases are considered jointly. They can find improvements in the 1-bit flip neighborhood by including edges (i. e., flipping zero-bits) and reducing the number of connected components or by reducing the number of edges from a fully connected subgraph by flipping one-bits. The expected number of 1-bit flips until the first phase is finished is bounded by  $O(m \log m)$  in expectation, as shown in [27] using fitness-level arguments. Using tail bounds for fitness levels [40], the number is  $c' m \ln m$  for some constant  $c' > 0$  with probability  $1 - o(1)$ . As long as fitness improvements by 1-bit flips are still available, the expected time until choosing rate 1 is at most  $1/\ell_1 = o(\log m)$ , so both in expectation and with high probability after  $o(m \log^2 m)$  steps a local optimum with respect to the 1-bit neighborhood has been reached and the current search point constitutes a (not necessarily minimum) spanning tree. If the constant  $c$  from  $C_1$  is chosen sufficiently large, then the failure counter  $c_1$  does not exceed the threshold  $C_1$  during the considered number of steps at rate 1 with probability at least  $1 - o(1/m^2)$ . Moreover, by a union bound, the probability that mutation rates higher than 2 are ever chosen in the considered period is

$$O\left(m \log^2 m\right) \sum_{t=3}^m \ell_t = o\left(1/m^3\right),$$

which we assume not to happen. Then no rates larger than 2 are in the archive  $A^{(t)}$  at the end of the period. If rate 2 is not yet present at that time, then the global counter  $G^{(t)}$  exceeds its threshold after  $O(m \ln^2 m)$  steps – and rate 2 enters the archive after another  $O(C_1/\ell_1) = O(m \ln^2 m)$  expected number of steps at rate 1, all of which are unsuccessful. Again, the probability of rates larger than 2 entering the archive during that time is  $o(1/m^3)$  by the same arguments as above.

After the current search point constitutes a spanning tree for the first time, the final phase begins. Here only steps flipping an even number of bits can be accepted since disconnected graphs and graphs having more than  $n-1$  edges lead to inferior fitness values. We now concentrate on 2-bit flips and aim at using multiplicative drift analysis using  $g(x) = \sum_{i=1}^m x_i r_i$  as potential function, where  $r_i$  is defined in the statement of the theorem. As shown in [30, Proof of Lemma 3], RLS flipping one or two bits uniformly at random behaves stochastically identically on the original fitness function  $f$  and the function  $g$ , assuming a spanning tree as starting point. Hence, our plan is to choose a sufficiently large constant  $c' > 0$  and to conduct the drift analysis conditioned on that only rates 1 and 2 are chosen in a phase of predefined length  $c'(1 - \ell_1)^{-1} m^2 \ln m = O(m^2 \ln^2 m)$ , using  $\ell_1 \leq 1/2$ . Note that in the phase, the steps of rate 1 cannot be successful. The probability of only choosing rates at most 2 throughout the phase is at least

$$1 - O\left(m^2 \ln^2 m\right) \sum_{t=3}^m \ell_t = 1 - O\left(1/m^2\right).$$

In the following, we assume the phase to fulfill this property. We choose  $c$  small enough to bound the probability of the error that  $c_2^{(\cdot)}$  exceeds its threshold  $C_2$  by  $O(m^{-2})$  and analyze the error case at the end of the proof.

Let  $X^{(t)} := g(x_t) - g(x_{\text{opt}})$  for the current search point  $x_t$  and an optimal search point  $x_{\text{opt}}$ . Since d-SD RLS under the assumption of choosing only rates 1 and 2 behaves stochastically the same on the original fitness function  $f$  and the potential function  $g$ , we obtain that  $E[X^{(t)} - X^{(t+1)} | X^{(t)}] \geq X^{(t)}/\binom{m}{2} \geq 2X^{(t)}/m^2$  since the  $g$ -value can be decreased by altogether  $g(x_t) - g(x_{\text{opt}})$  via a sequence of at most  $\binom{m}{2}$  disjoint two-bit flips; see also the proof of Theorem 15 in [13] for the underlying combinatorial argument. Let  $T$  denote the number of steps at strength 2 until  $g$  is minimized, assuming no larger strength to occur. Using the multiplicative drift theorem (Theorem A.3) we have  $E[T] \leq (m^2/2)(1 + \ln(r_1 + \dots + r_m)) \leq (m^2/2)(1 + \ln(m^2))$  and by the tail bounds for multiplicative drift (e. g., [24]) it holds that  $\Pr[T > (m^2/2)(\ln(m^2) + \ln(m^3))] \leq e^{-\ln(m^3)} = 1/m^3$ . We choose a phase length of  $4(1 - \ell_1)^{-1}m^2 \ln m$ , which includes at least  $m^2 \ln m$  steps choosing rate 2 with probability  $1 - e^{-\Omega(m^2 \ln m)}$  according to Chernoff bounds [6, Theorem 1.10.5]. Assuming this to happen, we have bounded the probability of not finding the optimum in Phase 2 of that length is  $o(m^{-2})$ . The total probability of any error as described above is still  $o(m^{-2})$ .

