



HAL
open science

Already Moderate Population Sizes Provably Yield Strong Robustness to Noise

Denis Antipov, Benjamin Doerr, Alexandra Ivanova

► **To cite this version:**

Denis Antipov, Benjamin Doerr, Alexandra Ivanova. Already Moderate Population Sizes Provably Yield Strong Robustness to Noise. GECCO '24: Genetic and Evolutionary Computation Conference, 2024, Melbourne VIC, Australia. pp.1524-1532, 10.1145/3638529.3654196 . hal-04649104

HAL Id: hal-04649104

<https://hal.science/hal-04649104>

Submitted on 16 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Already Moderate Population Sizes Provably Yield Strong Robustness to Noise

Denis Antipov
Optimisation and Logistics,
School of Computer and
Mathematical Sciences,
The University of Adelaide
Adelaide, Australia

Benjamin Doerr
Laboratoire d'Informatique (LIX)
CNRS, École Polytechnique
Institut Polytechnique de Paris
Palaiseau, France

Alexandra Ivanova
HSE University, Skoltech
Moscow, Russia

ABSTRACT

Experience shows that typical evolutionary algorithms can cope well with stochastic disturbances such as noisy function evaluations. In this first mathematical runtime analysis of the $(1 + \lambda)$ and $(1, \lambda)$ evolutionary algorithms in the presence of prior bit-wise noise, we show that both algorithms can tolerate constant noise probabilities without increasing the asymptotic runtime on the OneMax benchmark. For this, a population size λ suffices that is at least logarithmic in the problem size n . The only previous result in this direction regarded the less realistic one-bit noise model, required a population size super-linear in the problem size, and proved a runtime guarantee roughly cubic in the noiseless runtime for the OneMax benchmark. Our significantly stronger results are based on the novel proof argument that the noiseless offspring can be seen as a biased uniform crossover between the parent and the noisy offspring. We are optimistic that the technical lemmas resulting from this insight will find applications also in future mathematical runtime analyses of evolutionary algorithms.

CCS CONCEPTS

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

KEYWORDS

Noisy optimization, Runtime analysis, Theory, Population-based algorithms

ACM Reference Format:

Denis Antipov, Benjamin Doerr, and Alexandra Ivanova. 2024. Already Moderate Population Sizes Provably Yield Strong Robustness to Noise. In *Genetic and Evolutionary Computation Conference (GECCO '24)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3638529.3654196>

1 INTRODUCTION

The mathematical runtime analysis has accompanied and supported the design and analysis of evolutionary algorithms (EAs) for more than 30 years. It has led to a deeper understanding of many important aspects of evolutionary computation [3, 13, 25, 37, 49].



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

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

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

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

<https://doi.org/10.1145/3638529.3654196>

While this area has also studied how EAs cope with noise, that is, a stochastically disturbed access to the true problem instance, the rigorous understanding of this aspect of evolutionary computation is rather limited. Most previous works regard simple algorithms like the $(1 + 1)$ EA and show that these, without particular adjustments, can stand a moderate level of noise, but usually not more than one noisy fitness evaluation every roughly $n/\log(n)$ iterations for problems with bit-string representation of length n .

Some works have shown that larger population sizes improve the robustness to noise, but these still do not show a very satisfying picture. For example, in the work closest to ours, Gießen and Kötzing [23] showed that the $(1 + \lambda)$ EA can stand random one-bit prior noise with arbitrary rate $q \in [0, 1]$ when optimizing the ONEMAX benchmark. However, their result requires a relatively large population size of at least $\lambda = \Omega(\frac{1}{q}n \log n)$ and then gives a runtime guarantee of $O(\frac{1}{q}n^2\lambda)$ fitness evaluations, significantly above the guarantee $O(\max\{n\lambda \frac{\log \log \lambda}{\log \lambda}, n \log n\})$ for the noiseless setting [12]. We also note that both the required population size and the runtime guarantee contain a factor of $\frac{1}{q}$, that is, they become worse for lower noise rates. This counter-intuitive dependence on the noise intensity together with the weak runtime guarantee of order at least $n^3 \log n$ suggest that this problem is not yet fully understood, and this is why we undertake a new attempt to analyze how the $(1 + \lambda)$ EA solves the ONEMAX problem in the presence of noise.

As a main result, we prove that the $(1 + \lambda)$ EA with any population size $\lambda \geq C \ln(n)$, C a suitable constant, can optimize the ONEMAX problem in the presence of bit-wise prior noise with up to constant noise probability per iteration, in asymptotically the same time as when no noise is present. We note that we regard a different noise model than in [23], namely independent bit-wise prior noise. This model is generally regarded as more realistic because the true and the noisy fitness can deviate by arbitrary amounts. There is no reason to believe that the $(1 + \lambda)$ EA should suffer less from noise in this noise model. In fact, we are convinced that results analogous to ours hold in the one-bit noise model regarded in [23], and that such results can be proven with a variant of our general proof method. We do not conduct these proofs since, as often in runtime analyses, the precise proofs rely on the details of the particular algorithm, objective function, and noise model. For that reason, even though our general approach seems to apply also to one-bit noise, the proofs would differ in many small details. We therefore leave this task for future work.

This result is the first tight analysis of a standard population-based EA in the presence of noise in a standard model (in fact, it is the first analysis of the $(1 + \lambda)$ EA in the presence of bit-wise

noise). Our result shows that already moderate population sizes can yield an enormous robustness. We note that for $\lambda = \Theta(\log n)$, our runtime guarantee is $O(n \log n)$, that is, the same as for the simple $(1 + 1)$ EA in the noiseless setting. Hence the larger population size used here to obtain robustness does not lead to an increase in the runtime. We also recall the general lower bound of $\Omega(n \log n)$ valid for all unary unbiased black-box algorithm [33], which shows that our $O(n \log n)$ bound is asymptotically tight and that a better performance is not possible in the realm of mutation-based unbiased evolutionary algorithms.

We extend our analysis to the non-elitist $(1, \lambda)$ EA and show the same results for this algorithm. In the noiseless setting, the $(1 + \lambda)$ EA and the $(1, \lambda)$ EA are known to have a similar performance when the population size is at least logarithmic, basically because the $(1, \lambda)$ EA becomes near-elitist since with high probability at least one offspring is equal to the parent. That this similarity of the two algorithms extends to noisy settings was, a priori, not obvious. We note that this is the first runtime analysis of the $(1, \lambda)$ EA in the presence of prior noise.

Our results are based on a novel proof argument, namely that the noiseless offspring can be seen as a biased crossover between the parent and the noisy offspring. This allows to obtain probabilistic information on the offspring given the parent and the noisy offspring. Note that the noisy offspring is visible to the algorithm (and thus in a sense also to our proofs) via its fitness, whereas the non-noisy offspring is not visible, but is of course what really counts for the further run of the algorithm. We refer to Sections 4 and 5 for more details.

Overall, our work gives two main insights, namely (1) that using EAs with at least a moderate population size can lead to a strong robustness against noise, and this without performance losses even when compared with using the optimal population size in the noiseless setting, and (2) that such processes can be analyzed with mathematical means, in particular, with the tools we developed to gain from the noisy offspring probabilistic information on the noiseless offspring.

We complement our theoretical analysis with a small experimental study, aiming at answering two questions which our asymptotic runtime analysis naturally could not answer. We observe that, as predicted by theoretical considerations in the noiseless case, also in the noisy setting there is no significant performance difference between the $(1 + \lambda)$ EA and the $(1, \lambda)$ EA. Also, we show that the asymptotic runtime advantage of the $(1 + \lambda)$ EA over the $(1 + 1)$ EA at constant noise rates is clearly visible already for moderate population sizes.

