



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/143805/>

Version: Accepted Version

Article:

Corus, D., Oliveto, P.S. and Yazdani, D. (2020) When hypermutations and ageing enable artificial immune systems to outperform evolutionary algorithms. *Theoretical Computer Science*, 832. pp. 166-185. ISSN: 0304-3975

<https://doi.org/10.1016/j.tcs.2019.03.002>

Article available under the terms of the CC-BY-NC-ND licence
(<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

When Hypermutations and Ageing Enable Artificial Immune Systems to Outperform Evolutionary Algorithms*

Dogan Corus, Pietro S. Oliveto, Donya Yazdani

Rigorous Research, Department of Computer Science, University of Sheffield

Sheffield, UK

S1 4DP

Abstract

We present a time complexity analysis of the Opt-IA artificial immune system (AIS). We first highlight the power and limitations of its distinguishing operators (i.e., hypermutations with mutation potential and ageing) by analysing them in isolation. Recent work has shown that ageing combined with local mutations can help escape local optima on a dynamic optimisation benchmark function. We generalise this result by rigorously proving that, compared to evolutionary algorithms (EAs), ageing leads to impressive speed-ups on the standard CLIFF_d benchmark function both when using local and global mutations. Unless the *stop at first constructive mutation* (FCM) mechanism is applied, we show that hypermutations require exponential expected runtime to optimise any function with a polynomial number of optima. If instead FCM is used, the expected runtime is at most a linear factor larger than the upper bound achieved for any random local search algorithm using the artificial fitness levels method. Nevertheless, we prove that algorithms using hypermutations can be considerably faster than EAs at escaping local optima. An analysis of the complete Opt-IA reveals that it is efficient on the previously considered functions and highlights problems where the use of the full algorithm is crucial. We complete the picture by presenting a class of functions for which Opt-IA fails with overwhelming probability while standard EAs are efficient.

Keywords: Artificial Immune Systems, Opt-IA, Runtime Analysis, Evolutionary Algorithms, Hypermutation, Ageing

1. Introduction

Artificial immune systems (AIS) are a class of bio-inspired computing techniques that take inspiration from the immune system of vertebrates [2]. Burnet's clonal selection the-

*An extended abstract of this paper has been published at the 2017 Genetic and Evolutionary Computation Conference [1].

Email addresses: d.corus@sheffield.ac.uk (Dogan Corus), p.oliveto@sheffield.ac.uk (Pietro S. Oliveto), dyazdani1@sheffield.ac.uk (Donya Yazdani)

ory [3] has inspired various AIS for function optimisation. The most popular ones are Clonalg [4], the B-Cell Algorithm [5] and Opt-IA [6].

After numerous successful applications of AIS were reported, a growing body of theoretical work has gradually been built to shed light on the working principles of AIS. While initial work derived conditions that allowed to prove whether an AIS converges or not [7], nowadays rigorous time complexity analyses of AIS are available. Initial runtime analyses focused on studying the performance of typical AIS operators in isolation to explain when and why they are effective. Such studies have been extensively performed for the contiguous somatic hypermutation operator employed by the B-Cell algorithm [8, 9], the inversely proportional hypermutation operator of Clonalg [10, 11] and the ageing operator used by Opt-IA [12, 13, 14]. These studies formed a foundational basis which allowed the subsequent analysis of the complete B-Cell algorithm as used in practice for standard combinatorial optimisation [15, 16].

Compared to the relatively well understood B-Cell algorithm, the theoretical understanding of other AIS for optimisation is particularly limited. In this paper we consider the complete Opt-IA algorithm [17, 6]. This algorithm has been shown experimentally to be successful at optimising instances of problems such as protein structure prediction [6], graph colouring [18] and hitting set [19]. The main distinguishing features of Opt-IA compared to other AIS is their use of an ageing operator and of hypermutations with mutation potentials. In this work we will first analyse the characteristics of these operators respectively in isolation and afterwards consider a simple, but complete, Opt-IA algorithm. The aim is to highlight function characteristics for which Opt-IA and its main components are particularly effective, hence when it may be preferable to standard Evolutionary Algorithms (EAs).

The idea behind the ageing operator is that old individuals should have a lower probability of surviving compared to younger ones. Ageing was originally introduced as a mechanism to maintain diversity. Theoretical analyses have strived to justify this initial motivation because the new random individuals (introduced to replace old individuals) typically have very low fitness and die out quickly. On the other hand, it is well understood that ageing can be used as a substitute for a standard restart strategy if the whole population dies at the same generation [12] and to escape local optima if most of the population dies except for one survivor that, at the same generation, moves out of the local optimum [14]. This effect was shown for a Random Local Search (RLS) algorithm equipped with ageing on the BALANCE dynamic optimisation benchmark function. An evolutionary algorithm (EA) using standard bit mutation (SBM) [20, 21] and ageing would not be able to escape the local optima due to their very large basin of attraction. Herein, we carefully analyse the ability of ageing to escape local optima on the more general CLIFF_d benchmark function and show that using the operator with both RLS and EAs can make a difference between polynomial and exponential runtimes.

Hypermutation operators are inspired by the high mutation rates occurring in the immune system. In Opt-IA the *mutation potential* is linear in the problem size, and in different algorithmic variants may be static or either increase by a factor that is proportional to the fitness of the solution (i.e., the b-cell) undergoing the mutation or decrease by a factor that is inversely proportional to the fitness. The theoretical understanding of hypermutations

with mutation potential is very limited. To the best of our knowledge the only runtime analysis available is [22], where inversely proportional hypermutations were considered, with and without the *stop at first constructive mutation (FCM)* strategy². The analysis revealed that, without FCM, the operator requires exponential runtime to optimise the standard ONEMAX function, while by using FCM the algorithm is efficient. We consider a different hypermutation variant using static mutation potentials and argue that it is just as effective if not superior to other variants. We first show that the use of FCM is essential by rigorously proving that a (1+1) EA equipped with hypermutations and no FCM requires exponential expected runtime to optimise any function with a polynomial number of optima. We then consider the operator with FCM for any objective function that can be analysed using the artificial fitness level (AFL) method [20, 21] and show an upper bound on its runtime that is at most by a linear factor larger than the upper bound obtained for any RLS algorithm using AFL. To achieve this, we present a theorem that allows to derive an upper bound on the runtime of sophisticated hypermutation operators by analysing much simpler RLS algorithms with an arbitrary neighbourhood. As a result, all existing results achieved via AFL for RLS may be translated into upper bounds on the runtime of static hypermutations. Finally, we use the standard CLIFF_d and JUMP_k benchmark functions to show that hypermutations can achieve considerable speed-ups for escaping local optima compared to well studied EAs.

We then concentrate on the analysis of the complete Opt-IA algorithm. The standard Opt-IA uses both hypermutations and hypermacromutation (both with FCM) mainly because preliminary experimental studies for trap functions indicated that this setting led to the best results [17, 6]. Our analysis reveals that it is unnecessary to use both operators for Opt-IA to be efficient on trap functions. To this end, we will consider the simple version using only static hypermutations as in [17]. We will first consider the algorithm with the simplification that we allow genotypic duplicates in the population, to simplify the analysis and enhance the probabilities of ageing to create copies and escape from local optima. Afterwards we extend the analysis to the standard version using a genotype diversity mechanism. Apart from proving that the algorithm is efficient for the previously considered functions, we present a class of functions called HIDDENPATH , where it is necessary to use both ageing and hypermutations in conjunction, hence where the use of Opt-IA in its totality is crucial. Having shown several general settings where Opt-IA is advantageous compared to standard EAs, we conclude the paper by pointing out limitations of the algorithm. In particular, we present a class of functions called HYPERTRAP that is deceptive for Opt-IA while standard EAs optimise it efficiently with overwhelming probability.

Compared to its conference version [1], this paper has been improved in several ways. Firstly, we have extended our analyses of the ageing operator and Opt-IA to include the genotype diversity mechanism as in the algorithm proposed in the literature [17, 6]. Another addition is the introduction of a class of functions where Opt-IA fails to find the optimum

²An analysis of high mutation rates in the context of population-based evolutionary algorithms was performed in [23]. Increasing the mutation rate above the standard $1/n$ value has gained interest in recent years [24, 25, 26, 27, 28].

Algorithm 1 Opt-IA* [6]. Subroutines are described in Algorithms 2, 3, 4 and 5.

- 1: $t = 0$,
- 2: initialise $P^{(t)} = \{x_1, \dots, x_\mu\}$, a population of μ b-cells uniformly at random and set $x_i^{age} = 0$ for $i = \{1, \dots, \mu\}$.
- 3: **while** termination condition is not reached **do**
- 4: $P^{(clo)} = \text{Cloning}(P^{(t)}, dup)$,
- 5: $P^{(hyp)} = \text{Hypermutation}(P^{(clo)}, c)$,
- 6: $P^{(macro)} = \text{Hypermacromutation}(P^{(clo)})$,
- 7: Ageing $(P^{(t)}, P^{(hyp)} \cup P^{(macro)}, \tau, \mu)$,
- 8: $P^{(t+1)} = \text{Selection}(P^{(t)}, P^{(hyp)} \cup P^{(macro)}, \mu, 1)$.
- 9: $t = t + 1$.
- 10: **end while**

Algorithm 2 Cloning $(P^{(t)}, dup)$

- 1: $P^{(clo)} = \emptyset$.
- 2: **for** all $x_i \in P^{(t)}$ **do**
- 3: copy x_i dup times,
- 4: add the copies to $P^{(clo)}$.
- 5: **end for**

efficiently, allowing us to complete the picture by highlighting problem characteristics where Opt-IA succeeds and where it does not. Finally, this paper includes some proofs which were omitted from the conference version due to page limitations.

The rest of the paper is structured as follows. In Section 2, we introduce and define Opt-IA and its operators. In Section 3, we present the results of our analyses of the static hypermutation operator in a simple framework to shed light on its power and limitations in isolation. In Section 4, we present our analyses of the ageing operator in isolation and highlight its ability to escape from local optima. In Section 5, we present the results of our analyses of the complete algorithm. In Section 6, we extend the analyses to include the genotype diversity mechanism as applied in the standard Opt-IA [17, 6]. Finally, we conclude the paper with a discussion of the results and directions for future work.

2. Preliminaries

In this section we first present the standard Opt-IA as applied in [6] (called Opt-IA* from now on) for the maximisation of $f : \{0, 1\}^n \rightarrow \mathbb{R}$ and then a slightly different version which we will analyse.

The Opt-IA* pseudo-code is given in Algorithm 1. It is initialised with a population of μ b-cells, representing candidate solutions, generated uniformly at random with $age = 0$. In each generation, the algorithm creates a new parent population consisting of dup copies of each b-cell (i.e., Cloning) which will be the subject of variation. The pseudo-code of the Cloning operator is given in Algorithm 2.

The variation stage in Opt-IA* uses a *hypermutation operator with mutation potential* sometimes followed by *hypermacromutation* [6], sometimes not [17]. The Hypermacromutation operator is essentially the same as the well-studied contiguous somatic mutation operator of the B-Cell algorithm; it chooses two integers i and j at random such that $(i + 1) \leq j \leq n$, then mutates at most $j - i + 1$ values in the range of $[i, j]$. If both operators are applied, they act on the clone population (i.e., not in sequence) such that they generate μ mutants each. The number of bits M that are flipped by the hypermutation operator is determined by a function called *mutation potential*. Three different potentials have been considered in the literature: *static*, where the number of bits that are flipped is linear in the problem size and does not depend on the fitness function³, *fitness proportional* (i.e., a linear number of bits are always flipped but increasing proportionally with the fitness of the mutated b-cell) and *inversely fitness proportional*. The latter potential was previously theoretically analysed in [22]. What is unclear from the literature is whether the M bits to be flipped should be distinct or not and, when using FCM, whether a *constructive mutation* is a strictly improving move or whether a solution of equal fitness suffices. In this paper we will consider the static hypermutation operator with pseudo-code given in Algorithm 3. In particular, the M flipped bits will always be distinct and both kinds of constructive mutations will be considered. At the end of the variation stage all created individuals have $age = 0$ if their fitness is higher than that of their parent cell, otherwise they inherit their parent's age. Then the whole population (i.e., parents and offspring) undergoes the ageing process in which the age of each b-cell is increased by one. Additionally, the ageing operator removes old individuals. Three methods have been proposed in the literature for such an operator: *static ageing*, which deterministically removes all individuals who exceed age τ ; *stochastic ageing*, which removes each individual at each generation with probability p_{die} ; and the recently introduced *hybrid ageing* [14], where individuals have a probability p_{die} of dying only once they reach an age of τ . In [14] it was shown that the hybrid version allows to escape local optima, hence we employ this version in this paper and give its pseudo-code in Algorithm 4.

The generation ends with a selection phase for which the pseudo-code is given in Algorithm 5. If the total number of b-cells that have survived the ageing operator is larger than μ , then a standard $(\mu + \lambda)$ selection scheme is used with the exception that genotype duplicates are not allowed. If the population size is less than μ , then a birth phase fills the population up to size μ by introducing random b-cells of $age = 0$.

In this paper we give evidence that disallowing genotypic duplicates may be detrimental because, as we will show, genotypic copies may help the ageing operator to escape local optima more efficiently. Considering the results of the investigations made in Sections 3 and 4, the Opt-IA we will analyse does not apply hypermacromutation and uses static hypermutation coupled with FCM as variation operator, hybrid ageing and standard $(\mu + \lambda)$ selection which allows genotypic duplicates. The pseudo-code is given in Algorithm 6 for clarity.