To bound the total expected optimization time, we note that we already have proved an upper bound  $o(m^2 \log m)$  on the expected time to reach a spanning tree conditional on that no error occurs. If an error occurs, we wait for an accepted one-bit flip, which happens after expected time  $O(m \log m)$  and repeat the argumentation. The expected number of repetitions is  $o(1)$ .

We now bound the expected time for the final phase. Hence, we assume a spanning tree as current search point. If an error occurs, we have no real control over the archive  $A^{(t)}$ . However, after an expected number of at most  $(1/\ell_2)m^2 = O(m^4 \ln m)$  steps, an accepted two-bit mutation adds rate 2 to the archive and progress can be made by flipping two bits as above, with the caveat that the stochastically equivalent behavior as on the function  $g$  cannot be assumed any longer. Hence, we use a weaker bound on the expected optimization time. More precisely, we replace the fitness function  $f$  by the fitness function  $f'$  under which the  $2^m$  search points have the same ranking as on  $f$ , but the maximum weight is  $O(m^m)$  [35]. Since the flex-EA is ranking based, it has the same stochastic behavior on  $f$  and  $f'$ . On the process  $X'_t := f'(x_t)$ , using the classical analysis by [27], there is a multiplicative drift with parameter  $\delta = \Omega(q/m^2)$ , where  $q$  is a lower bound on a step flipping two bits, while steps flipping more bits do not harm the analysis. Clearly,  $q = \Omega(1/m)$  since no more than  $m$  rates can be in the archive. Hence, we obtain an expected runtime of  $O(m \cdot m^2 \ln(m^m)) = O(m^4 \ln m)$  in the failure case. By the law of total probability, using the failure probability  $o(m^{-2})$ , the total expected runtime is still  $(1 - \ell_1)^{-1} + o(1))(m^2/2)(1 + \ln(r_1 + \dots + r_m))$ .  $\square$

## A FUNCTION WHERE THE ARCHIVE STORES SEVERAL RATES SIMULTANEOUSLY (SECTION 8)

We now give the full proofs of Lemma 8.2 and Theorem 8.1 that were omitted from the main paper.

**PROOF OF LEMMA 8.2.** To reach  $a + d$  one-bits by flipping  $d + 2i$  bits, exactly  $d + i$  zero-bits and  $i$  one-bits must be flipped. Hence, we assume that  $d + i \leq n - a$  and  $i \leq a$ . By elementary combinatorics and expanding the binomial coefficients, we have

$$\frac{p(n, a, d, i)}{p(n, a, d, 0)} = \frac{\binom{n-a}{d+i} \binom{a}{i}}{\binom{n-a}{d}} = \frac{(d+2i)!}{i!(d+i)!} \cdot \frac{(n-a)^{d+i} \cdot a^i}{n^{d+2i}} \cdot \frac{n^d}{(n-a)^d} = \frac{(d+2i)!}{i!(d+i)!} \cdot \frac{(n-a-d)^i \cdot a^i}{(n-d)^{2i}},$$

where  $b^{\underline{c}} = b(b-1) \dots (b-c+1)$  is the falling factorial (sometimes called *Pochhammer* symbol).

We estimate the two fractions in the last expression separately and start with the second one. By our assumption that  $d + 2i \leq n^{1/3}$ , we have  $(n-d)^{2i} \geq (n-d-2i)^{2i} \geq (n-n^{1/3})^{2i} = (1-o(1))n^{2i}$  and therefore

$$\frac{(n-a-d)^i \cdot a^i}{(n-d)^{2i}} \leq \frac{(n-a)^i n^i}{(1-o(1))n^{2i}} \leq (1+o(1)) \left(\frac{n-a}{n}\right)^i.$$

For the first fraction, we have

$$\frac{(d+2i)!}{i!(d+i)!} = \frac{(d+2i)^{\underline{d+2i}}}{i!} \leq \frac{(d+2i)^i}{(i/e)^i} = \left(2 + \frac{d}{i}\right)^i \cdot e \leq \left(e^{1+d/i}\right)^i e^i = e^{d+2i},$$

where the first inequality used the well-known estimate  $i! \geq (i/e)^i$  following from Stirling's formula and the second inequality used that  $1+x \leq e^x$  for all  $x \in \mathbb{R}$ .