This work is organized as follows. In Section 2, we review the most relevant previous works. We describe the benchmark, noise model, and algorithms in Section 3. Our main technical tool, the analysis of the relation between parent, true offspring, and noisy offspring, is developed in Section 4. We use this tool in Section 5 to conduct the runtime analyses leading to the main results of this work. Our experimental study can be found in Section 6. The paper ends with a short conclusion and outlook.

2 RELATED WORKS

One of the most recent and detailed overviews of mathematical analyses of evolutionary algorithms in noisy environments can be found in the recent paper [32]. We refer the reader to this work and discuss now only the results most relevant to ours.

For the $(1 + 1)$ EA, several results show that it can tolerate a small amount of noise, but becomes highly inefficient for larger noise rates. In [23], it was shown that on ONEMAX with either prior bit-wise noise flipping each bit with probability $\frac{q}{n}$ or one-bit noise flipping exactly one random bit with probability q , the $(1 + 1)$ EA keeps its $O(n \log(n))$ runtime from the noiseless setting when $q = O(\frac{1}{n})$. For $q = O(\frac{\log(n)}{n})$, the runtime is still polynomial, but for $q = \omega(\frac{\log(n)}{n})$ the runtime is super-polynomial. That is, the $(1 + 1)$ EA can cope with noise only if it happens at most, on average, every $\Theta(\frac{n}{\log(n)})$ iterations. A similar effect was observed on LEADINGONES and made very precise in [44], where for the same noise models for all $q \leq \frac{1}{2}$ a runtime of order $n^2 \exp(\Theta(\min\{qn^2, n\}))$ was shown. Hence here already from $q = 1/n^2$ on the performance drops rapidly with increasing noise rate.

Posterior additive Gaussian noise was studied in [23, 40]. In [23] the authors showed that the $(1 + 1)$ EA can solve ONEMAX and LEADINGONES in its noiseless time when the variance σ^2 of the Gaussian distribution is at most $\frac{1}{4 \log(n)}$ and $\frac{1}{12en^2}$, respectively. In [40] it was shown that for large variance, namely 1 and n^2 , respectively, the runtime on these problems is exponential.

A more general approach based on estimating the probability that the noise inverts the comparison of two individuals was used in [6]. This led to a polynomial runtime guarantee on ONEMAX for every noise model for which the probability to invert a comparison via noise is at most $\frac{c \ln(n)}{n}$ for some constant c . Also, this approach allowed to make the results of [23] more precise, among others, giving that the $O(n \log(n))$ runtime bound remains valid up to $q = c' \log(\log(n))/n$ with c' a constant specified in [6].

It is also worth mentioning the papers [38–40], which showed that resampling an individual sufficiently often can help to determine its true fitness. This can make the $(1 + 1)$ EA robust to noise, but often requires large (more than n^3) numbers of samples. This strategy, however, does not work on LEADINGONES with high-rate prior noise. Additionally, in [14] it was shown that using a median instead of the mean of the samples can significantly reduce the number of resamplings necessary to yield a reasonable performance, e.g., to $\Theta(\log(n))$ if the noise is not too strong.

For population-based EAs there are considerably fewer results. In [23] it was shown that for the $(1 + \lambda)$ EA and the $(\mu + 1)$ EA, the population sizes $\mu = \Omega(\frac{\log(n)}{q})$ and $\lambda = \Omega(\frac{1}{q} n \log(n))$ can lead to polynomial runtimes on ONEMAX with one-bit noise occurring with probability q . The runtimes in this case are $O(\mu n \log(n))$ for the $(\mu + 1)$ EA, mildly above its $O(\mu n)$ noiseless runtime guarantee [47], and $O(\frac{1}{q} n^2 \lambda)$ for the $(1 + \lambda)$ EA, significantly above its noiseless runtime guarantee of $O(n \lambda \frac{\log \log \log \lambda}{\log \log \lambda})$ [12]. This result definitely suggests that the larger populations, in particular, larger parent populations, can be beneficial to cope with noise. However, these results counter-intuitively require larger population sizes to cope with smaller noise rates.

Another result for population-based EAs was obtained in [39], where the authors considered a symmetric noise model. The authors showed that with logarithmic population sizes both the $(1 + \lambda)$ EA and the $(\mu + 1)$ EA have an $\tilde{O}(n)$ runtime on ONEMAX, and also that smaller population sizes for the $(1 + \lambda)$ EA lead to exponential runtimes. The downside of this result is that the symmetric noise model is quite artificial. In particular, for the $(\mu + 1)$ EA it implies that if we have a single individual with best fitness, it is removed from the population only if the noise affects all individuals in the population (which is very unlikely when μ is at least logarithmic). For the $(1 + \lambda)$ EA it implies that we lose fitness only when all $\Theta(\lambda)$ copies of the current individual are affected by the noise, which is also unlikely for logarithmic or larger values of λ . With these particular properties, these results appear hard to extend to the more common noise models.

In [6] the authors write that the $(1 + \lambda)$ EA can optimize ONEMAX under a one-bit noise in $O(n \log(n))$ runtime, if $\lambda = \Theta(\log(np))$, where p is the probability that the noise flips a bit. However, this result is only described informally, there is no theorem stating it, nor a proof or a proof sketch.

The study of non-elitist algorithms in the presence of noise is confined to the three works [4, 31, 32]. In [4], it is shown that a non-elitist EA constructing the next population by λ times independently selecting two parents and taking the mutant of the better one into the next population, is very robust to prior noise with up to constant noise rates, keeping its noise-free runtime apart from a $\log \log n$ factor. Besides many other results, this result was sharpened and extended to more general noise models in [32], however, to the best of our understanding only for mutation rate below $1/n$. That paper, as well as the subsequent work [31], also give results for the symmetric noise model, but as discussed above, we are not optimistic that such results can be extended to more realistic noise models.

Other slightly less relevant results include studies of ant colony optimizers (ACO) [20], estimation of distribution algorithms such as the cGA [21] and the UMDA [30], and voting algorithms [41] in the presence of noise. For an ACO with a fitness-proportional update it was shown in [20] that with a small evaporation factor $\rho = O(\frac{1}{n^3 \log(n)})$ this algorithm can solve ONEMAX in $O(n^2 \log(n)/\rho)$ time for both Gaussian or 1-bit noise with any noise rate (note that this runtime is of order larger than n^5). For the cGA on ONEMAX with Gaussian noise with variance σ^2 it was shown that with population size $K = \omega(\sigma \sqrt{n} \log(n))$ the optimum is found in $O(K \sigma^2 \sqrt{n} \log(Kn))$ time [21]. With an automated choice of the algorithm parameter K , the runtime stemming from the optimal choice of K can be obtained [48], and this without knowing in advance the noise variance σ^2 . In [30], the UMDA was studied on LEADINGONES with a one-bit noise with rate less than one. It was shown that with parameters μ and λ which are at least logarithmic in n , the runtime is $O(n \lambda \log(\lambda) + n^2)$. The voting algorithm was studied on ONEMAX with Gaussian noise, and it was shown that if the noise rate σ^2 is at least $\frac{3n}{8}$, then the optimum is correctly found after $O(\sigma^2 \log(n))$ samples.

Frozen noise models, where noise is applied to the objective function once before the algorithm is run, were studied in [22, 28]. The main result is that RLS, the $(1 + 1)$ EA, and the $(1 + \lambda)$ EA are

not efficient on those rugged landscapes, but the compact genetic algorithm and the $(1, \lambda)$ EA are.

Noisy optimization has also been studied for multi-objective EAs [5, 7], but the very different population dynamics render it difficult to compare these results with the single-objective setting.