³In [17] the mutation potential is declared to be a constant $0 < c \leq 1$. This is obviously a typo: the authors intended the mutation potential to be cn , where $0 < c \leq 1$.

Algorithm 3 Static hypermutation ($P^{(clo)}, c$)

```
1:  $M = cn$ ,
2:  $P^{(hyp)} = \emptyset$ .
3: for all  $x_i \in P^{(clo)}$  do
4:   if FCM is not used then
5:     create  $y_i$  by flipping  $M$  distinct bits selected uniformly at random.
6:   else
7:     create  $y_i$  by flipping at most  $M$  distinct bits selected uniformly at random one after
       another until a constructive mutation happens.
8:   end if
9:   if  $f(y_i) > f(x_i)$  then
10:     $y_i^{age} = 0$ .
11:   else
12:     $y_i^{age} = x_i^{age}$ .
13:   end if
14:   add  $y_i$  to  $P^{(hyp)}$ .
15: end for
```

Algorithm 4 Hybrid ageing ($P^{(t)}, P^{(hyp)}, \tau, \mu$)

```
1: for all  $x_i \in (P^{(t)} \cup P^{(hyp)})$  do
2:    $x_i^{age} = x_i^{age} + 1$ ,
3:   if  $x_i^{age} > \tau$  then
4:     remove  $x_i$  with probability  $p_{die} = 1 - 1/\mu$ .
5:   end if
6: end for
```

3. Static Hypermutation

The aim of this section is to highlight the power and limitations of the static hypermutation operator in isolation. For this purpose we embed the operator into a minimal AIS framework that uses a population of only one b-cell and creates exactly one clone per generation. The resulting (1+1) IA^{hyp}, depicted in Algorithm 7, is essentially a (1+1) EA that applies the static hypermutation operator instead of using SBM. We will first show that, without the use of FCM, hypermutations are inefficient variation operators for virtually any optimisation function of interest. From there on we will only consider the operator equipped with FCM. Then we will prove that the (1+1) IA^{hyp} has a runtime that is at most a linear factor larger than that obtained for any RLS algorithm using the artificial fitness levels method. If no improvement is found in the first step, then the operator will perform at most cn useless fitness function evaluations before one hypermutation process is concluded. We formalise this result in Theorem 2 for two cases: when FCM only accepts strict improvements as constructive solutions (we formally call such algorithm (1+1) IA_>^{hyp}) and for the case when FCM also accepts points of equal fitness as constructive solutions (we

Algorithm 5 Selection ($P^{(t)}, P^{(hyp)}, \mu, div$)

- 1: $P^{(t+1)} = (P^{(t)} \cup P^{(hyp)})$.
- 2: **if** $div=1$ **then**
- 3: remove any offspring with the same genotype as individuals in $P^{(t)}$.
- 4: **end if**
- 5: **if** $|P^{(t+1)}| > \mu$ **then**
- 6: remove the $(|P^{(t+1)}| - \mu)$ individuals with the lowest fitness breaking ties uniformly at random.
- 7: **end if**
- 8: **if** $|P^{(t+1)}| < \mu$ **then**
- 9: add $(\mu - |P^{(t+1)}|)$ individuals initialised uniformly at random.
- 10: **end if**

Algorithm 6 Opt-IA

- 1: $t = 0$,
- 2: initialise $P^{(t)} = \{x_1, \dots, x_\mu\}$, a population of μ b-cells uniformly at random and set $x_i^{age} = 0$ for $i = \{1, \dots, \mu\}$.
- 3: **while** the optimum is not found **do**
- 4: $P^{(clo)} = \text{Cloning}(P^{(t)}, dup)$,
- 5: $P^{(hyp)} = \text{Static hypermutation}(P^{(clo)}, c)$,
- 6: Hybrid ageing ($P^{(t)}, P^{(hyp)}, \mu, \tau$),
- 7: Selection ($P^{(t)}, P^{(hyp)}, \mu, div$),
- 8: $t = t + 1$.
- 9: **end while**

name such algorithm (1+1) IA $_{\geq}^{hyp}$).

We prove in Theorem 2 that the (1+1) IA $_{\geq}^{hyp}$ cannot be too slow compared to the standard RLS₁ (i.e., flipping one bit per iteration). We show that the presented results are tight for some standard benchmark functions by proving that the (1+1) IA hyp has expected runtimes of $\Theta(n^2 \log n)$ for ONEMAX and $\Theta(n^3)$ for LEADINGONES, respectively, versus the expected $\Theta(n \log n)$ and $\Theta(n^2)$ fitness function evaluations required by RLS₁ [20]. Nevertheless, we conclude the section by showing for the standard benchmark functions JUMP_k and CLIFF_d that the (1+1) IA hyp can be particularly efficient on functions with local optima that are generally difficult to escape from.

We start by highlighting the limitations of static hypermutation when FCM is not used. Since $M = cn$ distinct bits have to be flipped at once, the outcome of the hypermutation operator is characterised by a uniform distribution over the set of all solutions which have Hamming distance M to the parent. Since M is linear in n , the size of this set of points is exponentially large and thus the probability of a particular outcome is exponentially small. In the following theorem, we formalise this limitation.

Theorem 1. *For any function with a polynomial number of optima, the (1+1) IA hyp without*

Algorithm 7 (1+1) IA^{hyp}

```
1:  $t = 0$ ,
2: initialise a solution uniformly at random and assign it to  $P^{(0)}$ .
3: while an optimum is not found do
4:   Cloning ( $P^{(t)}$ , 1),
5:   Static hypermutation ( $P^{(clo)}$ ,  $c$ ),
6:   if  $f(P^{(hyp)}) \geq f(P^{(t)})$  then
7:      $P^{(t+1)} = P^{(hyp)}$ .
8:   else
9:      $P^{(t+1)} = P^{(t)}$ .
10:  end if
11:   $t = t + 1$ .
12: end while
```

FCM needs expected exponential time to find any of the optima.

Proof. We will first consider the probability that the initial solution is optimal and then the probability of finding an optimal solution in a single step given that the current solution is suboptimal. Note that if $c = 1$, sampling the complementary bit string of an optimal solution allows the hypermutation to flip all bits in the next iteration and find an optimal solution. However, since each optimal solution has a single complementary bit string, the total number of solutions which are either optimal or complementary to an optimal solution is also polynomial in n . Since the initial solution is sampled uniformly at random among 2^n possible bit strings, the probability that one of these polynomially many solutions is sampled is $poly(n)/2^n = 2^{-\Omega(n)}$.

We now analyse the expected time of the last step before an optimal solution is found given that the current solution is neither an optimal solution nor the complementary bit string of an optimal solution. When $c = 1$ the probability of finding the optima is zero since the hypermutation deterministically samples the complementary bit strings of current b-cells. For $c < 1$, we optimistically assume that all the optima are at Hamming distance $M = cn$ from the current b-cell. Otherwise, if none of the optima are at Hamming distance M , the probability of reaching an optimum would be zero. Then, given that the number of different points at Hamming distance cn from any point is $\binom{n}{cn}$ and they all are reachable with equal probability, the probability of finding this optimum in the last step, for any $c \neq 1$, is $p \leq \frac{poly(n)}{\binom{n}{cn}} \leq \frac{poly(n)}{e^{\Omega(n)}} = e^{-\Omega(n)}$. By a simple waiting time argument, the expected time to reach any optimum is at least $e^{\Omega(n)}$. \square

The theorem explains why poor results were achieved in previous experimental work both on benchmark functions and real world applications such as the hard protein folding problem [17, 6]. The authors indeed state that “*With this policy, however, and for the problems which are faced in this paper, the implemented IA did not provide good results*” [6]. Theorem 1 shows that this is the case for any optimisation function of practical interest. In [22] it had already been shown that inversely proportional hypermutations cannot optimise

ONEMAX in less than exponential time (both in expectation and w.o.p.). Although static hypermutations are the focus of this paper, we point out that Theorem 1 can easily be extended to both the inversely proportional hypermutations considered in [22] and to the proportional hypermutations from the literature [17]. From now on we will only consider hypermutations coupled with FCM. We start by showing that hypermutations cannot be too slow compared to local search operators. We first state and prove the following helper lemma.

Lemma 1. *The probability that the static hypermutation applied to $x \in \{0, 1\}^n$ either evaluates a specific $y \in \{0, 1\}^n$ with Hamming distance $k \leq cn$ to x (i.e., event E_y), or that it stops earlier on a constructive mutation (i.e., event E_c) is lower bounded by $\binom{n}{k}^{-1}$ (i.e., $\Pr\{E_y \vee E_c\} \geq \binom{n}{k}^{-1}$). Moreover, if there are no constructive mutations with Hamming distance smaller than k , then $\Pr\{E_y\} = \binom{n}{k}^{-1}$.*

Proof. Since the bits to be flipped are picked without replacement, each successive bit-flip increases the Hamming distance between the current solution and the original solution by one. The lower bound is based on the fact that the first k bit positions to be flipped have $\binom{n}{k}$ different and equally probable outcomes. Since the only event that can prevent the static hypermutation to evaluate the k -th solution in the sequence is the discovery of a constructive mutation in one of the first $k - 1$ evaluations, $\Pr\{E_y \vee E_c\}$ is at least $\binom{n}{k}^{-1}$. If no such constructive mutation exists (i.e., $\Pr\{E_c\} = 0$) then, $\Pr\{E_y\}$ is exactly equal to $\binom{n}{k}^{-1}$. \square

We are ready to show that static hypermutations cannot be too slow compared to any upper bound obtained by applying the AFL method on the expected runtime of the $(1+1)$ RLS $_k$, which flips exactly k bits to produce a new solution and applies non-strict elitist selection. AFL requires a partition of the search space \mathcal{X} into m mutually exclusive sets $\bigcup_{s \in \{1, \dots, m\}} A_s = \mathcal{X}$ such that $\forall i < j, x \in A_i \wedge y \in A_j \implies f(x) < f(y)$. The expected runtime of a $(1 + 1)$ algorithm \mathcal{A} with variation operator $\text{HM}(x) : \mathcal{X} \rightarrow \mathcal{X}$ to solve any function defined on \mathcal{X} can be upper bounded via AFL by $E(T) \leq \sum_{i=1}^m \frac{1}{p_i}$, where

$$p_i = \min_{x \in A_i} \left(\Pr\{\text{HM}(x) \in \bigcup_{j=i+1}^m A_j\} \right) [20, 21].$$

Theorem 2. *Let $E(T_A^{\text{AFL}})$ be any upper bound on the expected runtime of algorithm A established via the artificial fitness levels method. Then, $E\left(T_{(1+1) \text{ I}A_{>}^{\text{hyp}}}\right) \leq cn \cdot E\left(T_{(1+1) \text{ RLS}_k}^{\text{AFL}}\right)$. Additionally, for the special case of $k = 1$, $E\left(T_{(1+1) \text{ I}A_{\geq}^{\text{hyp}}}\right) \leq cn \cdot E\left(T_{(1+1) \text{ RLS}_1}^{\text{AFL}}\right)$.*

Proof. Let the function $c_k(x)$ for solution $x \in A_i$ return the number of solutions which are at Hamming distance k away from x and belong to set A_j for some $j > i$. The upper bound on the expected runtime of the $(1+1)$ RLS $_k$ to solve any function obtained by applying the AFL method is $E\left(T_{(1+1) \text{ RLS}_k}^{\text{AFL}}\right) \leq \sum_{i=1}^m \frac{1}{p_i}$, where $p_i = \min_{x \in A_i} (c_k(x) / \binom{n}{k})$. Since the hypermutation operator wastes at most cn bit mutations when it fails to improve, to prove

the first claim it is sufficient to show that for any current solution $x \in A_i$, the probability that the $(1+1)$ $IA_{>}^{\text{hyp}}$ finds an improvement is at least $c_k(x)/\binom{n}{k}$. This follows from Lemma 1 and the definition of a constructive mutation for the $(1+1)$ $IA_{>}^{\text{hyp}}$, since for each one of the $c_k(x)$ search points, the probability of either finding it or finding a constructive mutation is lower bounded by $\binom{n}{k}^{-1}$.

Note that, for the $(1+1)$ IA_{\geq}^{hyp} , the probability of improving the current solution $x \in A_i$ can be smaller than $c_k(x)/\binom{n}{k}$ since it is also necessary that the first $k-1$ sampled solutions have strictly worse fitness than x to allow the hypermutation operator to flip at least k bits before stopping. However, for the special case of $k=1$, the hypermutation cannot be prematurely stopped. The probability that static hypermutation produces a particular Hamming neighbour of the input solution in the first mutation step is $1/n$, which is equal to the probability that the RLS_1 produces the same solution. Considering the fact that static hypermutation wastes at most cn fitness evaluations in every failure to obtain an improvement in the first step, the second claim follows. \square

In the following we show that the upper bounds of the previous theorem are tight for well-known benchmark functions.

Theorem 3. *The expected runtime of the $(1+1)$ $IA_{>}^{\text{hyp}}$ and of the $(1+1)$ IA_{\geq}^{hyp} to optimise $\text{ONEMAX}(x) := \sum_{i=1}^n x_i$ is $\Theta(n^2 \log n)$.*

Proof. The upper bounds for the $>$ and \geq FCM selection versions follow from Theorem 2 since it is easy to derive an upper bound of $O(n \log n)$ for RLS using AFL [20]. For the lower bound of the $>$ FCM selection version, we follow the analysis in Theorem 3 of [22] for inversely proportional hypermutation (IPH) with FCM to optimise ZEROMIN . The proof there relies on IPH wasting cn function evaluations every time it fails to find an improvement. This is obviously also true for static hypermutations (albeit for a different constant c), hence the proof also applies to our algorithm. The lower bound also holds for the \geq FCM selection algorithm as this algorithm cannot be faster on ONEMAX by accepting solutions of equal fitness (i.e., the probability of finding an improved solution does not increase). \square

We now turn to the LEADINGONES benchmark function, which simply returns the number of consecutive 1-bits before the first 0-bit.