Multiplying the two estimates gives the claim of the lemma.  $\square$

**PROOF OF THEOREM 8.1.** We start with the upper bound for the flex-EA.

*Reaching the first hurdle.* If the initial search point has less than  $s := (7/8)n - \sqrt{n}$  one-bits, then by Corollary 5.2, the level of at least  $s$  one-bits is reached within expected time  $O(n \log n / \ell_1) = O(n \log n)$  and rate 1 is in the archive with probability  $1 - O(1/n^2)$  by choosing the constant in  $C_1$  appropriately. Moreover, by applying Markov's inequality and repeating independent phases, the time to reach level at least  $s$  is bounded by  $O(n^4)$  with probability  $1 - 2^{-\omega(n)}$ .

We now consider the first time where the current search point has  $s$  or more one-bits and assume pessimistically exactly  $s$  one-bits. The smallest feasible rate at this point is  $g$  and the probability of a success, i. e., increasing the number of one-bits by  $g$ , at this rate is  $(1 + o(1))n^{-2}$  according to (1). Using Lemma 8.2, the probability of a success at rate  $g + 2i$ , where  $i > 0$ , is bounded from above by

$$(1 + o(1))n^{-2} \left( \left( 2 + \frac{g}{i} \right) \cdot e \cdot \left( \frac{1}{8} + \frac{1}{\sqrt{n}} \right) \right)^i \leq (1 + o(1))n^{-2} \left( \left( 2 + \frac{g}{i} \right) \cdot e \cdot \frac{1}{7.9} \right)^i$$

using  $a \geq s = 7n/8 - \sqrt{n}$  and assuming  $n$  large enough. The bracket attains its maximum at  $i \approx 0.2868g$  (which can be verified by a standard computer algebra system) and is bounded from above by  $1.89^{0.2868g} \leq 2^{g/3} = n^{1/6}$  then. Altogether, the success probability for a hurdle of size  $g$  even at optimal rate is no more than  $(1 + o(1))n^{-11/6}$ .

Therefore with probability  $1 - o(1)$  by a union bound over  $G = O(n \log n / \ell_1)$  steps, the global counter  $g$  (not to be confused with the gap size) will reach its threshold before a hurdle is jumped over and the archive will be reset to contain rate 1 only. We assume all this to happen.

*Bounding the archive size.* When the current search point has  $s$  one-bits, then successes are only possible at rates  $2g$  or larger. The aim is to bound the size of the archive while the hurdles of width  $g$  and  $2g$  are overcome. We first rule out that extremely large rates yield a success and thereby enter the archive. Applying a mutation flipping  $d$  bits uniformly at random to a search point with  $a \geq s$  one-bits, the number of flipping one-bits follows a hypergeometric distribution and the expected increase in one-bits equals  $ad/n - (n - a)d/n < 0$ . Applying a Chernoff bound for the hypergeometric distribution (Theorem 1.10.25 in [6]), the probability of increasing the number of one-bits by any positive amount when flipping a certain number  $d \geq n^{1/3}$  of bits in one step is bounded from above by  $2^{-\Omega(n^{1/3})}$ . By a union bound over all  $d \geq n^{1/3}$ , the probability of increasing the number of one-bits using any rate at least  $n^{1/3}$  is still  $2^{-\Omega(n^{1/3})}$ . Hence, we work under the assumption that rates larger than  $n^{1/3}$  are never successful in polynomial time; the error probability is still  $2^{-\Omega(n^{1/3})}$  by a union bound.