3 PRELIMINARIES

3.1 OneMax with Bit-wise Noise

ONEMAX is one of the most common benchmark functions used in theoretical studies of EAs. It returns the number of one-bits in its argument, which is formally defined by

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

for all $x \in \{0, 1\}^n$. Despite its simplicity, this benchmark has had an enormous impact on the field. It was the basis of the first mathematical runtime analyses at all [19, 35, 43], the first analyses of EAs in the presence of noise and dynamic changes [16, 17], the first analyses of population-based algorithms [1, 26, 47], or the first analyses of estimation-of-distribution algorithms and ant-colony optimizers [18, 24, 36]. These works triggered the development of many analysis methods that are used regularly since then. The study of how EAs optimize ONEMAX also led to the discovery of many new algorithmic ideas, e.g., various ways to dynamically adjust the parameters of an EA [8, 15, 29].

In this paper we study the optimization of ONEMAX under *bit-wise prior noise*. In this noise model, we do not learn the correct fitness of x , but the fitness of some bit string obtained from x by flipping each bit independently with probability $\frac{q}{n}$. We call $\frac{q}{n}$ the *noise rate*. In this paper we consider $q = O(1)$, so that the noise rate is at most $O(\frac{1}{n})$, which means that we might have a constant probability that the noise occurs and we learn a possibly wrong fitness value.

3.2 The $(1 + \lambda)$ EA and the $(1, \lambda)$ EA

The two most simple EAs with a non-trivial offspring population are the $(1 + \lambda)$ EA and $(1, \lambda)$ EA. Both algorithms work according to the following scheme. They keep a bit-string x , which is called the *current individual* and which is initialized with a random bit string. Then in each iteration they create λ individuals, independently, by applying standard bit mutation with rate $\frac{\chi}{n}$ to the current individual x (that is, each bit in a copy of x is flipped with probability $\frac{\chi}{n}$ independently from other bits). A best of these offspring is selected as the mutation winner y . Then the only difference in two algorithm occurs. The $(1 + \lambda)$ EA compares y with x and if y is not worse, then it replaces x in the next iteration. Otherwise x stays the same. In the $(1, \lambda)$ EA we use a non-elitist selection, and y always replaces x , even if it is worse. In this paper we use the standard assumption that $\chi = \Theta(1)$, since smaller mutation rates usually reduce the optimization speed and larger rates can lead to an exponential runtime even on monotone functions [11]. The pseudocodes of the $(1, \lambda)$ EA and the $(1 + \lambda)$ EA are shown in Algorithm 1, where the only difference between them is in lines 9 and 10.

In the context of noisy optimization, it is important to note that for the $(1 + \lambda)$ EA we re-compute the fitness of the current individual in each iteration and do not reuse the value computed

Algorithm 1: The $(1, \lambda)$ EA and the $(1 + \lambda)$ EA maximizing a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$.

```

1 Initialization: Sample  $x \in \{0, 1\}^n$  uniformly at random and
  evaluate  $f(x)$ ;
2 Optimization: for  $t = 1, 2, 3, \dots$  do
3   for  $i = 1, \dots, \lambda$  do
4      $x^{(i)} \leftarrow$  copy of  $x$ ;
5     Flip each bit in  $x^{(i)}$  with probability  $\frac{\chi}{n}$ ;
6     evaluate  $f(x^{(i)})$ ;
7   end
8    $y \leftarrow \arg \max\{f(x^{(i)}) \mid i \in [\lambda]\}$ ;
9   if running  $(1, \lambda)$  EA then  $x \leftarrow y$ ;
10  if running  $(1 + \lambda)$  EA and  $f(y) \geq f(x)$  then  $x \leftarrow y$ ;
11 end

```

previously. This allows to correct the situation that we have a current individual for which we believe its fitness to be significantly higher (due to a noisy fitness evaluation in the past). In such a situation the EA may get stuck for a long time, simply because all offspring appear to be worse than this parent. That this problem is real and can be seen, e.g., by comparing the results of [45] (without reevaluations) and [10] (with reevaluations).

To take into account the specifics of noisy optimization, we use the following point of view and notation when we consider an iteration of the $(1 + \lambda)$ EA or the $(1, \lambda)$ EA. We denote the λ offspring of x created via standard bit mutation by $x^{(i)}$ (for $i \in [1.. \lambda]$). Then when we apply noise to each of λ individuals independently, we obtain λ noisy individuals $\tilde{x}^{(i)}$, where for each $i \in [1.. \lambda]$ individual $\tilde{x}^{(i)}$ is obtained from $x^{(i)}$. When we talk about an arbitrary offspring and its noisy version, we denote them by x' and \tilde{x}' correspondingly. After creating noisy offspring we evaluate their ONEMAX values and we choose an individual \tilde{y} from the λ noisy ones with the best fitness value (the ties are broken uniformly at random) as the mutation winner and we denote its non-noisy parent by y . Then for the $(1 + \lambda)$ EA we also apply noise to x before we compare it with \tilde{y} , and thus we get a noisy individual \tilde{x} which competes with \tilde{y} .

3.3 Drift Analysis

Drift analysis is a rapidly developing set of tools which are widely used in the analysis of random search heuristics. They help to transform easy-to-obtain information about the expected progress into bounds on the expected runtime. In this paper we use the following variable drift theorem, which was first introduced in [27, 34]. We use its simplified version for processes over integer values from [9].

THEOREM 1 (THEOREM 6 IN [9]). *Let $(X_t)_{t \in \mathbb{N}}$ be a sequence of random variables in $[0..n]$ and let T be the random variable that denotes the earliest point in time $t \geq 0$ such that $X_t = 0$. Suppose that there exists a monotonically increasing function $h : [1..n] \mapsto \mathbb{R}_0^+$ such that*

$$E[X_t - X_{t+1} \mid X_t] \geq h(X_t)$$

Table 1: Distribution of a bit value in the offspring x' given the observed value in the noisy offspring \tilde{x}' .

Event A	Event B	$\Pr[A \mid B]$	Approximate Value
$x'_i = x_i$	$\tilde{x}'_i = x_i$	$\frac{(1-\frac{\chi}{n})(1-\frac{q}{n})}{(1-\frac{\chi}{n})(1-\frac{q}{n})+\frac{q\chi}{n^2}}$	$1 - \frac{q\chi}{n^2} - O\left(\frac{1}{n^3}\right)$
$x'_i = x_i$	$\tilde{x}'_i \neq x_i$	$\frac{(1-\frac{\chi}{n})\frac{q}{n}}{(1-\frac{\chi}{n})\frac{q}{n}+\frac{\chi}{n}(1-\frac{q}{n})}$	$\frac{q}{q+\chi} (1 \pm o(1))$
$x'_i \neq x_i$	$\tilde{x}'_i = x_i$	$\frac{\frac{\chi q}{n^2}}{(1-\frac{\chi}{n})(1-\frac{q}{n})+\frac{q\chi}{n^2}}$	$\frac{q\chi}{n^2} + O\left(\frac{1}{n^3}\right)$
$x'_i \neq x_i$	$\tilde{x}'_i \neq x_i$	$\frac{(1-\frac{q}{n})\frac{\chi}{n}}{(1-\frac{\chi}{n})\frac{q}{n}+\frac{\chi}{n}(1-\frac{q}{n})}$	$\frac{\chi}{q+\chi} (1 \pm o(1))$

holds for all $t < T$ (as an inequality of random variables). Then

$$E[T \mid X_0] \leq \sum_{i=1}^{X_0} \frac{1}{h(i)}.$$

4 OFFSPRING DISTRIBUTION

In this section we discuss the relationship between the parent individual x , its offspring x' obtained via standard bit mutation, and the individual \tilde{x}' obtained from the offspring via bit-wise noise.