Theorem 4. *The expected runtime of the $(1+1)$ IA_{\geq}^{hyp} on $\text{LEADINGONES} := \sum_{i=1}^n \prod_{j=1}^i x_j$ is $\Theta(n^3)$.*

Proof. The upper bound is implied by Theorem 2 because AFL gives an $O(n^2)$ runtime of RLS for LEADINGONES [20]. Let $E(f_i)$ be the expected number of fitness function evaluations until an improvement is made, considering that the initial solution has i leading 1-bits. The initial solutions consist of i leading 1-bits, followed by a 0-bit and $n-i-1$ bits which each can be either one or zero with equal probability. Let events E_1 , E_2 and E_3 be that the first mutation step flips one of the leading ones (with probability i/n), the first 0-bit (with probability $1/n$) or any of the remaining bits (with probability $(n-i-1)/n$), respectively. If E_1 occurs, then the following mutation steps cannot reach any solution with




Figure 1: The JUMP_k function for $k = 10$ and $n = 50$.

fitness value i or higher and all cn mutation steps are executed. Since no improvements have been achieved, the remaining expected number of evaluations will be the same as the initial expectation (i.e., $E(f_i|E_1) = cn + E(f_i)$). If E_2 occurs, then a new solution with higher fitness value is acquired and the mutation process stops (i.e., $E(f_i|E_2) = 1$). However if E_3 occurs, since the number of leading 1-bits in the new solution is i , the hypermutation operator stops without any improvement (i.e., $E(f_i|E_3) = 1 + E(f_i)$). According to the law of total expectation: $E(f_i) = \frac{i}{n}(cn + E(f_i)) + \frac{1}{n} \cdot 1 + \frac{n-i-1}{n}(1 + E(f_i))$. When, this equation is solved for $E(f_i)$, we obtain, $E(f_i) = icn + n - i$. Since the expected number of consecutive 1-bits that follow the leftmost 0-bit is less than two [29], the probability of not skipping a level i is $\Omega(1)$. The initial solution on the other hand will have more than $n/2$ leading ones with probability at most $2^{-n/2}$. Thus, we obtain a lower bound on the expectation, $(1 - 2^{-n/2}) \sum_{i=n/2}^n (f_i) = \Omega(1) \sum_{i=n/2}^n (icn + n - i) = \Omega(n^3)$ by summing over fitness levels starting from level $i = n/2$. \square

We now focus on establishing that hypermutations may produce considerable speed-ups if local optima need to be overcome. The JUMP_k function, introduced in [29], consists of a ONEMAX slope with a gap of length k bits that needs to be overcome for the optimum to be found. The function is formally defined as:

$$\text{JUMP}_k(x) := \begin{cases} k + \sum_{i=1}^n x_i & \text{if } \sum_{i=1}^n x_i \leq n - k \\ & \text{or } \sum_{i=1}^n x_i = n, \\ n - \sum_{i=1}^n x_i & \text{otherwise,} \end{cases}$$

for $n > 1$ and $k \in \{1 \dots n\}$. Figure 1 illustrates this function.

Mutation-based EAs require $\Theta(n^k)$ expected function evaluations to optimise JUMP_k and recently a faster upper bound by a linear factor has been proved for standard crossover-based steady-state GAs [30]. Hence, EAs require increasing runtimes as the length of the gap increases, from superpolynomial to exponential as soon as $k = \omega(1)$. The following theorem




Figure 2: The CLIFF_d function for $d = 10$ and $n = 50$.

shows that hypermutations allow speed-ups by an exponential factor of $(e/k)^k$, when the jump is hard to perform. A similar result has been shown for the recently introduced fast-GA [31].

Theorem 5. *Let $cn > k$. Then the expected runtime of the $(1+1)$ IA^{hyp} to optimise JUMP_k is at most $O\left(\frac{n^{k+1} \cdot e^k}{k^k}\right)$.*

Proof. The $(1+1)$ IA^{hyp} reaches the fitness level $n - k$ (i.e., local optima) in $O(n^2 \log n)$ steps according to Theorem 3. All local optima have Hamming distance k to the optimum and the probability that static hypermutation finds the optimum is lower bounded in Lemma 1 by $\binom{n}{k}^{-1}$. Hence, the total expected time to find the optimum is at most $O(n^2 \log n) + cn \cdot \binom{n}{k} = O\left(\frac{n^{k+1} \cdot e^k}{k^k}\right)$. \square

Obviously, hypermutations can jump over large fitness valleys also on functions with other characteristics. For instance the CLIFF_d function was originally introduced to show when non-elitist EAs may outperform elitist ones [32]. This function is formally defined as follows:

$$\text{CLIFF}_d(x) = \begin{cases} \sum_{i=1}^n x_i & \text{if } \sum_{i=1}^n x_i \leq n - d, \\ \sum_{i=1}^n x_i - d + 1/2 & \text{otherwise.} \end{cases}$$

Figure 2 shows an illustration of CLIFF_d . Similarly to JUMP_k , this function has a ONE-MAX slope with a gap of length d bits that needs to be overcome for the optimum to be found. Differently to JUMP_k though, the local optimum is followed by another ONE-MAX slope leading to the optimum. Hence, algorithms that accept a move jumping to the bottom of the cliff, can then optimise the following slope and reach the optimum. While elitist mutation-based EAs obviously have a runtime of $\Theta(n^d)$ (i.e., they do not accept the jump to the bottom of the cliff), the following corollary shows how hypermutations lead to speed-ups that increase exponentially with the distance d between the cliff and the optimum.

Corollary 1. *Let $cn > d$. Then the expected runtime of the $(1+1)$ IA^{hyp} to optimise CLIFF_d is $O\left(\frac{n^{d+1} \cdot e^d}{d^d}\right)$.*

Algorithm 8 $(\mu+1)$ RLS $_p^{\text{ageing}}$

- 1: $t = 0$,
 - 2: initialise $P^{(t)} = \{x_1, \dots, x_\mu\}$, a population of μ individuals uniformly at random and set $x_i^{\text{age}} = 0$ for $i = \{1, \dots, \mu\}$.
 - 3: **while** the optimum is not found **do**
 - 4: select $x \in P^{(t)}$ uniformly at random,
 - 5: with probability $1/2 < 1 - p \leq 1$ create y by flipping one bit of x , otherwise $y = x$,
 - 6: Hybrid ageing $(P^{(t)}, \{y\}, \tau, \mu)$,
 - 7: Selection $(P^{(t)}, \{y\}, \mu, 0)$,
 - 8: $t = t + 1$.
 - 9: **end while**
-

The analysis can also be extended to show an equivalent speed-up compared to the (1+1) EA for crossing the fitness valleys of arbitrary length and depth recently introduced in [33]. In the next section we will prove that the ageing operator can lead to surprising speed-ups for CLIFF $_d$ and other functions with similar characteristics.

4. Ageing

It is well-understood that the ageing operator can allow algorithms to escape from local optima. This effect was shown on the BALANCE function from dynamic optimisation for an RLS algorithm embedding a hybrid ageing operator [14]. However, for that specific function, an SBM operator would fail to escape, due to the large basin of attraction of the local optima. In this section we highlight the capabilities of ageing in a more general setting (i.e., the standard CLIFF $_d$ benchmark function) and show that ageing may also be efficient when coupled with SBM.

Ageing allows to escape from a local optimum if one not locally optimal b-cell is created and survives while all the other b-cells die. For this to happen, it is necessary that all the b-cells are old and have similar age. This is achieved on a local optimum by creating copies of the locally optimal b-cell (the b-cells will inherit the age of their parent). Hence, the ability of a mutation operator to create copies enhances this particular capability of the ageing operator. To this end we first consider a modified RLS algorithm that with some constant probability p does not flip any bit and implements the ageing operator presented in Algorithm 4. We call this algorithm $(\mu+1)$ RLS $_p^{\text{ageing}}$ and present its pseudo-code in Algorithm 8. Apart from making ageing more effective, this slight modification to the standard RLS considerably simplifies the proofs of the statements we wish to make. In Section 6 we will generalise the result to an RLS algorithm that does not allow genotype duplicates as in the standard Opt-IA.

The CLIFF $_d$ benchmark function is generally used to highlight circumstances when non-elitist EAs outperform elitist ones. Algorithms that accept inferior solutions can be efficient for the function by jumping to the bottom of the cliff and then optimising the ONEMAX slope. This effect was originally shown for the $(1, \lambda)$ EA that can optimise the function

in approximately $O(n^{25})$ fitness function evaluations if the population size λ is neither too large nor too small [32]. This makes the difference between polynomial and exponential expected runtimes compared to elitist EAs (i.e., $\Theta(n^d)$) if the cliff is located far away from the optimum. A smaller, but still exponential, speed-up was recently shown for the population-genetics-inspired SSWM (Strong Selection Weak Mutation) algorithm with runtime of at most $n^d/e^{\Omega(d)}$ [34]. The following theorem proves a surprising result for the considered $(\mu+1)$ RLS_{*p*}^{ageing} for CLIFF_{*d*}. Not only is the algorithm very fast, but our upper bound becomes lowest (i.e., $O(n \log n)$) when the function is most difficult (i.e., when the cliff is located at distance $d = \Theta(n)$ from the optimum). In this setting the algorithm is asymptotically as fast as any evolutionary algorithm using SBM can be on any function with unique optimum [35]. A similar result has recently been shown also for a hyperheuristic which switches between elitist and non-elitist selection operators [36].

Theorem 6. *For $\mu = O(\log n)$, $p < 1$ a constant and $\tau = \Theta(n \log n)$, the $(\mu+1)$ RLS_{*p*}^{ageing} optimises CLIFF_{*d*} in expected time $O\left(\frac{\mu^2 n^3 \log n}{d^2}\right)$ if $d < n/4 - \epsilon n$ for any constant $0 < \epsilon < 1/4$.*

Proof. We follow the proof of Theorem 10 in [14] of the $(\mu + 1)$ RLS for the BALANCE function and adapt the arguments therein to the ONEMAX landscape and to the RLS operator we use. Given that there are $i < \mu$ individuals with j 1-bits in the population, the probability of creating a new individual with j 1-bits is at least $(i/\mu)p$ because the RLS operator creates a copy with a constant probability p . Hence we follow the proof in [14] to show that in $O(\mu n + n \log n)$ expected steps the population climbs up the ONEMAX slope (i.e., samples a solution with $n - d$ 1-bits) and subsequently the whole population will be taken over by the local optimum. Given that there are already k copies of the best individual in the population, the probability that one is selected for mutation is k/μ and the probability that a new copy is added to the population is kp/μ . Thus in at most $\sum_{k=1}^{\mu} \mu/kp = O(\mu \log \mu)$ generations in expectation after the first local optima is sampled, the population is taken over by the local optima. Now we can apply Lemma 5 in [14] to show that in expected $O(\mu^3)$ steps the whole population will have the same age. As a result, after another at most $\tau = \Theta(n \log n)$ generations the whole population will reach age τ simultaneously because no improvements may occur unless the optimum is found. Overall, the total expected time until the population consists only of local optima with age τ is at most $O(\mu n + n \log n + \mu \log \mu + \mu^3 + \tau) = O(\mu n + n \log n)$. Now we calculate the probability that in the next step one individual jumps to the bottom of the cliff and the rest die in the same generation. The first event happens with probability $(1 - p)(d/n) = \Omega(d/n)$ (i.e., an offspring solution with $n - d + 1$ 1-bits is created by flipping one of the d 0-bits in the parent solution). The probability that the rest of the population dies is $1/\mu \cdot (1 - 1/\mu)^\mu$. We now require that the survivor creates an offspring with higher fitness (i.e., with *age* = 0) by flipping one of its 0-bits (i.e., it climbs one step up the second slope). This event happens with probability at least $(1 - p)(d - 1)/(\mu n) = \Omega\left(\frac{d}{\mu n}\right)$ and in the same generation with probability $(1 - 1/\mu) = \Omega(1)$ the parent of age $\tau + 1$ dies due to ageing. Finally, the new solution (i.e., the *safe* individual) takes over the population in $O(\mu \log \mu)$ expected generations by standard arguments. It takes at least $\Omega(n)$ generations before any of the new

random individuals has ϵn more 0-bits than it had after initialisation since FCM terminates after each improvement. Using Markov's inequality, we can show that the probability that the takeover happens after more than $\Omega(n)$ steps is at most $O((\mu \log \mu)/n)$.