Clearly, the minimum successful rate is  $g$  while crossing the hurdles in the interval  $[s, (7/8)n]$  of one-bits. We now estimate the success probabilities of rates in  $[g, n^{1/3}]$ . By flipping  $2g + i$  zero-bits and  $i$  one-bits at rate  $2g + 2i$ , also larger rates than  $2g$  can overcome hurdles of size  $2g$ . Using Lemma 8.2 using  $d = 2g$ ,  $i = c' \log n$  for a sufficiently large constant  $c'$ ,  $a \geq s$  and our assumption of rates at most  $n^{1/3}$ , we have that the probability of increasing the number of one-bits by  $d$  at rate  $d + 2i$  is no larger than a factor of

$$(1 + o(1))e^{\log n} e^{2i} \left( \frac{1}{8} + \frac{1}{\sqrt{n}} \right)^i \leq (1 + o(1))e^{\log n} 0.93^i$$

larger than the probability of doing this at rate  $d$ . Hence, for a sufficiently large constant  $c' > 0$ , rates larger than  $2g + 2c' \log n$  only have a success probability of  $O(n^{-5})$ . By a union bound, the size of the archive is bounded by  $O(\log n)$  within a period of  $\Theta(n^{4.5})$  steps with probability  $1 - O(n^{-1/2})$ . We assume that the error case of a larger archive does not happen either.

*Crossing the hurdles and reaching the optimum.* In the following, under the above assumptions, we first show that a phase of  $c''n^{4.5}$  steps is sufficient to cross all hurdles by reaching at least  $7n/8$  one-bits with high probability. No matter whether there is a current hurdle of width  $g$  or  $2g$  one-bits to the next improvement, the rate having the highest success probability (which is at least  $(1 - o(1))n^{-2}$  and  $(1 - o(1))n^{-4}$ , respectively, but possibly larger since it may be beneficial to flip more bits than the gap size to overcome the gap), enters the archive after  $O(\ell_{2g}n^4)$  steps since it is sufficient to select the best rate and have a success (using that  $\ell_g \leq \ell_{2g}$ ). By Chernoff bounds for geometrically distributed random variables [6, 1.10.32], the time is  $O(\ell_{2g}n^4)$  with overwhelming probability. Once both the optimal rate for gap  $g$  and for gap  $2g$  are in the archive, the probability of choosing the rate fitting the next hurdle is  $\Omega(1/\log n)$  by our assumption on the archive size and the time for an improvement is  $O(n^4 \log n)$  with overwhelming probability and in expectation (using that gap  $2g$  is the worst case). Since at most  $\sqrt{n}/g = O(\sqrt{n}/\log n)$  improvements jumping hurdles are sufficient to reach level  $7n/8$  or larger and the remaining part of the optimization is completed within expected time  $O((n \log n)/\ell_1)$  in the same way as described above, the optimization time is  $O(n^{4.5})$  in expectation and also with high probability by Chernoff bounds for geometrically distributed random variables. In the error case of a larger archive, we argue with the maximum archive size  $n$  and have a success probability of  $\Omega(1/n^5)$ . The expected time in the error case is therefore  $O(n^{5.5}/\log n)$ . Finally, the probability of exceeding a counter is superpolynomially small since  $C_g \geq \binom{n}{\log n}$ ; if this happens nevertheless, then we wait for the optimum rates for gap sizes  $g$  and  $2g$  to enter the archive as above. By the law of total probability, this contributes a superpolynomially small extra term to the expected runtime. Altogether, the unconditional expected optimization time of flex-EA is  $O(n^{4.5})$ .

*Lower Bounds for stagnation detection and Fast (1+1) EA.* We are left with the bounds for the other algorithms. With overwhelming probability  $1 - 2^{-\Omega(n)}$ , their first search point, which is drawn uniformly, is before the first hurdle, i. e., has at most  $s$  one-bits. Hence, in the algorithms SD-RLS<sup>f</sup> and SD-RLS<sup>m</sup>, the rate must increase from its initial value 1 to at least  $g$  as least once. With the standard parameter choices, this takes  $\Omega(\sum_{i=1}^{g-1} \binom{n}{i} \ln R) = n^{\Omega(\log n)}$  steps.

The Fast (1+1) EA, under the assumption on the first point, must increase the number of one-bits by a total of at least  $\sqrt{n}$  to cross all hurdles. We show that the algorithm does not cross more than  $h^*$  hurdles at once in a period of  $O(n^5)$  steps with high probability, where  $h^*$  is a sufficiently large constant. To cross  $h^*$  hurdles in one mutation, it is necessary that at least  $cn$  zero-bits flip at once for some constant  $c > 0$ . By the results in [14], for any  $k \in [n/2]$ , the probability of flipping exactly  $k$  bits is maximized at mutation probability  $k/n$  (recall that the

Fast (1+1) EA uses standard bit mutation). Hence, the probability of flipping  $k \geq cn$  zero-bits and thereby crossing at least  $h^*$  hurdles at once is at most

$$\binom{n-s}{k} \left(\frac{k}{n}\right)^k \leq \left(\frac{(n-s)e}{k}\right)^k \left(\frac{k}{n}\right)^k \leq \left(\frac{1}{8} + \frac{1}{\sqrt{n}}\right)^k e^k \leq \left(\frac{e}{8} + \frac{e}{\sqrt{n}}\right)^{cg} = O(n^{-6})$$

if  $h^*$ , and thereby  $c$ , are large enough. By a union bound over all  $k \geq cn$ , the probability of crossing at least  $h^*$  hurdles at once is still  $O(n^{-5})$ . In the following, we assume that this does not happen in  $O(n^{4.6})$  steps.