We start with the distribution of \tilde{x}' , when we have no information about the true offspring x' . Each bit of the noisy offspring \tilde{x}' is different from the bit in the same position in x , if and only if either it was flipped by the mutation, but not by the noise, or it was flipped by the noise, but not by the mutation. The probability of this event is

$$\frac{\chi}{n} \left(1 - \frac{q}{n}\right) + \frac{q}{n} \left(1 - \frac{\chi}{n}\right) = \frac{\chi}{n} + \frac{q}{n} - 2 \cdot \frac{\chi q}{n^2} = \frac{1}{n} \left(\chi + q - \frac{2q\chi}{n}\right).$$

Since this event is independent for all bits, \tilde{x}' is distributed in the same way as if it was created from x via standard bit mutation with rate $\frac{r}{n}$, where $r = \chi + q - \frac{2q\chi}{n}$. Since in this paper we assume $\chi = \Theta(1)$ and $q = O(1)$, we have $r = \Theta(1)$ as well.

When we know the parent x and also the noisy offspring \tilde{x}' , we can estimate the distribution of the noiseless offspring x' , which is an intermediate step when we get \tilde{x}' from x . The following lemma shows how the bit values in x' are distributed when we know those bit values in x and \tilde{x}' . Since, naturally, these estimates depend not on the particular bit value in x , but only on whether the corresponding bits agree or differ with this value, we also formulate our result in this symmetric fashion.

LEMMA 2. *Consider an arbitrary bit position i . The distribution of this bit value x'_i in x' , conditional on the values of this bit in x and \tilde{x}' , are as shown in Table 1.*

PROOF. We start with the first row of Table 1, that is, we compute $\Pr[x'_i = x_i \mid \tilde{x}'_i = x_i]$. By the definition of conditional probabilities we have

$$\Pr[x'_i = x_i \mid \tilde{x}'_i = x_i] = \frac{\Pr[x'_i = x_i = \tilde{x}'_i]}{\Pr[x_i = \tilde{x}'_i]}.$$

The bit in position i has the same values in all three individuals only if it has not been flipped, neither by mutation, nor by noise. The probability of this event is $(1 - \frac{\chi}{n})(1 - \frac{q}{n})$. This bit is equal in

the parent x and in the noisy offspring \tilde{x}' when we either have not flipped it at all, or we flipped it twice. The probability of this event is $(1 - \frac{\chi}{n})(1 - \frac{q}{n}) + \frac{q\chi}{n^2}$. Thus, we have

$$\begin{aligned} \Pr[x'_i = x_i \mid \tilde{x}'_i = x_i] &= \frac{\left(1 - \frac{\chi}{n}\right)\left(1 - \frac{q}{n}\right)}{\left(1 - \frac{\chi}{n}\right)\left(1 - \frac{q}{n}\right) + \frac{q\chi}{n^2}} \\ &= 1 - \frac{\frac{q\chi}{n^2}}{\left(1 - \frac{\chi}{n}\right)\left(1 - \frac{q}{n}\right) + \frac{q\chi}{n^2}} = 1 - \frac{q\chi}{n^2} \cdot \frac{1}{\left(1 - O\left(\frac{1}{n}\right)\right)} \\ &= 1 - \frac{q\chi}{n^2} - O\left(\frac{1}{n^3}\right). \end{aligned}$$

To prove the second row of Table 1, we use similar arguments to show that the event $x_i = x'_i \neq \tilde{x}'_i$ happens only if we flip this bit by noise, but do not flip it by mutation. The probability of this is $(1 - \frac{\chi}{n})\frac{q}{n}$. We have $x_i \neq \tilde{x}'_i$, when we flip the i -th bit only once, the probability of which is $(1 - \frac{\chi}{n})\frac{q}{n} + (1 - \frac{q}{n})\frac{\chi}{n}$. Hence, we have

$$\begin{aligned} \Pr[x'_i = x_i \mid \tilde{x}'_i \neq x_i] &= \frac{\Pr[x'_i = x_i \neq \tilde{x}'_i]}{\Pr[x_i \neq \tilde{x}'_i]} = \frac{\left(1 - \frac{\chi}{n}\right)\frac{q}{n}}{\left(1 - \frac{\chi}{n}\right)\frac{q}{n} + \frac{\chi}{n}\left(1 - \frac{q}{n}\right)} \\ &= \frac{q}{q + \chi\left(\frac{1 - \frac{q}{n}}{1 - \frac{\chi}{n}}\right)} = \frac{q}{q + \chi} \cdot \frac{q + \chi}{q + \chi \pm o(1)} = \frac{q}{q + \chi} (1 \pm o(1)), \end{aligned}$$

since $\chi = \Theta(1)$ and $q = O(1)$.

By regarding complementary events, we obtain the remaining two rows of Table 1:

$$\begin{aligned} \Pr[x'_i \neq x_i \mid \tilde{x}'_i = x_i] &= 1 - \Pr[x'_i = x_i \mid \tilde{x}'_i = x_i] \\ &= \frac{\frac{\chi q}{n^2}}{\left(1 - \frac{\chi}{n}\right)\left(1 - \frac{q}{n}\right) + \frac{q\chi}{n^2}} = \frac{q\chi}{n^2} \left(1 + O\left(\frac{1}{n}\right)\right) = \frac{q\chi}{n^2} + O\left(\frac{1}{n^3}\right), \end{aligned}$$

$$\begin{aligned} \Pr[x'_i \neq x_i \mid \tilde{x}'_i \neq x_i] &= 1 - \Pr[x'_i = x_i \mid \tilde{x}'_i \neq x_i] \\ &= \frac{\chi}{q + \chi} (1 \pm o(1)). \quad \square \end{aligned}$$

Lemma 2 can be interpreted in the way that given x and \tilde{x}' , the true offspring x' is, asymptotically, a biased crossover between x and \tilde{x}' , taking bit values from x with probability $\frac{q}{q + \chi}$ and from \tilde{x}' otherwise. Indeed, if x and \tilde{x}' agree in a bit value, then with high probability x' also has this value in this bit (the first and the third lines of Table 1). Where x and \tilde{x}' differ, the offspring x' takes the value from x with probability $\frac{q}{q + \chi} \pm o(1)$, see the second line of Table 1, and otherwise from \tilde{x}' .

By linearity of expectation, this observation can be lifted to distances, which is what we do in the following lemma. It implies, in particular, that if we observe some progress towards some target solution with respect to the noisy offspring, then a constant fraction of this progress is real, that is, witnessed by the true offspring. This insight will be the central tool in our later analyses. We are optimistic that it will be useful in other runtime analyses of EAs in the presence of noise as well.

LEMMA 3. *Consider some arbitrary, but fixed parent x . Let $d(\cdot)$ denote the Hamming distance to some arbitrary point x^* in the search*




Figure 1: Illustration of the four groups of bits in the proof of Lemma 3.

space. Then for any bit string z we have

$$\begin{aligned} E[d(x) - d(x') \mid \tilde{x}' = z] &\geq (d(x) - d(z)) \cdot \frac{\chi}{q + \chi} \cdot (1 \pm o(1)) - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right). \end{aligned}$$

PROOF. W.l.o.g. we assume that x^* is the all-ones bit string¹, thus $d(\cdot)$ stands for the number of zero-bits in its argument. Consider some arbitrary z . We divide the bits into four groups depending on their values in x and z as illustrated in Figure 1. Let A be the number of bits which are both ones in x and z , let B be the number of bits which are ones in x and zeros in z , let C be the number of bits which are zeros in x and ones in z , and let D be the number of bits with are both zeros in x and z . Note that in each of these four groups the number of zero-bits in x' follows a binomial distribution, with success probability as given in Table 1. We denote the precise probabilities from rows 3 and 4 of Table 1 by p_3 and p_4 respectively (the ones in column $\Pr[A \mid B]$). Then the probabilities from rows 1 and 2 are $(1 - p_3)$ and $(1 - p_4)$ respectively. Hence, for any arbitrary z we have

$$\begin{aligned} E[d(x') \mid \tilde{x}' = z] &= p_3 A + p_4 B + (1 - p_4) C + (1 - p_3) D \\ &= D + C + p_3(A - D) + p_4(B - C) \\ &= (D + C) + (A - D) \left(\frac{q\chi}{n^2} + O\left(\frac{1}{n^3}\right) \right) + (B - C) \frac{\chi}{q + \chi} (1 \pm o(1)) \\ &\leq d(x) + n \left(\frac{q\chi}{n^2} + O\left(\frac{1}{n^3}\right) \right) + (d(z) - d(x)) \frac{\chi}{q + \chi} (1 \pm o(1)) \\ &= d(x) + \frac{q\chi}{n} + O\left(\frac{1}{n^2}\right) + (d(z) - d(x)) \frac{\chi}{q + \chi} (1 \pm o(1)), \end{aligned}$$

since $C + D = d(x)$, $A - D \leq n$, and $B - C = d(z) - d(x)$. By moving $d(x)$ to the left hand side and multiplying both sides by -1 , we finish the proof. \square

5 RUNTIME ANALYSIS

We use the results from Section 4 to estimate upper bounds on the runtime of the $(1 + \lambda)$ EA and the $(1, \lambda)$ EA. The main result of this section is the following theorem.