Hence, the overall probability of this series of consecutive events is $\Omega\left(\frac{d}{n}\right) \cdot \frac{1}{\mu} \left(1 - \frac{1}{\mu}\right)^\mu \cdot \Omega\left(\frac{d}{\mu n}\right) \cdot \left(1 - \frac{1}{\mu}\right) \cdot \left(1 - O\left(\frac{\mu \log \mu}{n}\right)\right) = \Omega\left(\frac{d^2}{n^2 \mu^2}\right)$ and the expected number of trials (i.e., climb-ups and restarts) until we get a survivor which is safe at the bottom of the cliff is $O\left(\frac{n^2 \mu^2}{d^2}\right)$. Every time the set of events fails to happen, we wait for another $O(\mu n + n \log n)$ fitness evaluations until the population reaches a configuration where all individuals are locally optimal and have age τ . Once a safe individual has taken over the population, the expected time to find the global optimum will be at most $O(\mu n + n \log n)$. Overall, the total expected time to optimise CLIFF_d conditional on the best individual never dying when climbing up the slopes is $E(T_{total}) \leq O(\mu n + n \log n) \cdot O\left(\frac{n^2 \mu^2}{d^2}\right) + O(\mu n + n \log n) = O\left(\frac{\mu^2 n^3 \log n}{d^2}\right)$. Finally, we consider the probability that the best individual in the population never dies when climbing up the slopes due to ageing. After any higher fitness level is discovered, it takes $O(\mu \log \mu)$ generations in expectation and at most $n^{1/2}$ generations with overwhelming probability (w.o.p.⁴) until the whole population takes over the level. For the first $n - \log n$ levels, the probability of improving a solution is at least $\Omega(\log n/n)$ and the probability that this improvement does not happen in $\tau - n^{1/2} = \Omega(n \log n)$ generations is at most $(1 - \Omega(\log n/n))^{\Omega(n \log n)} = e^{-\Omega(\log^2 n)} = n^{-\Omega(\log n)}$. For the remaining fitness levels, the probability of reaching age τ before improving is similarly $(1 - \Omega(1/n))^{\Omega(n \log n)} = e^{-\Omega(\log n)} = n^{-\Omega(1)}$. By the union bound over all levels, the probability that the best solution is never lost due to ageing is at least $1 - o(1)$. We pessimistically assume that the whole optimisation process has to restart if the best individual reaches age τ . However, since at most $1/(1 - o(1)) = O(1)$ restarts are necessary in expectation, our bound on the expected runtime holds. \square

We conclude the section by considering the $(\mu+1)$ $\text{EA}^{\text{ageing}}$ which differs from the $(\mu+1)$ $\text{RLS}_p^{\text{ageing}}$ by using SBM with mutation rate $1/n$ instead of flipping exactly one bit. SBM allows copying individuals but, since it is a global operator, it can jump back to the local optima from anywhere in the search space with non-zero probability. Nevertheless, the following theorem shows that, for not too large populations, the algorithm is still very efficient when the cliff is at linear distance from the optimum. Its proof follows similar arguments to those of Theorem 6. The main difference in the analysis is that it has to be shown that once the solutions have jumped to the bottom of the cliff, they have a good probability of reaching the optimum before jumping back to the top of the cliff.

Theorem 7. *The $(\mu+1)$ $\text{EA}^{\text{ageing}}$ optimises CLIFF_d in expected $O(n^{1+\epsilon})$ time if $d = (1/4)(1-c)n$ for some constant $0 < c < 1$, $\tau = \Theta(n \log n)$ and $\mu = \Theta(1)$, where ϵ is an arbitrarily small positive constant.*

⁴In the rest of the paper we consider events to occur “with overwhelming probability” (w.o.p.) meaning that they occur with probability at least $1 - 2^{-\Omega(n)}$.

Proof. The probability that the SBM operator increases the number of 1-bits in a parent solution with $\Omega(n)$ 0-bits is at least $\Omega(n)/(ne) = \Omega(1)$. Following the same arguments as in the proof of Theorem 6 while considering $n/4 > d = \Theta(n)$ and $\mu = O(1)$ we can show that with constant probability and in expected $O(n \log n)$ time the algorithm reaches a configuration where there is a single individual at the bottom of the cliff with $n - d + 2$ 1-bits and age zero while the rest of the population consists of solutions which have been randomly sampled in the previous iteration.

We will now show that the newly generated individuals have worse fitness than the solution at the bottom of the cliff when they are initialised. Since $d = (1/4)(1 - c)n$, the fitness value of the solutions with more than $n - d$ 1-bits is at least $n - (1/4)(1 - c)n - (1/4)(1 - c)n = (n/2)(1 + c)$. Due to Chernoff bounds, the newly created individuals have less than $(n/2)(1 + (c/2)) < (n/2)(1 + c)$ 1-bits with overwhelming probability.

Next we prove that for any constant ϵ , there exist some positive constant c^* , such that the best solution at the bottom of the cliff (the leading solution) will be improved consecutively for $c^* \log n$ iterations with probability at least $n^{-\epsilon}$. The leading solution with i 0-bits is selected for mutation and SBM flips a single 0-bit with probability $p_i = (1/\mu) \cdot (i/n)(1 - n^{-1})^{n-1}$. With probability $\prod_{i=n-d+2}^{n-d+1+c^* \log n} p_i$, $c^* \log n$ consecutive improvements occur. We can bound this probability from below by the final improvement probability p_f raised to the power of $c^* \log n$ since the improvement probability is inversely proportional to the number of 0-bits. Considering that $p_f \geq (n - d + 1 + c^* \log n)/(n\mu e) = \Omega(1)$, we can set $c^* := -\epsilon \log_{p_f} 2 = \Omega(1)$, which yields $p_f^{c^* \log n} = n^{-\epsilon}$. Here, we note that immediately after $c^* \log n$ consecutive improvements, all the individuals in the population have at least $n - d + 1 + c^* \log n - \mu$ 1-bits. More precisely there will be one and only one individual in the population with j bits for all j in $[n - d + 1 + c^* \log n - \mu, n - d + 2 + c^* \log n]$. Since μ is constant and SBM flips at least k bits with probability $n^{-k} \binom{n}{k}$, the probability that a single operation of SBM decreases the number of 1-bits in any individual below $n - d + 1$ is in the order of $O(1/(\log n)!)$. Since this probability is not polynomially bounded, with probability at least $1 - O(1/n)$ it will not happen in any polynomial number of iterations. After the consecutive improvements occur, it takes $O(n \log n)$ until the second slope is climbed and the optimum is found.

The asymptotic bound on the runtime is obtained by considering that algorithm will reach the local optimum n^ϵ times in expectation before the consecutive improvements are observed. Since the time to restart after reaching the local optimum is in the order of $O(n \log n)$, the expected time is $O(n^{1+\epsilon} \log n)$. Since ϵ is an arbitrarily small constant, the order $O(n^{1+\epsilon} \log n)$ is equivalent to the order $O(n^{1+\epsilon})$. \square

5. Opt-IA

After having analysed the operators separately, in this section we consider the complete Opt-IA. The considered Opt-IA, shown in Algorithm 6, uses static hypermutation coupled with FCM as variation operator, hybrid ageing and an standard $(\mu + \lambda)$ selection which allows genotype duplicates. Also, a mutation is considered constructive if it results in creating an equally fit solution or a better one.

In this section we first show that Opt-IA is efficient for all the functions considered previously in the paper. Then, in Subsection 5.1 we present a problem where the use of the whole Opt-IA is crucial. In Subsection 5.2 we show limitations of Opt-IA by presenting a class of functions where standard EAs are efficient while Opt-IA is not w.o.p. We conclude the section with Subsection 5.3 where we present an analysis for trap functions which disproves previous claims in the literature about Opt-IA’s behaviour on this class of problems.

The following theorem proves that Opt-IA can optimise any benchmark function considered previously in this paper. The theorem uses that the ageing parameter τ is set large enough such that no individuals die with high probability before the optima are found.

Theorem 8. *Let τ be large enough. Then the following upper bounds on the expected runtime of Opt-IA hold:*

$$\begin{aligned} E(T_{\text{ONEMAX}}) &= O(\mu \cdot \text{dup} \cdot n^2 \log n) \text{ for } \tau = \Omega(n^2), \\ E(T_{\text{LEADINGONES}}) &= O(\mu \cdot \text{dup} \cdot n^3) \text{ for } \tau = \Omega(n^2), \\ E(T_{\text{JUMP}_k}) &= O\left(\mu \cdot \text{dup} \cdot \frac{n^{k+1} \cdot e^k}{k^k}\right) \text{ for } \tau = \Omega(n^{k+1}), \\ \text{and } E(T_{\text{CLIFF}_d}) &= O\left(\mu \cdot \text{dup} \cdot \frac{n^{d+1} \cdot e^d}{d^d}\right) \text{ for } \tau = \Omega(n^{d+1}). \end{aligned}$$

Proof. The claims use that if τ is large enough (i.e., $\Omega(n^2)$ for ONEMAX and LEADINGONES, $\Omega(n^{k+1})$ for JUMP_k and $\Omega(n^{d+1})$ for CLIFF_d), then with probability $1 - o(1)$ the current best solution will never reach age τ and die due to ageing before an improvement is found. For the Opt-IA to lose the current best solution due to ageing, it is necessary that the best solution is not improved for τ generations consecutively. If the improvement probability for any non-optimal solution is at least p_{\min} and if the age τ is set to be $p_{\min}^{-1}n$, then the probability that a solution will reach age τ before creating an offspring with higher fitness is at most $(1 - p_{\min})^\tau \leq e^{-n}$. By the union bound it is also exponentially unlikely that this occurs in a polynomial number of fitness levels that need to be traversed before reaching the optimum. Since the suggested τ for each function is larger than the corresponding $p_{\min}^{-1}n$, the upper bounds of the (1+1) IA^{hyp} (which does not implement ageing) for ONEMAX, LEADINGONES, JUMP_k and CLIFF_d are valid for Opt-IA when multiplied by $\mu \cdot \text{dup}$ to take into account the population and clones. \square

Theorem 10 shows that the presented upper bound for ONEMAX is tight for sub-logarithmic population and clone sizes (i.e., $\mu = o(\log n)$ and $\text{dup} = o(\log n)$). Before stating this theorem, we state the following helper theorem which was already used in [22] to prove lower bounds on the expected runtime of inversely proportional hypermutations.

Theorem 9 (Ballot Theorem [37]). *“Suppose that, in a ballot, candidate P scores p votes and candidate Q scores q votes, where $p > q$. The probability that throughout the counting there are always more votes for P than for Q equals $(p - q)/(p + q)$ ”.*

This theorem allows us to derive an upper bound on the probability that at some point of the hypermutation we have picked more 0-bits than 1-bits, which implies an improvement by hypermutation for ONEMAX.

Theorem 10. *Opt-IA needs at least $cn^2(\frac{\log(n/3)}{2} - \frac{\mu \cdot dup}{3})$ expected fitness function evaluations for any mutation potential c to optimise ONEMAX.*

Proof. By Chernoff bounds, the initial individuals have at least $i = n/3$ 0-bits with overwhelming probability. To calculate the probability of an improvement (i.e., flipping equal or more 0-bits than 1-bits during one mutation operation), we use the Ballot theorem (i.e., Theorem 9) in a similar way to [22]. Considering the number of 0-bits as $i = q$ and the number of 1-bits as $n - i = p$, the probability of an improvement is at most $1 - (p - q)/(p + q) = 1 - (n - 2i)/n = 2i/n$ according to the Ballot theorem⁵. Hence, the probability that at least one out of $dup \cdot \mu$ individuals succeeds is $P \leq dup \cdot \mu 2i/n$ by the union bound. We optimistically assume that the rest of the individuals also improve their fitness after such event happens. Recall that the mutation operator wastes cn fitness function evaluations every time it does not improve the fitness. Therefore, the expected time to see an improvement is $E(T_{improve}) \geq \left(\frac{n}{dup \cdot 2\mu i} - 1\right) \cdot \mu cn \cdot dup$. Since the mutation operator stops at the first constructive mutation (i.e., when the number of 1-bits is increased by one), it is necessary to improve at least $n/3$ times. So the total expected time to optimise ONEMAX is $E(T_{total}) \geq \sum_{i=1}^{n/3} E(T_{improve}) = cn^2(\frac{\log n/3}{2} - \frac{\mu \cdot dup}{3})$.

If individuals were to be removed because of ageing, then the new randomly generated individuals that replace them will have to improve at least $n/3$ times all over again w.o.p. Hence, the runtime may only increase in such an event. \square

5.1. Opt-IA Can Be More Efficient

In this section, we present the function HIDDENPATH to illustrate a problem where the use of static hypermutation and ageing together is crucial. When either of these two characteristic operators of Opt-IA is not used, we will prove that the expected runtime is at least superpolynomial. HIDDENPATH: $\{0, 1\}^n \rightarrow \mathbb{R}$ can be described by a series of modifications to the well-know ZEROMAX function. The distinguishing solutions are those with five 0-bits and those with $n - 1$ 0-bits, along with $\log n - 3$ solutions of the form $1^{n-k}0^k$ for $5 \leq k \leq \log n + 1$ (called SP points). The solutions with exactly $n - 1$ 0-bits constitute the local optima of HIDDENPATH (called LOCALOPT), and the solutions with exactly five 0-bits form a gradient with fitness increasing with more 0-bits in the rightmost five bit positions. Given that $|x|_0$ and $|x|_1$ respectively denote the number of 0-bits and 1-bits in a bit string x , HIDDENPATH is formally defined as in Definition 1. HIDDENPATH is illustrated in Figure 3 where OPT shows the global optimum.

Definition 1. *Given the definitions of SP and LOCALOPT as above, for any positive constant $\epsilon < 1$, the HIDDENPATH function is defined for all $x \in \{0, 1\}^n$ by*

⁵Like in [22] we consider that using $cn < n$ mutations, the probability can only be lower than the result stated in the Ballot theorem.