Under our assumptions, to reach the optimum, at least  $\sqrt{n}/(cg)$  increases of the number of one-bits by at least  $cg$  in search points at level  $s$  or larger are necessary, where  $c$  may be a large constant. We pessimistically (for the perspective of a lower bound) assume that the heavy-tailed mutation operator always flips a number of bits  $2g + 2i$  that maximizes the probability of overcoming a hurdle of width  $2g$  (and in the same way for the hurdles of width  $g$ ). As mentioned above, the Fast (1+1) EA flips  $2g + 2i$  bits with the highest probability at mutation probability  $(2g + 2i)/n$ . Hence, the probability of crossing a hurdle of width  $2g$  is at most

$$\binom{n-s}{2g+i} \binom{s}{i} \left(\frac{2g+2i}{n}\right)^{2g+2i} \left(1 - \frac{2g+2i}{n}\right)^{n-2g-2i} = p(n, s, 2g, i) \cdot \left(\frac{2g+2i}{n}\right)^{2g+2i} \left(1 - \frac{2g+2i}{n}\right)^{n-2g-2i} \binom{n}{2g+2i}.$$

Inspecting the product following  $p(n, s, 2g, i)$ , which is the ratio of success probabilities of the Fast (1+1) EA and of the flex-EA, we have

$$\left(\frac{2g+2i}{n}\right)^{2g+2i} \left(1 - \frac{2g+2i}{n}\right)^{n-2g-2i} \binom{n}{2g+2i} \leq \left(\frac{2g+2i}{n}\right)^{2g+2i} e^{-2g-2i+o(1)} \frac{n^{2g+2i}}{(2g+2i)!},$$

where the  $o(1)$  comes from the assumption of  $2g+2i \leq n^{1/3}$ . Using Stirling's formula, we have  $(2g+2i)! \geq \frac{1}{\sqrt{2\pi}} \sqrt{2g+2i} \cdot e^{-2g-2i} (2g+2i)^{2g+2i}$ . Plugging this into the previous formula, the bound on the ratio of success probabilities simplifies to

$$(1 + o(1)) \sqrt{2\pi} \frac{1}{\sqrt{2g+2i}},$$

so the probability of the Fast (1+1) EA crossing a hurdle is by a factor of  $\Omega(\sqrt{2g+2i}) = \Omega(\sqrt{\log n})$  smaller than when the flex-EA has chosen the optimum radius.

Finally, we analyze the probability of the Fast (1+1) EA choosing a successful rate in the heavy-tailed mutation operator. If its rate is at most  $c' \log n$  for a constant  $c' > 0$ , the probability of crossing a hurdle of width  $g$  is at most

$$\binom{n-s}{g} \left(\frac{c \ln n}{n}\right)^g \leq \left(\frac{c' \ln n}{\ln n}\right)^{\ln n} = O(1/n^5)$$

choosing  $c'$  small enough. Hence, by a union bound, no hurdle is crossed at rate less than  $c' \log n$  in a period of  $\Omega(n^{4.5} \sqrt{\ln n})$  steps, which we assume to happen. On the other hand, the probability of the heavy-tailed mutation choosing rate at least  $c' \ln n$  is at most  $O((1/\ln n)^\beta)$ . This probability is by a factor  $O((1/\ln n)^{\beta-1})$  smaller than  $1/\log n$ , the asymptotic worst-case probability of the flex-EA choosing the best radius from an archive of  $O(\log n)$  bits. Since both algorithms have to overcome  $\Theta(\sqrt{n}/\log n)$  hurdles with high probability and the success probability of the flex-EA for crossing a hurdle is bigger by a factor of  $\Omega(\sqrt{\log n})$ , we obtained the claimed asymptotic speed-up of  $\Omega((\ln n)^{\beta-1/2})$  for the flex-EA compared to the Fast (1+1) EA.  $\square$