THEOREM 4. *Consider a run of the $(1, \lambda)$ EA or the $(1 + \lambda)$ EA with mutation rate $\frac{\chi}{n}$ where $\chi = \Theta(1)$ on the ONEMAX problem with*

¹To generalize this proof for an arbitrary bit string x^* , it is sufficient to replace “one-bits” with “bits which have the same value in x^* ” and “zero-bits” with “bits which have a different value in x^* ”. We avoid doing so to improve the readability of the proof.

bit-wise noise with rate $\frac{q}{n}$ where $q = O(1)$. If the population size λ is at least $C \ln(n)$ for some constant C depending on χ and q , then the expected number $E[T_F]$ of fitness evaluations until the optimum is sampled is

$$O\left(n \log(n) + n\lambda \frac{\log \log(\lambda)}{\log(\lambda)}\right).$$

To prove Theorem 4, we use the variable drift theorem (Theorem 1). The process we apply it to is $\{d_t\}_{t \in \mathbb{N}_0}$, which is the distance of the current individual (of either the $(1, \lambda)$ EA or the $(1 + \lambda)$ EA) to the optimum after iteration t . Since we consider the ONEMAX function, d_t is equal to the number of zero-bits in that individual.

To bound the drift of d_t , we study what happens in one iteration of the two considered algorithms. We use the notation defined in Section 3.2, and additionally we denote the individuals accepted as the parent for the next iteration by x_{com} for the $(1, \lambda)$ EA and by x_{plus} for the $(1 + \lambda)$ EA. Note that x_{com} is always equal to y (the winning offspring), while x_{plus} can be either y or x , depending on their noisy comparison. We also denote the distances to the optimum from different individuals as follows. By d and \tilde{d} we denote the distance to the optimum from the parent x and the noisy parent \tilde{x} , respectively. By d_y and \tilde{d}_y we denote the distance to the optimum from the mutation winner y and its noisy version \tilde{y} , respectively. By d_{com} and d_{plus} we denote distances from next generation parents x_{com} and x_{plus} to the optimum, respectively.

With this notation the drift in one iteration is defined as

$$\Delta^{\text{com}}(d) := E[d - d_{\text{com}}] \text{ for the } (1, \lambda) \text{ EA, and}$$

$$\Delta^{\text{plus}}(d) := E[d - d_{\text{plus}}] \text{ for the } (1 + \lambda) \text{ EA.}$$

We estimate these drifts in the following lemma.

LEMMA 5. *Let*

$$\begin{aligned} \Delta^+(d) &:= \sum_{i=1}^d \Pr[\tilde{d}_y \leq d - i], \text{ and} \\ \Delta^-(d) &:= \sum_{j=1}^{n-d} \Pr[\tilde{d}_y \geq d + j]. \end{aligned}$$

Then the drift of the current individual x towards the optimum in one iteration is at least

$$\Delta^{\text{com}}(d) \geq (1 - o(1)) \frac{\chi}{q + \chi} (\Delta^+(d) - \Delta^-(d)) - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right) \quad (1)$$

for the $(1, \lambda)$ EA, and it is at least

$$\begin{aligned} \Delta^{\text{plus}}(d) &\geq (1 - o(1)) \frac{\chi}{q + \chi} ((1 - o(1))e^{-q}\Delta^+(d) - \Delta^-(d)) \\ &\quad - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right) \end{aligned} \quad (2)$$

for the $(1 + \lambda)$ EA.

Before we prove this lemma we note that $\Delta^+(d)$ and $\Delta^-(d)$ are the positive and the negative components of the drift of \tilde{y} from x respectively, since we have $E[d - \tilde{d}_y] = \Delta^+(d) - \Delta^-(d)$.

PROOF. In this proof we consider a fixed parent individual x , and therefore d is also fixed (that is, it is not a random variable). For the $(1, \lambda)$ EA we have $x_{\text{com}} = y$, hence we have

$$\begin{aligned} \Delta^{\text{com}}(d) &= E[d - d_{\text{com}}] = E[d - d_y] \\ &= \sum_{Y \in \{0,1\}^n} \Pr[\tilde{y} = Y] E[d - d_y \mid \tilde{y} = Y]. \end{aligned}$$

For any bit string $z \in \{0, 1\}^n$ we denote the distance to the optimum by d_z . Then by Lemma 3 we have

$$\begin{aligned} \Delta^{\text{com}}(d) &\geq \sum_{z \in \{0,1\}^n} \Pr[\tilde{y} = z] \left((d - d_z) \frac{\chi}{q + \chi} (1 \pm o(1)) - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right) \right) \\ &= \left((1 - o(1)) \frac{\chi}{q + \chi} \sum_{z \in \{0,1\}^n} \Pr[\tilde{y} = z] (d - d_z) \right) - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right) \\ &= (1 - o(1)) \frac{\chi}{q + \chi} E[d - \tilde{d}_y] - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right). \end{aligned}$$

Noting that $E[d - \tilde{d}_y] = \Delta^+(d) - \Delta^-(d)$ by the definition of $\Delta^+(d)$ and $\Delta^-(d)$ completes the proof of eq. (1).

Estimating drift for the $(1 + \lambda)$ EA requires more effort, since we are not guaranteed that $x_{\text{plus}} = y$. This happens if and only if \tilde{y} appears no worse than \tilde{x} (the individual we obtain from x via noise). Consequently, we have

$$d - d_{\text{plus}} = (d - d_y) \mathbb{I}[\tilde{d} \geq \tilde{d}_y], \quad (3)$$

where all random variables are over the joint probability space of noise on x , noise on offspring and the mutation, $\mathbb{I}[\cdot]$ is an indicator random variable, and eq. (3) is an identity of random variables functions. Therefore, by the law of total probability, the drift for the $(1 + \lambda)$ EA is

$$\begin{aligned} \Delta^{\text{plus}}(d) &= E[d - d_{\text{plus}}] = E\left[(d - d_y) \mathbb{I}[\tilde{d} \geq \tilde{d}_y]\right] \\ &= \sum_{z \in \{0,1\}^n} \Pr[\tilde{y} = z] \cdot \left(\Pr[\tilde{d} \geq \tilde{d}_y \mid \tilde{y} = z] \right. \\ &\quad \cdot E\left[d - d_y \mid \tilde{y} = z, \tilde{d} \geq \tilde{d}_y\right] + \Pr[\tilde{d} < \tilde{d}_y \mid \tilde{y} = z] \cdot 0 \left. \right). \end{aligned}$$

Note that when we fix x and \tilde{y} , then event $\tilde{d} \geq \tilde{d}_y$ depends only on the noise which affects \tilde{x} . This noise does not affect any offspring of x , hence the random variable $d - d_y$ conditioned on $\tilde{y} = z$ is independent of $\tilde{d} \geq \tilde{d}_y$, which together with Lemma 3 implies that

$$\begin{aligned} E\left[d - d_y \mid \tilde{y} = z, \tilde{d} \geq \tilde{d}_y\right] &= E\left[d - d_y \mid \tilde{y} = z\right] \\ &\geq (1 \pm o(1)) \frac{\chi}{q + \chi} (d - d_z) - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right). \end{aligned}$$