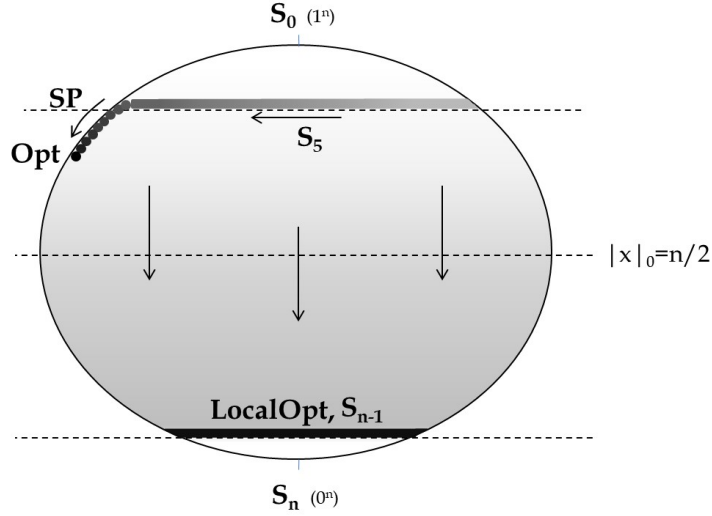


Figure 3: The HIDDENPATH function (Definition 1). The set S_i corresponds to the set of solutions with i 0-bits while the set SP refers to the set of solutions in the form $1^{n-k}0^k$ for $5 \leq k \leq \log n + 1$. Fitness increases with darker shades of gray.

$$\text{HIDDENPATH}(x) = \begin{cases} n - \epsilon + \frac{\sum_{i=n-4}^n (1-x_i)}{n} & \text{if } |x|_0 = 5 \text{ and } x \neq 1^{n-5}0^5, \\ 0 & \text{if } |x|_0 < 5 \text{ or } |x|_0 = n, \\ n - \epsilon + \epsilon k / \log n & \text{if } x \in \text{SP} := \{x \mid x = 1^{n-k}0^k \text{ and } 5 \leq k \leq \log n + 1\}, \\ n & \text{if } x \in \text{LOCALOPT} := \{x \mid |x|_0 = n - 1\}, \\ |x|_0 & \text{otherwise.} \end{cases}$$

Since the all 0-bits string returns fitness value zero, there is a drift towards solutions with $n - 1$ 0-bits while the global optimum (OPT) is the $1^{n-\log n-1}0^{\log n+1}$ bit string. The solutions with exactly five 0-bits work as a net that stops any static hypermutation that has an input solution with less than five 0-bits. The path to the global optimum consists of $\log n - 3$ Hamming neighbours and the first solution on this path has five 0-bits.

Theorem 11. For $c = 1$, $\text{dup} = 1$, $\mu = O(\log n)$ and $\tau = \Omega(n^2 \log n)$, Opt-IA needs expected $O(\tau\mu n + \mu n^{7/2})$ fitness function evaluations to optimise HIDDENPATH.

Proof. For convenience we will call any solution with i 0-bits (except the SP solutions) an S_i solution. After $O(n \log n)$ generations in expectation, an S_{n-1} solution is found by optimising ZEROMAX. Assuming the global optimum is not found first, consider the generation when an S_{n-1} solution is found for the first time. Another S_{n-1} solution is created and accepted by Opt-IA with probability at least $1/n$ since it is sufficient to flip the single 1-bit in the first mutation step and any 0-bit in the second step will be flipped next with probability 1. Thus, S_{n-1} solutions take over the population in expected $O(n\mu)$ generations. Since, apart from the optimum, no other solution has higher fitness than S_{n-1} solutions,

the population consists only of S_{n-1} solutions after the takeover occurs. A solution reaches age τ before the takeover only with probability $2^{\Omega(\sqrt{n})}$, due to Markov's inequality applied iteratively for $\Omega(\sqrt{n})$ consecutive phases of length $\Theta(n\mu)$. We now bound the expected time until the entire population has the same age. Considering that the probability of creating another S_{n-1} solution is $\Theta(1/n)$, the probability of creating two copies in one single generation is $\binom{\mu}{2}O\left(\frac{1}{n^2}\right) = O\left(\frac{\log^2 n}{n^2}\right)$. With constant probability this event does not happen in $o(n^2/\log^2 n)$ generations. Conditional on that at most one additional S_{n-1} solution is created in every generation, we can follow a similar argument as in the proof of Theorem 6. Hence, we can show that in expected $O(\mu^3 n)$ iterations after the takeover, the whole population reaches the same age. When the population of S_{n-1} solutions with the same age reaches age τ , with probability $1/\mu \cdot (1 - (1/\mu))^{2\mu-1} = \Theta(1/\mu)$ a single new clone survives while the rest of the population dies. With probability $1 - O(1/n)$ the survived clone has hypermutated all n bits (i.e., the survived clone is an S_1 solution). In the following generation, the population consists of an S_1 solution and $\mu - 1$ randomly sampled solutions. With probability 1, the S_1 solution produces an S_5 solution via hypermutation. On the other hand, with overwhelming probability the randomly sampled solutions still have fitness value $n/2 \pm O(\sqrt{n})$, hence the S_1 solution is removed from the population while the S_5 b-cell is kept. Overall, after $\Theta(\mu) + o(1)$ expected restarts an S_5 solution will be found in a total expected runtime of $O(\mu \cdot (n \log n + \mu n + \mu^3 n + \tau + 1)) = O(\mu\tau)$ generations. We momentarily ignore the event that the S_5 solution reaches an S_{n-1} point via hypermutation.

Now, the population consists of a single S_5 solution and $\mu - 1$ solutions with at most $n/2 + O(\sqrt{n})$ 0-bits. We want to bound the expected time until S_5 solutions take over the population conditional on no S_{n-1} solutions being created. Clones of any S_5 solutions are also S_5 solutions after hypermutation if one of the first two bits to be flipped is a 0 and the other is a 1, which happens with $O(1/n)$ probability. Moreover, if the outcome of hypermutation is neither an S_5 , an SP nor an S_{n-1} solution, then it is an S_{n-5} solution since all n bit-flips will have been executed. Since S_{n-5} solutions have higher fitness value than the randomly sampled solutions, they stay in the population. In the subsequent generation, if hypermutation does not improve an S_{n-5} solution (which happens with probability $O(1/n)$), it executes n bit-flips to create yet another S_5 solution unless the SP path is found.

This feedback causes the number of S_{n-5} and S_5 solutions to double in each generation with constant probability until they collectively take over the population in $O(\log \mu)$ generations in expectation. Then, with constant probability all the S_{n-5} solutions produce an S_5 solution via hypermutation and consequently the population consists only of S_5 solutions. Since the takeover happens in expected $O(\log \mu)$ generations, the probability that it fails to complete in $O(\log \mu)$ generations for consecutive $\Omega(\sqrt{n})$ times is exponentially small. Hence, in $O(\sqrt{n} \log \mu)$ generations w.o.p. (by applying Markov's inequality iteratively [20]) the entire population are S_5 solutions conditional on S_{n-4} solutions not being created before. Since the probability that the static hypermutation creates an S_{n-4} solution from an S_{n-5} solution is less than $(4/n) \cdot \mu$, and the probability that it happens in $O(\sqrt{n} \log \mu)$ generations is less than $(4/n) \cdot \mu \cdot O(\sqrt{n} \log \mu) = o(1)$, the takeover occurs with high probability, i.e., $1 - o(1)$.

After the whole population consists of only S_5 solutions, except for the global optimum and the local optima, only other points on the gradient would have higher fitness. The probability of improving on the gradient is at least $1/n^2$ which is the probability of choosing two specific bits to be flipped (a 0-bit and a 1-bit) leading towards the b-cell with the best fitness on the gradient. Considering that the total number of improvements on the gradient is at most 5, in $O(n^2 \cdot 5) = O(n^2)$ generations in expectation the first point of SP will be found. Now we consider the probability of jumping back to the local optima from the gradient, which is $\Theta(1)/\binom{n}{n-4} = O(n^{-4})$, before finding the first point of the SP. Due to the Markov's inequality applied iteratively over $\Omega(\sqrt{n})$ consecutive phases of length $\Theta(n^2)$ with an appropriate constant, the probability that the time to find the first point of SP is more than $\Omega(n^{5/2})$ is less than $2^{-\Omega(\sqrt{n})}$. The probability of jumping to a local optimum in $O(n^{5/2})$ steps is at most $n^{(-4)}O(n^{5/2}) = O(n^{-3/2})$ by the union bound. Therefore, SP is found before any local optima in at most $O(n^{5/2})$ generations with probability $1 - o(1)$. Hence, the previously excluded event has now been taken into account.

After $1^{n-5}0^5$ is added to the population, the best solution on the path is improved with probability $\Omega(1/n)$ by hypermutation and in expected $O(n \log n)$ generations the global optimum is found. Since all SP and S_5 solutions have a Hamming distance smaller than $n - 4$ and larger than $n/2$ to any S_{n-1} solution, the probability that a local optimum is found before the global optimum is at most $O(n \log n) \cdot \mu n \binom{n}{4}^{-1} = o(1)$ by the union bound. Thus with probability $1 - o(1)$ the time to find the optimum is $O(n \log n)$ generations. We pessimistically assume that we start over when this event does not occur, which implies that the whole process until that point should be repeated $1/(1 - o(1))$ times in expectation. Overall the dominating term in the runtime is $O(\tau + n^{5/2})$ generations. By multiplying this time with the maximum possible wasted fitness evaluations per generation (μcn), the upper bound is proven. \square

In the following two theorems we show that hypermutations and ageing used in conjunction are essential.

Theorem 12. *Opt-IA without ageing (i.e., $\tau = \infty$) with $\mu = O(\log n)$ and $dup = O(1)$ cannot optimise HIDDENPATH in less than $n^{\Omega(\log n)}$ expected fitness function evaluations.*

Proof. The only points with higher fitness than solutions with more than $n/2$ 0-bits are SP points and S_5 points. If all solutions in the population have less than $n - c_1 \log n$ 0-bits or less than $n - c_1 \log n$ 1-bits for some $c_1 > 1$, then the Hamming distance of any solution in the population to any SP solution or an S_5 solution is at least $\Omega(\log n)$. Since there are no other improving solution, the probability that an SP or an S_5 solution will be discovered in a single hypermutation operation is exponentially small according to the last statement of Lemma 1. As a result, conditional on not seeing these points, the algorithm will find a solution with $n - c_1 \log n$ 0-bits in $O(\mu \cdot dup \cdot n^2 \log n)$ fitness function evaluations by optimising ZEROMAX (i.e., Theorem 8). In the rest of the proof we will calculate the probability of reaching an S_{n-1} solution before finding either an SP point or S_5 point, or a complementary solution of an SP point. The latter solutions, if accepted, with high probability would hypermutate into SP points. We call an event *bad* where any of the mentioned points are discovered.

Any solution in the search space can have at most two Hamming neighbours on SP or at most two Hamming neighbours that their complementary bit strings are on SP. The probability of sampling any of these neighbours is in the order of $O(1/n)$ since it is necessary to flip a specific bit position. Hamming distance to the remaining SP points and their complementary bit strings are at least two. The probability of reaching a particular solution with Hamming distance at least two is at most $\binom{n}{2}^{-1} = O(1/n^2)$, since two particular bits need to be flipped. Excluding the two Hamming neighbours (which are sampled with probability $O(1/n)$), the probability of reaching any other SP point is $O(1/n^2) \cdot O(\log n) = O(\log n/n^2)$ by the union bound. Now, taking the Hamming neighbours into consideration, the probability that any of the $\log n - 3$ SP points (or their complementary bit strings) are discovered is $O(1/n) + O(1/n^2) \cdot O(\log n) = O(1/n)$. For any initial solution with less than $n - 10$ 0-bits, the probability of finding a solution with five 0-bits is at most $\binom{n}{6}^{-1} \binom{n}{5} = O(1/n)$. Since the probability of reaching an S_5 point from solutions with more than $n - 10$ 0-bits may be much larger, we first calculate the probability of finding $n - 11$ 0-bits before a bad event occurs.

The probability of finding a solution which improves the ZEROMAX value is at least $\Theta(1/n)$ (even when we exclude the potential improvements whose complementary bit strings are on SP). Since at every static hypermutation, the probabilities of finding an S_5 solution, an SP solution or its complementary bit string are all in the order of $O(1/n)$ (up to $n - 10$ 0-bits), the conditional probability that a solution which improves the current best ZEROMAX value is found before any other improvement is in the order of $\Theta(1/n)/(\Theta(1/n) + O(\text{dup} \cdot \mu/n)) = \Omega(1/(\text{dup} \cdot \mu))$. The probability that it happens $c_1 \log n$ times immediately after the first solution with more than $n - c_1 \log n$ 0-bits is added to the population is at least $\Omega(\text{dup} \cdot \mu)^{-c_1(\log n)} = (\text{dup} \cdot \mu)^{-O(\log n)}$. This sequence of $c_1 \log n$ improvements implies that a solution with $n - 11$ 0-bits is added to the population. We now consider the probability that one individual finds the local optimum (i.e., an S_{n-1} solution) by improving its ZEROMAX value in 10 consecutive generations and that none of the other individuals improve in these generations. The probability that the current best individual improves in the next step is at least $1/n$ and the probability that the other individuals do not improve is at least $(1 - c/n)^\mu \geq 1/e$. Hence, the probability that this happens consecutively in the next 10 steps is at least $(1/ne)^{10} = \Omega(n^{-10})$.