Hence, we have

$$\begin{aligned} \Delta^{\text{plus}}(d) &\geq (1 - o(1)) \frac{\chi}{q + \chi} \\ &\quad \cdot \sum_{z \in \{0,1\}^n} \Pr[\tilde{y} = z] \Pr[\tilde{d} \geq \tilde{d}_y \mid \tilde{y} = z] (d - d_z) \\ &\quad - \frac{q\chi}{n} - O\left(\frac{1}{n^2}\right). \end{aligned} \quad (4)$$

When $d_z > d$, we estimate the conditional probability $\Pr[\tilde{d} \geq \tilde{d}_y \mid \tilde{y} = z]$ by one. Otherwise, if $d_z < d$, this probability is at least the probability that $\tilde{x} = x$, that is, that the noise has not flipped any bit in x when we compared it with \tilde{y} . This probability is $(1 - \frac{q}{n})^n = (1 - o(1))e^{-q}$. With this observation, similar to the $(1, \lambda)$ EA, we can rewrite the sum above as follows

$$\begin{aligned} & \sum_{z \in \{0,1\}^n} \Pr[\tilde{y} = z] \Pr[\tilde{d} \geq \tilde{d}_y \mid \tilde{y} = z] (d - d_z) \\ & \geq \sum_{i=1}^d (1 - o(1))e^{-qi} \Pr[\tilde{d}_y = d - i] - \sum_{j=1}^{n-d} j \Pr[\tilde{d}_y = d + j] \\ & = (1 - o(1))e^{-q} \Delta^+(d) - \Delta^-(d). \end{aligned}$$

Putting this into eq. (4) completes the proof of eq. (2). \square

Above we obtained a weaker drift estimate for the $(1 + \lambda)$ EA than the $(1, \lambda)$ EA. This is counter-intuitive – one would feel that the $(1 + \lambda)$ EA should profit to some extent from the property that the current-best solution is participating in the selection of the next parent. The reason for our weaker bound in the estimate following eq. 4, where we pessimistically estimated that inferior mutation winners y are always accepted, but superior ones only when the parent is not subject to noise. We feel that this cannot be avoided in the general case (note that Lemma 5 is valid in general and is not specific for the optimization of ONEMAX). We are sure that when exploiting properties of ONEMAX, stronger bounds could be shown. We refrain from this both because we like our general, problem-independent approach and because all we can gain are constant factors, which are generally ignored in an asymptotic analysis as ours (note that we lose constant factors and lower order terms in the later part of the analysis anyway).

In the following lemmas we show estimates for $\Delta^+(d)$ and $\Delta^-(d)$, which naturally are specific to the ONEMAX problem. Lemmas 6-8 estimate the positive component of the drift for different distances. We omit their proofs for reasons of space². We also note that implicitly these results exist in the proofs in [12], but distilling them from there is rather complicated.

LEMMA 6 (POSITIVE DRIFT, LARGE DISTANCE). *If $\lambda = \omega(1)$ then for any distance $d \geq \frac{n}{\ln \lambda}$ we have $\Pr[d - \tilde{d}_y \geq \lfloor \frac{\ln \lambda}{2 \ln \ln \lambda} \rfloor] \geq \frac{1}{4}$ and the positive component of the drift is at least*

$$\Delta^+(d) \geq \frac{1}{4} \left\lfloor \frac{\ln \lambda}{2 \ln \ln \lambda} \right\rfloor.$$

LEMMA 7 (POSITIVE DRIFT, MEDIUM DISTANCE). *For $d \in [\frac{n}{\lambda}, \frac{n}{\ln \lambda}]$ we have $\Pr[d - \tilde{d}_y \geq 1] \geq \frac{r}{\lambda e^r} (1 - o(1))$ and the positive drift is at least $\Delta^+(d) \geq \frac{r}{2e^r} (1 - o(1)) = \Theta(1)$, where $r = q + \chi - \frac{2q\chi}{n}$.*

LEMMA 8 (POSITIVE DRIFT, SMALL DISTANCE). *For $d \leq \frac{n}{\lambda}$ we have $\Pr[d - \tilde{d}_y \geq 1] \geq \frac{r}{\lambda e^r} (1 - o(1))$ and the positive drift is at least $\Delta^+(d) \geq \frac{\lambda dr}{2ne^r} (1 - o(1))$.*

We proceed with bounding the negative component of the drift. A similar estimate for this negative part of the drift of the $(1, \lambda)$ EA on the noiseless ONEMAX was shown in [42]. However, there this drift was bounded with a constant, which is not enough in our situation, so we give a stronger bound.

²All proofs can be found in the full version of this paper available at arXiv [2].

LEMMA 9 (NEGATIVE DRIFT). *If $\lambda > C \ln(n)$ for some sufficiently large constant C which depends on r , then for all $d \in [1..n]$ we have $\Delta^-(d) \leq \frac{1}{n}$.*

PROOF. If \tilde{d}_y is at least d , then for all λ noisy offspring \tilde{x}' their distances to the optimum \tilde{d}' are at least d . Then we have

$$\sum_{j=1}^{n-d} \Pr[d - \tilde{d}_y \geq j] = \sum_{j=1}^{n-d} (\Pr[d - \tilde{d}' \geq j])^\lambda.$$

Let $p_j := \Pr[d - \tilde{d}' \geq j]$. Since p_1 is at most the probability that at least one 1-bit in x is flipped, we have $p_1 \leq 1 - (1 - \frac{r}{n})^{n-d} \leq 1 - e^{-r} + o(1)$ (the last inequality is by eq. 3.6.2 in [46]). Hence,

$$p_1^\lambda \leq (1 - e^{-r} + o(1))^\lambda \leq \frac{1}{n^2}$$

for $\lambda \geq C \ln n$ with $C = -\frac{2}{\ln(1 - e^{-r} + o(1))}$. For all $j \in \mathbb{N}$ we have $p_j \leq p_1$, and therefore $p_j^\lambda \leq \frac{1}{n^2}$. Consequently,

$$\Delta^-(d) = \sum_{j=1}^{n-d} (\Pr[d - \tilde{d}' \geq j])^\lambda = \sum_{j=1}^{n-d} p_j^\lambda \leq \frac{n-d}{n^2} \leq \frac{1}{n}. \quad \square$$

With Lemma 5 and with estimates given in Lemmas 6-9 we are now in position to prove our main result, Theorem 4. We only sketch the proof for reasons of space³.

By Lemmas 6-9, for all $d \in [1..n]$ the positive component $\Delta^+(d)$ of the drift is asymptotically larger than the negative component $\Delta^-(d)$. Together with Lemma 5 this observation allows us to show that the drift is always bounded from below by some $h(d) = \Theta(\Delta^+(d))$, both for the $(1 + \lambda)$ EA and for the $(1, \lambda)$ EA.

We then apply Theorem 1 to this $h(d)$ and split the sum into three ranges of distance d which correspond to Lemmas 6-8. The sum of $h^{-1}(d)$ over large distances is $O(n \frac{\log \log \lambda}{\log \lambda})$, over medium distances it is $O(\frac{n}{\log \lambda})$ and over small distances it is $O(\frac{n}{\lambda} \log \frac{n}{\lambda})$. We sum these three values and multiply them by the number of fitness evaluations performed by the algorithms in one iteration, that is, λ for the $(1, \lambda)$ EA and $\lambda + 1$ for the $(1 + \lambda)$ EA, which yields the stated bound.

We also note that for all $\lambda = O(\frac{\log(n) \log \log(n)}{\log \log \log(n)})$ the first term of the upper bound in Theorem 4 is dominating, hence the expected runtime is $O(n \log(n))$, that is, the same as in the noiseless case.