Once a locally optimal solution is found, with high probability it takes over the population before the optimum is found and the expected runtime conditional on the current population consisting only of solutions with $n - 1$ 0-bits is at least $\binom{n}{\log n} \geq \frac{(n - \log n)^{\log n}}{(\log n)^{\log n}}$. By the law of total expectation, the unconditional expected runtime is lower bounded by $E(T) \geq (1 - o(1)) n^{-10} (\text{dup} \cdot \mu)^{-O(\log n)} \frac{(n - \log n)^{\log n}}{(\log n)^{\log n}}$. This expression is in the order of $n^{\Omega(\log n)}$ for any $\mu = O(\log n)$ and $\text{dup} = O(1)$. \square

Theorem 13. *With Probability at least $1 - e^{-\Omega(n)}$, Opt-IA using SBM and ageing cannot optimise HIDDENPATH in less than $n^{\Omega(n)}$ fitness function evaluations.*

Proof. Since the number of 1-bits in an initial solution is binomially distributed with expectation $n/2$ the probability that an initial solution has less than $n/3$ 1-bits is bounded above

by $e^{-\Omega(n)}$ using Chernoff bounds. Hence, w.o.p., the population has a Hamming distance at least $n/3 - \log n - 1 > n/4$ from any solution on the path and any solution with five 0-bits ($\text{SP} \cup S_5$) by the union bound. Therefore, the probability of finding either any of the $\log n$ points on SP or one of the $\binom{n}{5}$ points on S_5 is $(\log n + \binom{n}{5}) \cdot 1/n^{n/4} (1 - 1/n)^{n-n/4} \leq n^{-\Omega(n)}$.

Since accepting any improvements, except for the $\text{SP} \cup S_5$ solutions, increases the distance to the $\text{SP} \cup S_5$ solutions, the probability of jumping to an $\text{SP} \cup S_5$ solution further decreases throughout the search process. \square

5.2. When Opt-IA is detrimental

In this section we present a function class for which Opt-IA is inefficient. The class of functions, which we call HYPERTRAP_γ , is defined formally in Definition 2. HYPERTRAP_γ is inspired by the SP-TARGET function introduced in [38] and used in [22] as an example where hypermutations outperform SBM. Compared to SP-TARGET , HYPERTRAP_γ has local and global optima inverted with the purpose of trapping hypermutations. Also there are some other necessary modifications to prevent the algorithm from finding the global optimum via large mutations. The parameter γ defines the distance between the local optima and the path to the global optimum.

In HYPERTRAP_γ , the solutions with $|x|_1 < n/2$ are evaluated by ONEMAX . Solutions of the form $1^i 0^{n-i}$ with $n/2 \leq i \leq n$, shape a short path which is called SP . The last point of SP , i.e., 1^n is the global optimum (OPT). Local optima of HYPERTRAP_γ (LOCALOPT) are formed by the points with $|x|_1 \geq 3n/4$ which have a Hamming distance of at least γn to all points in SP . Also, the points with $|x|_1 = n/2$ are ranked among each other such that bit strings with more 1-bits in the beginning have higher fitness. This ranking forms a gradient from $0^{n/2} 1^{n/2}$, which has the lowest fitness, to $1^{n/2-1} 0^{n/2+1}$, which has the highest fitness. Finally, the fitness of points which do not belong to any of the mentioned sub-spaces is evaluated by ZEROMAX (i.e., the fitness is $|x|_0$).

Definition 2. Given the definitions of SP , OPT and LOCALOPT as above and $H(x, \text{SP})$ showing the minimum Hamming distance of the individual x to all SP points, the HYPERTRAP_γ function with $0 < \gamma \leq 1/8$ is defined for all $x \in \{0, 1\}^n$ by

$$\text{HYPERTRAP}_\gamma(x) := \begin{cases} |x|_1 & \text{if } |x|_1 < n/2, \\ n/2 + \frac{\sum_{i=1}^n (n-i) \cdot (x_i)}{n} & \text{if } |x|_1 = n/2, \\ n^2 \cdot |x|_1 & \text{if } x \in \text{SP} := \{x \mid x = 1^i 0^{n-i} \text{ and } n/2 \leq i < n\}, \\ n^3 & \text{if } x \in \text{LOCALOPT} := \{x \mid |x|_1 \geq 3n/4 \text{ and } H(x, \text{SP}) \geq \gamma n\}, \\ n^4 & \text{if } x = 1^n = \text{OPT}, \\ |x|_0 & \text{if } |x|_1 > n/2 \text{ and } x \notin (\text{SP} \cup \text{LOCALOPT} \cup \text{OPT}). \end{cases}$$

This function is depicted in Figure 4. We will show that there exists a HYPERTRAP_γ such that Opt-IA with mutation potential cn gets trapped in the local optima.




Figure 4: HYPERTRAP_γ (Definition 2). SP corresponds to the points in the form of $1^i 0^{n-i}$ with $n/2 \leq i \leq n$. LocalOpt shows the locally optimal points which have $|x|_1 \geq 3n/4$ and Hamming distance of at least γn to all points in SP. Fitness increases with darker shades of gray.

Theorem 14. *With probability $1 - 2^{-\Omega(\sqrt{n})}$, Opt-IA with mutation potential cn cannot optimise $\text{HYPERTRAP}_{c/8}$ in less than exponential time.*

Proof. We will show that the population will first follow the ZEROMAX (or ONEMAX) gradient until it samples a solution with $n/2$ 0-bits and then the gradient of $\sum_{i=1}^n (n-i) \cdot (x_i)$ until it samples an SP point with approximately $n/2$ 0-bits. Afterwards, we will prove that large jumps on SP are unlikely. Finally, we will show that with overwhelming probability a locally optimal solution is sampled before a linear number of SP points are traversed. We optimistically assume that τ is large enough such that individuals do not die before finding the optimum.

With overwhelmingly high probability, all randomly initialised individuals have $(1/2 \pm \epsilon)n$ 0-bits for any arbitrarily small $\epsilon = \theta(1)$. Starting with such points, the probability of jumping to the global optimum is exponentially small according to Lemma 1. Being optimistic, we assume that the algorithm does not sample a locally optimal point for now. Hence, until an SP point is found, the current best fitness can only be improved if a point with either a higher ZEROMAX (or ONEMAX) or a higher $\sum_{i=1}^n (n-i) \cdot (x_i)$ value than the current best individual is sampled. Since there are less than n^2 different values of $\sum_{i=1}^n (n-i) \cdot (x_i)$, and less than n different values of ZEROMAX (or ONEMAX), the current best individual cannot be improved more than $n^2 + n$ times after initialisation without sampling an SP point.

Let SP_{lower} denote the SP points with less than $(1/2 + 2\epsilon)n$ 1-bits, and SP_{upper} denote $\text{SP} \setminus \text{SP}_{lower}$. We will now show that with overwhelmingly high probability an SP_{lower} point will be sampled before an SP_{upper} point. The search points between $(1/2 \pm \epsilon)n$ 1-bits have at least a distance of ϵn to the SP_{upper} points. Using Lemma 1, we can bound the probability of sampling an SP_{upper} by $\binom{n}{\epsilon n}^{-1}$ from any input solution with $(1/2 \pm \epsilon)n$ 1-bits. Excluding SP points, a solution with $(1/2 \pm \epsilon)n$ 1-bits has an improvement probability of at least $1/n^2$

(i.e., $\min\{1/n^2, \Theta(1)\} = 1/n^2$ with the second term being the probability of improving the ZEROMAX or ONEMAX value and the first term the probability of improving when the solution has exactly $n/2$ 0-bits). Thus, the conditional probability that the best solution is improved before any search point in the population is mutated into an SP_{upper} point is $(n^{-2})/(n^{-2} + dup \cdot \mu \cdot \binom{n}{cn}^{-1}) = 1 - 2^{-\Omega(n)}$. This implies that an SP_{lower} point is sampled before an SP_{upper} point with overwhelmingly high probability. Indeed, by the union bound, this event happens $n^2 + n$ times consecutively after initialisation with probability at least $1 - (n^2 + n) \cdot 2^{-\Omega(n)} = 1 - 2^{-\Omega(n)}$, thus a point in SP_{lower} is sampled before an SP_{upper} point with overwhelmingly high probability. Each point of SP has at least $n/2$ 1-bits at the beginning of its bit string. In order to improve by any value in one mutation operation, 1-bits should never be touched before the mutation operator stops. Hence, the probability of improving by X on SP is $p(X) < (1/2)^X$. This yields that the probability of improving by \sqrt{n} in one generation is at most $dup \cdot \mu \cdot (1/2)^{\sqrt{n}}$, as there are μ individuals in the population. We have now shown that it is exponentially unlikely that even an arbitrarily small fraction of the path points are avoided by jumping directly to path points with more 1-bits.

Let t be the first iteration when the current SP solution has at least $99n/100$ 1-bits for the first time. We will first lower bound the probability of improving in order to upper bound the conditional probability that more than \sqrt{n} new 1-bits are added given that an improving path point is sampled. An improvement is achieved if the 0-bit adjacent to the last 1-bit is flipped in the first mutation step, which happens with probability $1/n$. Thus, the conditional probability of improving more than $\Omega(\sqrt{n})$ points on SP is at most $2^{-\Omega(\sqrt{n})}/(1/n) = 2^{-\Omega(\sqrt{n})}$. Therefore, in generation t , the current best individual cannot have more than $(99n/100) + \Omega(\sqrt{n})$ 1-bits with probability $1 - dup \cdot \mu \cdot 2^{-\Omega(\sqrt{n})}$ by the union bound. Similarly, the probability that no jump of size $\Omega(\sqrt{n})$ happens in $poly(n)$ generations is at least $1 - dup \cdot \mu \cdot 2^{-\Omega(\sqrt{n})}$. So, we can conclude that the number of generations to add another $n/200$ 1-bits to the prefix of the current best solution is at least $(n/200)/O(\sqrt{n}) = \Omega(\sqrt{n})$ with the same probability.

Now, we show that during this number of generations, the algorithm gets trapped in the local optima with overwhelming probability with the optimistic assumption that only the best b-cell in the population can be mutated into a locally optimal solution. Considering the best current individual, in each generation a 1-bit is flipped with probability at least $99/100$ in the first mutation step. Thus, the probability of not touching a 1-bit in $\Omega(\sqrt{n})$ generations is less than $(1/100)^{\Omega(\sqrt{n})}$. Considering that μ and dup are both $poly(n)$, the probability of flipping a 1-bit from the prefix in this number of generations and of not locating the optimum in the same amount of time is

$$P \geq \left(1 - \left(\frac{1}{100}\right)^{\Omega(\sqrt{n})}\right) \cdot \left(1 - \frac{\sqrt{n}}{200} \cdot dup \cdot \mu \cdot 2^{-\Omega(\sqrt{n})}\right) \geq 1 - 2^{-\Omega(\sqrt{n})}.$$

After flipping a 1-bit, the mutation operator mutates at most cn bits until it finds an improvement. All sampled solutions in the first $n/4 - n/100$ mutation steps will have more than $3n/4$ 1-bits and thus satisfy the first condition of the local optima. If the Hamming

distance between one of the first $n/4 - n/100$ sampled solutions and all SP points is at least γn , then the algorithm is in the trap. We only consider the SP points with a prefix containing more than $3n/4 - \gamma n$ 1-bits since SP points with less than $3n/4 - \gamma n$ 1-bits have already Hamming distance more than γn to the local optima.

Similarly to the analysis done in [22], we consider the k th mutation step for $k := 4\gamma n(1 + 1/5)/(3 - 4\gamma) = \frac{6cn}{30-5c} \leq cn$ where the last expression is due to $\gamma = c/8$. After k steps, the expected number of bits flipped in the prefix of length $n(3/4 - \gamma)$ is at least $k(3/4 - \gamma) = (1 + 1/5)\gamma n$. For any mutation potential $0 < c \leq 1$, $\frac{6cn}{30-5c} \leq n/4 - n/100$. Using Chernoff bounds, we can show that the probability of having less than γn 0-bits in the prefix is $P(X \leq (1 - 0.2)E(X)) \leq e^{-\frac{\gamma n \cdot 6}{1000}} = e^{-\Omega(n)}$.

Altogether, with probability $(1 - e^{-\Omega(\gamma n)}) \cdot (1 - 2^{-\Omega(\sqrt{n})}) = 1 - e^{-\Omega(\sqrt{n})}$ a point in the local optima is sampled. In $O(\log \mu) = O(\log n)$ generations this individual takes over the population. Once in the local optima, the algorithm needs at least $\text{dup} \cdot \mu \cdot 1/\binom{n}{\gamma n}$ time to find the global optimum according to Lemma 1. \square

In the following theorem we see how the (1+1) EA using SBM optimises HYPERTRAP_γ in polynomial time.