6 EXPERIMENTS

In this section we describe the results of our empirical study, the goal of which is to give a better understanding of the noisy optimization process and to answer some questions for which our theoretical analysis was not detailed enough.

When we estimated the drift of the $(1 + \lambda)$ EA in Lemma 5, we pessimistically assumed that the elitism of this algorithm can be only harmful. Namely, in our proofs, the comparison with the parent does not save us from decreasing the fitness (since the parent's fitness might be decreased by the noise) and it might prevent us from increasing fitness (since the parent's fitness might be increased by the noise). This pessimistic view eased the proof without harming

³The full proof can be found in the arXiv version of this paper [2].




Figure 2: Mean runtimes (number of fitness evaluations, normalized by $n \ln n$) and their standard deviations over 100 runs of the $(1 + \lambda)$ EA and the $(1, \lambda)$ EA on ONEMAX with noise rates $q = 1$ with varying problem size n .

the asymptotic order of magnitude of the runtime, so it was fully appropriate for the theoretical analysis. Still, it raises the question which of the two algorithms is better when looking at the precise runtime rather than its asymptotics. To understand this question, we ran the $(1 + \lambda)$ EA and the $(1, \lambda)$ EA on ONEMAX with problem sizes $n = \{2^6, 2^7, \dots, 2^{14}\}$, with 100 repetitions for each value of n . We used a strong noise with $q = 1$ to maximize its effect, standard bit mutation with $\chi = 1$, and we chose $\lambda = \lceil C \ln(n) \rceil$ with $C = 14$, since for the chosen values of q and χ this C would satisfy the assumptions of Lemma 9. The mean runtimes of these runs and their standard deviations are shown in the plot in Figure 2. For better visual comparison, we normalize both runtimes by $n \ln(n)$, which is our upper bound on their asymptotical runtime and the well-known lower bound for all unary unbiased black-box algorithms [33]. In the plot we see that for all problem sizes the difference between the mean runtimes of two algorithms is very small compared to the standard deviations of their runtimes. Therefore, we conclude that in practice for this population size the elitism does not slow down the algorithm by more than some lower order terms (but neither can speed it up). We note that this is well-known for the noise-free setting, simply because with high probability some offspring equals the parent, so removing the parent in the $(1, \lambda)$ EA has no negative effect. For the noisy setting, as our proof shows, this is less obvious.

We also compare the $(1 + \lambda)$ EA with $\lambda = \lfloor \ln(n) \rfloor$ and the $(1 + 1)$ EA on ONEMAX in the presence of bit-wise noise with small but constant value of $q = 0.01$. Ours and the previous theoretical works showed that in this setting the asymptotic runtime of the $(1 + \lambda)$ EA is much better, but the question is if this translates to significant runtime differences already for common problem sizes. Our results, depicted in Figure 3, show that the difference between the algorithms is notable already on the smallest problem size $n = 2^6$. We also note that since we use a logarithmic scale for both axes, the convex plot of the $(1 + 1)$ EA indicates that its runtime is super-polynomial in problem size n , which lines up with the theoretical results in [23].

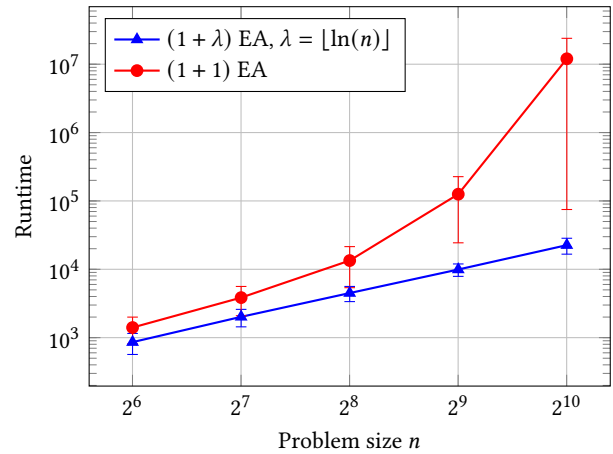


Figure 3: Mean runtimes (number of fitness evaluations) and their standard deviations over 100 runs of the $(1 + \lambda)$ EA and the $(1 + 1)$ EA on ONEMAX with noise rates $q = 0.01$ with varying problem size n .

7 CONCLUSION

In this work, we have proven that both the $(1 + \lambda)$ EA and $(1, \lambda)$ EA with at least logarithmic population sizes are very robust to noise, that is, with up to constant noise probabilities they optimize the ONEMAX benchmark in asymptotically the same time as if no noise was present. This significantly improves the state of the art for the $(1 + \lambda)$ EA, considerably reducing both the required population size and the runtime guarantee, and this is the first such result for the $(1, \lambda)$ EA.

The reason for this progress is the general observation that the noise-free offspring can be seen as a biased uniform crossover between the parent and the noisy offspring. From this we proved that the true progress is at least a constant fraction of the noisy progress. The latter can be analyzed with known methods because the noisy offspring, asymptotically, has the distribution of an offspring obtained from bit-wise mutation with a rate that is the sum of the mutation and the noise rate. We are optimistic that this or analogous arguments will find applications in future runtime analyses again.

Next steps to continue this line of research could include the first analysis of the (μ, λ) EA [1] in the presence of noise or an analysis of the algorithms studied in this work on the LEADING-ONES benchmark, which generally is more affected by noise.

ACKNOWLEDGMENTS

This research benefited from the support of the FMJH Program Gaspard Monge for optimization and operations research and their interactions with data science. It was also supported by Australian Research Council through grant DP190103894.

REFERENCES

- [1] Denis Antipov and Benjamin Doerr. 2021. A tight runtime analysis for the $(\mu + \lambda)$ EA. *Algorithmica* 83 (2021), 1054–1095.
- [2] Denis Antipov, Benjamin Doerr, and Alexandra Ivanova. 2024. Already Moderate Population Sizes Provably Yield Strong Robustness to Noise. (2024).