Theorem 15. *The (1+1) EA optimises HYPERTRAP_γ in $O(n^{3+\epsilon})$ steps w.o.p. $1 - e^{-\Omega(n^\epsilon)}$ for $\epsilon > 0$.*

Proof. With overwhelming probability, the algorithm is initialised within $n/2 \pm n/10$ 1-bits by Chernoff bounds. Since the distance to any trap point is linear and the probability that SBM flips at least \sqrt{n} bits in a single iteration ($\binom{n}{\sqrt{n}} n^{-\sqrt{n}} = 2^{-\Omega(\sqrt{n})}$) is exponentially small, so is the probability of mutating into a trap point. As the algorithm approaches the bit string with $n/2$ 1-bits, the distance to the trap remains linear. Conditional on not entering the trap, by standard AFL arguments it takes the (1+1) EA at most $O(\sqrt{n})$ steps in expectation to find a point with $n/2$ 1-bits, i.e., on the gradient. After finding such a point, the (1+1) EA improves on the gradient with probability at least $1/n^2 \cdot (1 - 1/n)^{n-2} \geq 1/(en^2)$ which is the probability of flipping the rightmost 1-bit and the leftmost 0-bit while leaving the rest of the bits untouched. To reach the first point of SP, there are a linear number of 1-bits that need to be shifted to the beginning of the bit string. It therefore takes $O(n^3)$ steps in expectation to find SP. The (1+1) EA improves on SP with probability at least $1/n \cdot (1 - 1/n)^{n-1} \geq 1/en$, which is the probability of flipping the leftmost 0-bit and not touching the other bits. Hence, the global optimum is found in $O(n^2)$ steps in expectation giving a total expected runtime of $O(n^3)$ conditional on not falling into the trap. By applying Markov's inequality iteratively over $\Omega(n^\epsilon)$ consecutive phases of length $\Theta(n^3)$ with an appropriate constant, the probability that the optimum is not found within $O(n^{3+\epsilon})$ steps is less than e^{-n^ϵ} with ϵ being an arbitrarily small constant. Since the probability of finding a local optimum from the gradient points or SP points is $n^{-\Omega(n)}$ in each step, the probability of not falling into the trap in $n^{3+\epsilon}$ steps is less than $(n^{3+\epsilon} \cdot n^{-c'n})$ by the union bound. Overall, the total probability of finding the optimum within $O(n^{3+\epsilon})$ steps is bigger than $(1 - e^{-n^\epsilon}) \cdot (1 - n^{3+\epsilon} \cdot n^{-c'n}) = 1 - e^{-\Omega(n^\epsilon)}$. \square

5.3. On Trap Functions

In [17], where Opt-IA was originally introduced, the effectiveness of the algorithm was tested for optimising the following simple trap function:

$$\text{SIMPLE TRAP}(x) := \begin{cases} \frac{a}{z}(z - |x|_1) & \text{if } |x|_1 \leq z, \\ \frac{b}{n-z}(|x|_1 - z) & \text{otherwise,} \end{cases}$$

where $z \approx n/4$, $b = n - z - 1$, $3b/2 \leq a \leq 2b$ and the optimal solution is the 0^n bit string.

The reported experimental results were averaged over 100 independent runs each with a termination criterion of reaching 5×10^5 fitness function evaluations. For all of the results, the population size (i.e., μ) was 10 and $\text{dup} = 1$. In these experiments, Opt-IA using either hypermutations or hypermacromutation never find the optimum of SIMPLE TRAP already for problem sizes $n > 50$. However, the following theorem shows that Opt-IA* indeed optimises SIMPLE TRAP efficiently.

Theorem 16. *Opt-IA needs $O(\mu n^2 \log n)$ expected fitness function evaluations to optimise SIMPLE TRAP with $\tau = \Omega(n^2)$ for $c = 1$ and $\text{dup} = 1$.*

Proof. Given that the number of 1-bits in the best solution is i , the probability of improving is at least $(n - i)/n$ if the best solution has more than z 1-bits and i/n otherwise. By following the proof of Theorem 3, at least one individual will reach 1^n or 0^n in $O(\mu n^2 \log n)$ fitness function evaluations in expectation as long as no individual dies due to ageing.

The age of an individual reaches n^2 only if the improvement fails to happen in n^2 generations which happens with probability at most $(1 - 1/n)^{n^2} = 2^{-\Omega(n)}$ since the improvement probability is at least $1/n$. This implies that the expected number of restarts by ageing is exponentially small. \square

Given that Opt-IA was tested in [17] also with the parameters suggested by Theorem 16 (i.e., $c = 1$, $\text{dup} = 1$, $\tau = \infty$), we speculate that either FCM was mistakenly not used or the stopping criterion (i.e., the total number of allowed fitness evaluations, i.e., 5×10^5) was too small. We point out that, for large enough τ also using hypermacromutation as mutation operator would lead to an expected runtime of $O(\mu n^2 \log n)$ for TRAP functions [8], with or without FCM. In any case, it is not necessary to apply both hypermutations and hypermacromutation together to efficiently optimise TRAP functions as reported in [17]. On the other hand, the inversely proportional hypermutation operator considered in [22] would fail to optimise this function efficiently because it cannot flip n bits when on the local optimum.

6. Not Allowing Genotype Duplicates

None of the algorithms considered in the previous sections use the genotype diversity mechanism. In this section, we do not allow genotype redundancies in the population as proposed in the original algorithm [17, 6]. This change potentially affects the behaviour of the algorithm. In the following, we will first consider the ageing operator in isolation with

Algorithm 9 $(\mu+1)$ RLS^{ageing} with genotype diversity (i.e., $div = 1$).

- 1: $t = 0$,
 - 2: initialise $P^{(t)} = \{x_1, \dots, x_\mu\}$, a population of μ individuals uniformly at random and set $x_i^{age} = 0$ for $i = \{1, \dots, \mu\}$.
 - 3: **while** the optimum is not found **do**
 - 4: select $x \in P^{(t)}$ uniformly at random,
 - 5: create y by flipping one bit of x ,
 - 6: Hybrid ageing $(P^{(t)}, \{y\}, \tau, \mu)$,
 - 7: Selection $(P^{(t)}, \{y\}, \mu, 1)$,
 - 8: $t = t + 1$.
 - 9: **end while**
-

genotype diversity (i.e., no genotypic duplicates are allowed in the population). Afterwards we will analyse the complete Opt-IA algorithm with the same diversity mechanism (as originally introduced in the literature).

6.1. $(\mu+1)$ RLS^{ageing} with genotype diversity

In this subsection, we analyse $(\mu+1)$ RLS^{ageing} (Algorithm 9) with genotype diversity for optimising CLIFF _{d} for which $(\mu+1)$ RLS _{p} ^{ageing} without genotype diversity was previously analysed in Section 4. The main difference compared to the analysis there is that taking over the population on the local optima is now harder since identical individuals are not allowed. The proof of the following theorem shows that the algorithm can still take over and use ageing to escape from the local optima as long as the population size is not too large.

Theorem 17. *For constant μ and $\tau = \Theta(n)$, the $(\mu+1)$ RLS^{ageing} with genotype diversity optimises CLIFF _{d} in expected $O(n \log n)$ fitness function evaluations for any linear $d \leq n/4 - \epsilon$.*

Proof. By Chernoff bounds, with overwhelming probability the initial individuals are sampled with $n(1/2 \pm \epsilon)$ 1-bits for any arbitrarily small $0 < \epsilon = \Theta(1)$. Since the population size is constant and there is a constant probability of improving in the first mutation step, a local optimum is found in at most $O(n)$ steps by picking the best individual and improving it $O(n)$ times.

If there is a single locally optimal solution in the population, then with probability $1/\mu$ this individual is selected as parent and produces an offspring with fitness $n - d - 1$ (i.e., one bit away) with probability $(n - d - i)/n \leq (n - d - \mu)/n = \Theta(1)$ where i is the number of individuals already with fitness $n - d - 1$. In the next generation, with probability $(i + 1)/\mu$ one of the individuals on the local optimum or one step away from it is selected as parent and produces an offspring on either the local optimum or one bit away with probability at least $(n - d - \mu)/n = \Theta(1)$. Hence, in expected time $\mu \cdot \Theta(1) = \Theta(\mu)$ all μ individuals are on the local optimum.

When the last inferior solution is replaced by an individual on the local optimum, the other individuals have ages in the order of $O(\mu)$. Thus, the probability that the rest

of the population does not die until the youngest individual reaches age τ , is at least $(1/\mu)^{(\mu-1)\cdot O(\mu)} = \Theta(1)$, the probability that $\mu - 1$ individuals above age τ survive $O(\mu)$ consecutive generations.

In the first generation when the last individual reaches age τ , with probability d/n an offspring is created at the bottom of the cliff (i.e., with a fitness value of $n - d + 1$) and with probability $1/\mu \cdot (1 - 1/\mu)^\mu = \Theta(1)$ all the parents die together at that step and the offspring survives. The rest of the proof follows the same arguments as the proof of Theorem 6.

Overall, the total expected time to optimise CLIFF_d is dominated by the time to climb the second ONEMAX slope which takes $O(n \log n)$ steps in expectation. \square

6.2. Opt-IA with genotype diversity

In this subsection, we analyse Opt-IA (Algorithm 6) with genotype diversity to optimise all the functions for which Opt-IA without genotype diversity was analysed in Section 5. The following are straightforward corollaries of Theorem 8 and Theorem 16. Since the ageing mechanism never triggers here and the proofs of those theorems do not depend on creating genotype copies, the arguments are still valid for Opt-IA with genotype diversity.

Corollary 2. *The upper bounds on the expected runtime of Opt-IA with genotype diversity and ageing parameter τ large enough for ONEMAX , LEADINGONES , JUMP_k and CLIFF_d are as follows:*

$$\begin{aligned} E(T_{\text{ONEMAX}}) &= O(\mu \cdot \text{dup} \cdot n^2 \log n), \\ E(T_{\text{LEADINGONES}}) &= O(\mu \cdot \text{dup} \cdot n^3), \\ E(T_{\text{JUMP}_k}) &= O\left(\mu \cdot \text{dup} \cdot \frac{n^{k+1} \cdot e^k}{k^k}\right), \\ \text{and } E(T_{\text{CLIFF}_d}) &= O\left(\mu \cdot \text{dup} \cdot \frac{n^{d+1} \cdot e^d}{d^d}\right). \end{aligned}$$

Corollary 3. *Opt-IA with genotype diversity needs $O(\mu n^2 \log n)$ expected fitness function evaluations to optimise SIMPLE TRAP with $\tau = \Omega(\mu n^{1+\epsilon})$, $c = 1$ and $\text{dup} = 1$.*

The following theorem shows the same expected runtime for Opt-IA with genotype diversity for HIDDENPATH as that of the Opt-IA without genotype diversity proven in Theorem 11. However, we reduce the population size to be constant. The proof follows the main arguments of the proof of Theorem 11. Here we only discuss the probability of events which should be handled differently considering that genotype duplicates are not allowed.

Theorem 18. *For $c = 1$, $\text{dup} = 1$, $\mu = \Theta(1)$ and $\tau = \Omega(n^2 \log n)$, Opt-IA with genotype diversity needs $O(\tau \mu n + \mu n^{7/2})$ expected fitness function evaluations to optimise HIDDENPATH .*

Proof. We follow the analysis of the proof of Theorem 11 for Opt-IA without genotype diversity. Although the analysis did not benefit from genotype duplicates, not allowing them potentially affects the runtime of the events where the population takes over. The potentially affected events are:

- For S_{n-1} solutions to take over the population, the probabilities are different here since a new S_{n-1} solution will not be accepted if it is identical to any current S_{n-1} solutions. Here, after finding the first S_{n-1} solution, the rest are created and accepted by Opt-IA with probability at least $\Omega(1/n \cdot (n - \mu)/(n - 1)) = \Omega(1/n)$. Therefore, the argument made in the proof of Theorem 11 does not change the runtime asymptotically.
- The arguments about the expected time needed for S_{n-1} solutions to reach the same age after the takeover are the same as in the proof of Theorem 11; without genotype duplicates, the probability of creating another S_{n-1} is still $\Omega(1/n)$. Hence, the probability of creating two copies in the same generation is still unlikely and we can adapt the arguments made in the proof of Theorem 11. Therefore, in expected $O(\mu^3 n)$ generations after the takeover of S_{n-1} , the population reaches the same age.
- In the proof of Theorem 11, the expected time for S_5 solutions to take over the population of recently initialised individuals is bounded relying on having multiple copies of one S_5 and one S_{n-5} solution. This proof strategy cannot be applied with the genotypic diversity mechanism which only allows unique solutions in the population.

In order to create a unique S_5 solution from another S_5 solution, it is sufficient that the first bit position to be flipped has value 1 in the parent bit string and the second position to be flipped has value 0 in all S_5 solutions currently in the population, including the parent. Such a mutation occurs with probability at least $(\frac{5}{n} \cdot \frac{n-5-\mu}{n-1}) \geq 4/n = \Omega(1/n)$. This lower bound in the probability implies that the number of S_5 individuals in the population is increased by a factor of $1 + \Omega(1/n)$ at every generation in expectation. For any constant c , the exponent t that satisfies $1 \cdot (1 + c/n)^t = \mu$ is in the order of $O(n \log \mu)$. Thus, in at most $O(n \log \mu)$ generations in expectation the S_5 solutions take over the population if no solutions with higher fitness have been added to the population before the takeover. Due to the assumption that $\mu = \Theta(1)$, the expected number of generations reduces to $O(n)$. By Markov's inequality, the probability that the expected time is in the order of $O(n)$ is $\Omega(1)$.

Only the solutions on SP and S_{n-1} solutions have better fitness than S_5 solutions. The rest of the proof of Theorem 11 can still be applied if the only remaining non S_5 solutions in the population are SP solutions. So, we will only show that it is unlikely that an S_{n-1} solution will be sampled before the population consists only of S_5 and SP solutions. The number of 0-bits in randomly initialised solutions is in the interval of $[n/2 - n^{2/3}, n/2 + n^{2/3}]$ with probability $1 - 2^{-\Omega(n^{1/3})}$ due to Chernoff bounds. We can then follow the strategy from Theorem 10 to show that the expectation of T^* , the time until an S_{n-1} solution descending from a randomly created solution is sampled, is in the order of $\Omega(n \log n)$. In order to bound the probability that this event will not happen in $O(n)$, we will bound the variance of this runtime. Since FCM stops after a solution with more 0-bits is sampled, we can pessimistically divide the runtime into phases of length T_i , $i \in \{2, \dots, n/2 + n^{2/3}\}$, where the best among the newly created solutions has i 1-bits. The length of phase T_i is distributed according to a geometric distribution with success probability at most $\mu \cdot 2i/n$ due to the Ballot theorem (i.e.,

Theorem 9) and the union bound summed up over μ candidate solutions which can be improved at every generation. Being geometrically distributed, the variance of T_i is at most $(1 - (d \cdot i/n))/(d \cdot i/n)^2$ for some constant $d > 2\mu$. Summing up over all phases of independently distributed lengths, we obtain the variance of T^* as:

$$\text{Var}(T^*) \leq \sum_{i=2}^{n/2+n^{2/3}} \frac{1 - (d \cdot i/n)}{(d \cdot i/n)^2} \leq \frac{n^2}{d^2} \sum_{i=2}^{n/2+n^{2/3}} \frac{1}{i^2} = \frac{n^2 \pi^2}{d^2 6} = \Theta(n^2).$$

Due to Chebyshev's Inequality, the probability that such an event happens in $O(n)$ generations instead of its expectation, which is in the order of $\Omega(n \log n)$, is at most $O(1/\log n)$. Conversely, the probability that such a failure does not occur is $1 - O(1/\log n)$.

The path solutions have between 5 to $\log n + 1$ 1-bits, thus the probability that the hypermutation operator yields an S_{n-1} solution as output given an SP or S_5 solution as input is at most $\binom{n}{n-4}^{-1} = O(n^{-4})$. The probability that such a mutation occurs in $O(n)$ generations is at most $O(n^{-3})$ by the union bound. Considering all the possible failure events, the takeover happens in $O(n)$ generations with $\Omega(1)$ probability before any S_{n-1} individual is added to the population.

The rest of the proof of Theorem 11 is not affected by genotype diversity. \square

7. Conclusion

We have presented an analysis of the standard Opt-IA artificial immune system. We first highlighted how both the ageing and hypermutation operators may allow to efficiently escape local optima that are particularly hard for standard evolutionary algorithms. Concerning hypermutations, we proved that FCM is essential to the operator and suggested considering a mutation *constructive* if the produced fitness is at least as good as the previous one. The reason is that far away points of equal fitness should be attractive for the sake of increasing exploration capabilities. Our analysis on the JUMP_k function suggests that hypermutation with FCM is generally preferable to the SBM operator when escaping the local optima requires a jump of size k such that $(k/e)^k \geq n$. This advantage is least pronounced when there is a single solution with better fitness than the local optima as in the case of JUMP_k and the performance difference with the SBM in terms of expected escape time scales multiplicatively with the number of acceptable solutions at distance k . Hence, the JUMP_k function may be considered as a worst-case scenario concerning the advantages of hypermutations over SBM for escaping local optima.

Concerning ageing, we showed for the first time that the operator can be very efficient when coupled with SBM and hypermutations. To the best of our knowledge, the operator allows the best known expected runtime (i.e., $O(n \log n)$) for hard CLIFF_d functions (this expected runtime has recently been matched by a simple hyperheuristic [36]). Afterwards, we presented a class of functions where both the characteristics of ageing and hypermutation are crucial, hence Opt-IA is efficient while standard evolutionary algorithms are inefficient even

if coupled with one extra AIS operator (either cloning, ageing, hypermutation or contiguous somatic mutation). Finally, we proved that all the positive results presented for the Opt-IA algorithm without genotype diversity also hold for Opt-IA with genotype diversity as used in the original algorithm. However, small population sizes may be required if ageing has to be triggered to escape from local optima with genotype diversity. To complete the picture we presented a class of problems where the use of hypermutations and ageing is detrimental while standard evolutionary algorithms are efficient.

Our analysis shows that for easy problems for which local search strategies are efficient, using hypermutations and ageing may be detrimental. We have shown this effect for the simple ONEMAX and LEADINGONES functions for which we have proven a linear slow-down in the expected runtime compared to local search strategies and simple evolutionary algorithms. While such a slow-down may be considered an acceptable expense in exchange for being efficient on more complicated optimisation problems, we have shown for HYPERTRAP_{c/8} that the consequences may be more drastic, by making the difference between polynomial and exponential runtimes. Indeed, the function is easy for local search algorithms with neighbourhoods of size greater than 1 and for simple EAs. However, hypermutations make Opt-IA fail with overwhelming probability and ageing does not help the algorithm to escape from the trap permanently. We point out that recently *Fast Hypermutation* operators have been presented that allow to hillclimb efficiently while still keeping the efficiency at escaping local optima [39]. However, such Fast AIS still suffer on HYPERTRAP_{c/8}.

On the other hand, for more complicated multimodal functions, with closer characteristics to the optimisation problems that occur in practice, we have shown several advantages of hypermutations and ageing to escape from local optima (i.e., JUMP_k and CLIFF_d). Furthermore, the combination of hypermutation and ageing may be advantageous to locate new promising basins of attraction that are hard to find via more traditional optimisation techniques such as local search or EAs using SBM. We have illustrated such effect in the analysis of HIDDENPATH, where the combination of hypermutations and ageing allow Opt-IA to locate a new basin of attraction that initially has lower fitness than the easy-to-find local optima. However, in the long run, having identified this new basin of attraction allows the algorithm to find the global optimum.

Overall, we believe this work is a significant contribution towards the understanding of which kind of problems it is advantageous to use artificial immune systems on rather than evolutionary algorithms and for which it is detrimental. Future work should focus on providing such advantages and disadvantages for classical combinatorial optimisation problems, in similar fashion to the recently presented results for the NP-hard PARTITION problem that have been achieved by building upon the present work [40].

Acknowledgement: The research leading to these results has received funding from the EPSRC under grant agreement no EP/M004252/1.

References

References

- [1] D. Corus, P. S. Oliveto, D. Yazdani, On the runtime analysis of the Opt-IA artificial immune system, in: Proc. of GECCO 2017, 2017, pp. 83–90.
- [2] L. N. de Castro, J. Timmis, Artificial Immune Systems: A New Computational Intelligence Paradigm, Springer-Verlag, Secaucus, NJ, USA, 2002.
- [3] F. M. Burnet, The Clonal Selection Theory of Acquired Immunity, Cambridge University Press, 1959.
- [4] L. N. de Castro, F. J. V. Zuben, Learning and optimization using the clonal selection principle, IEEE Transaction on Evolutionary Computation 6 (2002) 239–251.
- [5] J. Kelsey, J. Timmis, Immune inspired somatic contiguous hypermutation for function optimisation, in: Proc. of GECCO 2003, 2003, pp. 207–218.
- [6] V. Cutello, M. Pavone, J. Timmis, An immune algorithm for protein structure prediction on lattice models, IEEE Transactions on Evolutionary Computation 10 (2006) 844–861.
- [7] V. Cutello, G. Nicosia, M. Romeo, P. S. Oliveto, On the convergence of immune algorithms, in: Proc. of FOCI 2007, 2007, pp. 409–415.
- [8] T. Jansen, C. Zarges, Analyzing different variants of immune inspired somatic contiguous hypermutations, Theoretical Computer Science 412 (2011) 517–533.
- [9] D. Corus, J. He, T. Jansen, P. S. Oliveto, D. Sudholt, C. Zarges, On easiest functions for mutation operators in bio-inspired optimisation, Algorithmica (2016).
- [10] C. Zarges, Rigorous runtime analysis of inversely fitness proportional mutation rates, in: Proc. of PPSN X, 2008, pp. 112–122.
- [11] C. Zarges, On the utility of the population size for inversely fitness proportional mutation rates, in: Proc. of FOGA 2009, 2009, pp. 39–46.
- [12] T. Jansen, C. Zarges, On the role of age diversity for effective aging operators, Evolutionary Intelligence 4 (2011) 99–125.
- [13] C. Horoba, T. Jansen, C. Zarges, Maximal age in randomized search heuristics with aging, in: Proc. of GECCO 2009, 2009, pp. 803–810.
- [14] P. S. Oliveto, D. Sudholt, On the runtime analysis of stochastic ageing mechanisms, in: Proc. of GECCO 2014, 2014, pp. 113–120.
- [15] T. Jansen, P. S. Oliveto, C. Zarges, On the analysis of the immune-inspired B-Cell algorithm for the vertex cover problem, in: Proc. of ICARIS 2011, 2011, pp. 117–131.
- [16] T. Jansen, C. Zarges, Computing longest common subsequences with the B-Cell Algorithm, in: Proc. of ICARIS 2012, 2012, pp. 111–124.
- [17] V. Cutello, G. Nicosia, M. Pavone, Exploring the capability of immune algorithms: A characterization of hypermutation operators, in: Proc. of ICARIS 2004, 2004, pp. 263–276.
- [18] V. Cutello, G. Nicosia, M. Pavone, A hybrid immune algorithm with information gain for the graph coloring problem, in: Proc. of GECCO 2003, 2003, pp. 171–182.
- [19] V. Cutello, G. Nicosia, A clonal selection algorithm for coloring, hitting set and satisfiability problems, Neural Nets 3931 (2006) 324–337.
- [20] P. S. Oliveto, X. Yao, Runtime analysis of evolutionary algorithms for discrete optimisation, in: Theory of Randomized Search Heuristics: Foundations and Recent Developments, World Scientific, 2011, pp. 21–52.
- [21] T. Jansen, Analyzing Evolutionary Algorithms: The Computer Science Perspective, Springer, 2013.
- [22] T. Jansen, C. Zarges, Variation in artificial immune systems: Hypermutations with mutation potential, in: Proc. of ICARIS 2011, 2011, pp. 132–145.
- [23] P. S. Oliveto, P. K. Lehre, F. Neumann, Theoretical analysis of rank-based mutation-combining exploration and exploitation, in: Proc. of CEC 2009, 2009, pp. 1455–1462.
- [24] D. Corus, P. S. Oliveto, Standard steady state genetic algorithms can hillclimb faster than mutation-only evolutionary algorithms, IEEE Trans. Evol. Comput. 22 (2018) 720–732.

- [25] A. Lissovoi, P. S. Oliveto, J. A. Warwicker, On the runtime analysis of generalised selection hyper-heuristics for pseudo-boolean optimisation, in: Proc. of GECCO 2017, 2017, pp. 849–856.
- [26] B. Doerr, A. Lissovoi, P. S. Oliveto, J. A. Warwicker, On the runtime analysis of selection hyper-heuristics with adaptive learning periods, in: Proc. of GECCO 2018, 2018, pp. 1015–1022.
- [27] J. Lengler, A general dichotomy of evolutionary algorithms on monotone functions, in: Proc. of PPSN XV, 2018, pp. 3–15.
- [28] T. Friedrich, A. Göbel, F. Quinzan, M. Wagner, Heavy-tailed mutation operators in single-objective combinatorial optimization, in: Proc. of PPSN XV, 2018, pp. 134–145.
- [29] S. Droste, T. Jansen, I. Wegener, On the analysis of the $(1+1)$ evolutionary algorithm, *Theoretical Computer Science* 276 (2002) 51–81.
- [30] D.-C. Dang, T. Friedrich, T. Kötzing, M. S. Krejca, P. K. Lehre, P. S. Oliveto, D. Sudholt, A. M. Sutton, Escaping local optima using crossover with emergent diversity, *IEEE Transactions on Evolutionary Computation* 22 (2018) 484–497.
- [31] B. Doerr, H. P. Le, R. Makhmara, T. D. Nguyen, Fast genetic algorithms, in: Proc. of GECCO 2017, 2017, pp. 777–784.
- [32] J. Jägersküpfer, T. Storch, When the plus strategy outperforms the comma strategy and when not, in: Proc. of FOCI 2007, 2007, pp. 25–32.
- [33] P. S. Oliveto, T. Paixão, J. Pérez Heredia, D. Sudholt, B. Trubenová, How to escape local optima in black box optimisation: When non-elitism outperforms elitism, *Algorithmica* (2017). doi: 10.1007/s00453-017-0369-2.
- [34] T. Paixão, J. Pérez Heredia, D. Sudholt, B. Trubenová, Towards a runtime comparison of natural and artificial evolution, *Algorithmica* 78 (2017) 681–713.
- [35] D. Sudholt, A new method for lower bounds on the running time of evolutionary algorithms, *IEEE Transactions on Evolutionary Computation* 17 (2012) 418–435.
- [36] A. Lissovoi, P. S. Oliveto, J. A. Warwicker, On the time complexity of algorithm selection hyper-heuristics for multimodal optimisation, in: Proc. of AAAI 2019, 2019. To appear.
- [37] W. Feller, *An Introduction to Probability Theory and Its Applications*, John Wiley & Sons, 1968.
- [38] F. Neumann, D. Sudholt, C. Witt, Rigorous analyses for the combination of ant colony optimization and local search, in: Proc. of ANTS 2008, volume 5217, 2008, pp. 132–143.
- [39] D. Corus, P. S. Oliveto, D. Yazdani, Fast artificial immune systems, in: Proc. of PPSN XV, 2018, pp. 67–78.
- [40] D. Corus, P. S. Oliveto, D. Yazdani, Artificial immune systems can find arbitrarily good approximations for the NP-hard partition problem, in: Proc. of PPSN XV, 2018, pp. 16–28.