- arXiv:2404.02090
- [3] Anne Auger and Benjamin Doerr (Eds.). 2011. *Theory of Randomized Search Heuristics*. World Scientific Publishing.
 - [4] Duc-Cuong Dang and Per Kristian Lehre. 2015. Simplified runtime analysis of estimation of distribution algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2015*. ACM, 513–518.
 - [5] Duc-Cuong Dang, Andre Opris, Bahare Salehi, and Dirk Sudholt. 2023. Analysing the robustness of NSGA-II under noise. In *Genetic and Evolutionary Computation Conference, GECCO 2023*. ACM, 642–651.
 - [6] Raphaël Dang-Nhu, Thibault Dardinier, Benjamin Doerr, Gautier Izacard, and Dorian Nogneng. 2018. A new analysis method for evolutionary optimization of dynamic and noisy objective functions. In *Genetic and Evolutionary Computation Conference, GECCO 2018*. ACM, 1467–1474.
 - [7] Matthieu Dinot, Benjamin Doerr, Ulysse Hennebel, and Sebastian Will. 2023. Runtime analyses of multi-objective evolutionary algorithms in the presence of noise. In *International Joint Conference on Artificial Intelligence, IJCAI 2023*. ijcai.org, 5549–5557.
 - [8] Benjamin Doerr, Carola Doerr, and Franziska Ebel. 2015. From black-box complexity to designing new genetic algorithms. *Theoretical Computer Science* 567 (2015), 87–104.
 - [9] Benjamin Doerr, Carola Doerr, and Jing Yang. 2020. Optimal parameter choices via precise black-box analysis. *Theoretical Computer Science* 801 (2020), 1–34.
 - [10] Benjamin Doerr, Ashish Ranjan Hota, and Timo Kötzing. 2012. Ants easily solve stochastic shortest path problems. In *Genetic and Evolutionary Computation Conference, GECCO 2012*. ACM, 17–24.
 - [11] Benjamin Doerr, Thomas Jansen, Dirk Sudholt, Carola Winzen, and Christine Zarges. 2013. Mutation rate matters even when optimizing monotone functions. *Evolutionary Computation* 21 (2013), 1–21.
 - [12] Benjamin Doerr and Marvin Künnemann. 2015. Optimizing linear functions with the $(1 + \lambda)$ evolutionary algorithm—different asymptotic runtimes for different instances. *Theoretical Computer Science* 561 (2015), 3–23.
 - [13] Benjamin Doerr and Frank Neumann (Eds.). 2020. *Theory of Evolutionary Computation—Recent Developments in Discrete Optimization*. Springer. Also available at http://www.lix.polytechnique.fr/Labo/Benjamin.Doerr/doerr_neumann_book.html.
 - [14] Benjamin Doerr and Andrew M. Sutton. 2019. When resampling to cope with noise, use median, not mean. In *Genetic and Evolutionary Computation Conference, GECCO 2019*. ACM, 242–248.
 - [15] Benjamin Doerr, Carsten Witt, and Jing Yang. 2021. Runtime analysis for self-adaptive mutation rates. *Algorithmica* 83 (2021), 1012–1053.
 - [16] Stefan Droste. 2002. Analysis of the $(1+1)$ EA for a dynamically changing OneMax-variant. In *Congress on Evolutionary Computation, CEC 2002*. IEEE, 55–60.
 - [17] Stefan Droste. 2004. Analysis of the $(1+1)$ EA for a noisy OneMax. In *Genetic and Evolutionary Computation Conference, GECCO 2004*. Springer, 1088–1099.
 - [18] Stefan Droste. 2005. Not all linear functions are equally difficult for the compact genetic algorithm. In *Genetic and Evolutionary Computation Conference, GECCO 2005*. ACM, 679–686.
 - [19] Stefan Droste, Thomas Jansen, and Ingo Wegener. 2002. On the analysis of the $(1+1)$ evolutionary algorithm. *Theoretical Computer Science* 276 (2002), 51–81.
 - [20] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. 2016. Robustness of ant colony optimization to noise. *Evolutionary Computation* 24 (2016), 237–254.
 - [21] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. 2017. The compact genetic algorithm is efficient under extreme Gaussian noise. *IEEE Transactions on Evolutionary Computation* 21 (2017), 477–490.
 - [22] Tobias Friedrich, Timo Kötzing, Frank Neumann, and Aishwarya Radhakrishnan. 2022. Theoretical study of optimizing rugged landscapes with the cGA. In *Parallel Problem Solving from Nature, PPSN 2022, Part II*. Springer, 586–599.
 - [23] Christian Gießen and Timo Kötzing. 2016. Robustness of populations in stochastic environments. *Algorithmica* 75 (2016), 462–489.
 - [24] Walter J. Gutjahr. 2008. First steps to the runtime complexity analysis of ant colony optimization. *Computers & Operations Research* 35 (2008), 2711–2727.
 - [25] Thomas Jansen. 2013. *Analyzing Evolutionary Algorithms – The Computer Science Perspective*. Springer.
 - [26] Thomas Jansen, Kenneth A. De Jong, and Ingo Wegener. 2005. On the choice of the offspring population size in evolutionary algorithms. *Evolutionary Computation* 13 (2005), 413–440.
 - [27] Daniel Johannsen. 2010. *Random Combinatorial Structures and Randomized Search Heuristics*. Ph. D. Dissertation. Universität des Saarlandes.
 - [28] Joost Jorritsma, Johannes Lengler, and Dirk Sudholt. 2023. Comma selection outperforms plus selection on OneMax with randomly planted optima. In *Genetic and Evolutionary Computation Conference, GECCO 2023*. ACM, 1602–1610.
 - [29] Jörg Lässig and Dirk Sudholt. 2011. Adaptive population models for offspring populations and parallel evolutionary algorithms. In *Foundations of Genetic Algorithms, FOGA 2011*. ACM, 181–192.
 - [30] Per Kristian Lehre and Phan Trung Hai Nguyen. 2021. Runtime analyses of the population-based univariate estimation of distribution algorithms on LeadingOnes. *Algorithmica* 83 (2021), 3238–3280.
 - [31] Per Kristian Lehre and Xiaoyu Qin. 2023. Self-adaptation can improve the noise-tolerance of evolutionary algorithms. In *Foundations of Genetic Algorithms, FOGA 2023*. ACM, 105–116.
 - [32] Per Kristian Lehre and Xiaoyu Qin. 2024. More precise runtime analyses of non-elitist evolutionary algorithms in uncertain environments. *Algorithmica* 86 (2024), 396–441.
 - [33] Per Kristian Lehre and Carsten Witt. 2012. Black-box search by unbiased variation. *Algorithmica* 64 (2012), 623–642.
 - [34] Boris Mitavskiy, Jonathan E. Rowe, and Chris Cannings. 2009. Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *International Journal on Intelligent Computing and Cybernetics* 2 (2009), 243–284.
 - [35] Heinz Mühlenbein. 1992. How genetic algorithms really work: mutation and hillclimbing. In *Parallel Problem Solving from Nature, PPSN 1992*. Elsevier, 15–26.
 - [36] Frank Neumann and Carsten Witt. 2009. Runtime analysis of a simple ant colony optimization algorithm. *Algorithmica* 54 (2009), 243–255.
 - [37] Frank Neumann and Carsten Witt. 2010. *Bioinspired Computation in Combinatorial Optimization – Algorithms and Their Computational Complexity*. Springer.
 - [38] Chao Qian, Chao Bian, Wu Jiang, and Ke Tang. 2019. Running time analysis of the $(1 + 1)$ -EA for OneMax and LeadingOnes under bit-wise noise. *Algorithmica* 81 (2019), 749–795.
 - [39] Chao Qian, Chao Bian, Yang Yu, Ke Tang, and Xin Yao. 2021. Analysis of noisy evolutionary optimization when sampling fails. *Algorithmica* 83 (2021), 940–975.
 - [40] Chao Qian, Yang Yu, Ke Tang, Yaochu Jin, Xin Yao, and Zhi-Hua Zhou. 2018. On the effectiveness of sampling for evolutionary optimization in noisy environments. *Evolutionary Computation* 26 (2018), 237–267.
 - [41] Jonathan E. Rowe and Aishwaryaprajna. 2019. The benefits and limitations of voting mechanisms in evolutionary optimisation. In *Foundations of Genetic Algorithms, FOGA 2019*. ACM, 34–42.
 - [42] Jonathan E. Rowe and Dirk Sudholt. 2014. The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theoretical Computer Science* 545 (2014), 20–38.
 - [43] Günter Rudolph. 1997. *Convergence Properties of Evolutionary Algorithms*. Verlag Dr. Kováč.
 - [44] Dirk Sudholt. 2021. Analysing the robustness of evolutionary algorithms to noise: refined runtime bounds and an example where noise is beneficial. *Algorithmica* 83 (2021), 976–1011.
 - [45] Dirk Sudholt and Christian Thyssen. 2012. A simple ant colony optimizer for stochastic shortest path problems. *Algorithmica* 83 (2021), 940–975.
 - [46] Petar M. Vasić and Dragoslav S. Mitrović. 2012. *Analytic Inequalities*. Springer Berlin Heidelberg.
 - [47] Carsten Witt. 2006. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation* 14 (2006), 65–86.
 - [48] Weijie Zheng and Benjamin Doerr. 2023. From understanding genetic drift to a smart-restart mechanism for estimation-of-distribution algorithms. *Journal of Machine Learning Research* 24 (2023), 1–40.
 - [49] Zhi-Hua Zhou, Yang Yu, and Chao Qian. 2019. *Evolutionary Learning: Advances in Theories and Algorithms*. Springer.