Scalable Speed-ups for the SMS-EMOA from a Simple Aging Strategy

Mingfeng Li[1] , **Weijie Zheng**[1*] , **Benjamin Doerr**[2]
[1] School of Computer Science and Technology, National Key Laboratory of Smart Farm Technologies and Systems, International Research Institute for Artificial Intelligence,
Harbin Institute of Technology, Shenzhen
[2]Laboratoire d'Informatique (LIX), CNRS, École Polytechnique,
Institut Polytechnique de Paris, Palaiseau, France
limingfeng@stu.hit.edu.cn, zhengweijie@hit.edu.cn, doerr@lix.polytechnique.fr

Abstract

Different from single-objective evolutionary algorithms, where non-elitism is an established concept, multi-objective evolutionary algorithms almost always select the next population in a greedy fashion. In the only notable exception, a stochastic selection mechanism was recently proposed for the SMS-EMOA and was proven to speed up computing the Pareto front of the bi-objective jump benchmark with problem size n and gap parameter k by a factor of $\max\{1, 2^{k/4}/n\}$. While this constitutes the first proven speed-up from non-elitist selection, suggesting a very interesting research direction, it has to be noted that a true speed-up only occurs for $k \geq 4\log_2(n)$, where the runtime is super-polynomial, and that the advantage reduces for larger numbers of objectives as shown in a later work. In this work, we propose a different non-elitist selection mechanism based on aging, which exempts individuals younger than a certain age from a possible removal. This remedies the two shortcomings of stochastic selection: We prove a speed-up by a factor of $\max\{1, \Theta(k)^{k-1}\}$, regardless of the number of objectives. In particular, a positive speed-up can already be observed for constant k, the only setting for which polynomial runtimes can be witnessed. Overall, this result supports the use of non-elitist selection schemes, but suggests that aging-based mechanisms can be considerably more powerful than stochastic selection mechanisms.

1 Introduction

In many real-world optimization problems, several objectives have to be optimized simultaneously rather than in isolation, since improving one performance metric often leads to the degradation of another. For such multi-objective optimization problems, multi-objective evolutionary algorithms (MOEAs) have been employed with great success [Zhou *et al.*, 2011]. For example, the famous NSGA-II algorithm [Deb *et al.*, 2002], the most widely used MOEA, has received more than 50,000 citations on Google scholar.

Interestingly, different from single-objective optimization, these MOEAs almost always select the next population in a greedy fashion. In the NSGA-II framework, for example, the parent and offspring populations are divided into several fronts by non-dominated sorting and then the first fronts are selected into the next population. As tie-breaker for the critical front, the crowding distance is used, and again the individuals with largest crowding distance survive. Similarly, the NSGA-III [Deb and Jain, 2014] and SMS-EMOA [Beume *et al.*, 2007] extend this same two-stage, greedy method. These two algorithms greedily select according to non-dominated sorting, using references points or the hypervolume contribution as the secondary (greedy) selection criteria.

This greedy behavior in multi-objective evolutionary computation differs quite significantly from single-objective evolutionary computation, where concepts such as *selection pressure* are explicitly used to quantify the degree of greediness of the selection. Furthermore, classic selection operators such as tournament selection or roulette-wheel selection, are intentionally designed to be somewhat less greedy than truncation selection.

The first to substantially challenge the believe in greedy selection in multi-objective evolutionary computation were Bian et al. [2023]. In the recently very successful runtime analysis methodology [Zheng *et al.*, 2022; Bian and Qian, 2022; Doerr and Qu, 2023a; Doerr and Qu, 2023b; Doerr and Qu, 2023c; Dang *et al.*, 2023; Bian *et al.*, 2023; Dinot *et al.*, 2023; Wietheger and Doerr, 2023; Zheng and Doerr, 2023a; Zheng and Doerr, 2024a; Zheng and Doerr, 2024b; Zheng *et al.*, 2024b; Opris *et al.*, 2024; Ren *et al.*, 2024; Doerr *et al.*, 2025; Alghouass *et al.*, 2025; Doerr *et al.*, 2025; Deng *et al.*, 2025; Opris, 2025a; Opris, 2025b], they designed a stochastic population update for the SMS-EMOA. Instead of performing a greedy selection over the whole combined parent and offspring population, they let (roughly) a random half of the individuals survive irrespective of their quality and conduct the usual selection of the SMS-EMOA only in the other half. With this non-greedy selection mechanism, they obtained a proven speed-up by a factor of $\max\{1, 2^{k/4}/n\}$ for the bi-objective OJZJ benchmark with problem size n and gap parameter k (this speed-up was later improved to

*Corresponding author.

$\max\{1, 2^k/n\}$ in [Zheng and Doerr, 2024b]).[1]

Following the exciting results on the stochastic population update of Bian et al. [2023], Zheng and Doerr [2024b] (among other results) discussed how these results extend to more than two objectives. For the m-objective mOJZJ benchmark, they proved a speed-up factor of (roughly) $\max\{1, 2^k/n^{m/2}\}$. Hence the proven speed-up from the stochastic population update vanishes for larger numbers of objectives. Till now, it remains an open question whether a non-elitist survival strategy can result in scalable speed-ups for all number of objectives.

Our contributions: In this paper, we take such an attempt and design a simple aging strategy for the population update process of the SMS-EMOA. We initialize each individual with the age limit τ. Any newly generated individual has an age of 0. Only individuals with an age of at least τ will be subject to the original survival selection of the SMS-EMOA. After the removal, all remaining individuals will increase their ages by 1 and enter into the next generation. This simple aging strategy allows any newly generated individual, irrespective of its quality, to survive for at least τ iterations.

For the SMS-EMOA with this selection strategy, we prove that the runtime for covering the full Pareto front of the bi-objective OJZJ benchmark is $O(n^{k+1}/\Theta(k)^{k-1})$, see Theorem 5, comparing favorably with the best known guarantee of $O(n^{k+1})$ for the classic SMS-EMOA [Bian *et al.*, 2023]. For the m-objective version mOJZJ, we show a guarantee of $O(\overline{M}kmn^k/\Theta(k)^k)$, see Theorem 10, and note that the best known upper bound for the classic SMS-EMOA is $O(\overline{M}mn^k)$ by Wietheger and Doerr [2024], where $\overline{M}$ denotes the size of the largest set of pairwise non-dominating solutions of the problem. Hence, a speed-up factor of $\max\{1, \Theta(k)^{k-1}\}$ has been obtained for the m-objective mOJZJ benchmark.

Our experimental results further support these theoretical findings, demonstrating that even for small values of k in both OJZJ and mOJZJ, the aging strategy significantly accelerates the SMS-EMOA algorithm compared to both the original version and the variant with the stochastic population update strategy.

The rest of the paper is organized as follows. Section 2 introduces the basic concept in multi-objective optimization and presents the SMS-EMOA algorithm, including its variant with the stochastic population update strategy. Section 3 introduces our simple aging strategy and gives the basic behavior of this variant when optimizing an m-objective optimization problem. In Section 4 and Section 5, we conduct the runtime analyses of the SMS-EMOA with the aging strategy on OJZJ and mOJZJ. Section 6 presents the experiments, and finally Section 7 concludes our paper.

[1]Another non-elitist attempt was studied by Zheng et al. [2024b]. Inspired by the theoretical advantage of the Metropolis algorithm for the single-objective DLB benchmark [Wang *et al.*, 2024], they discussed how to use the Metropolis algorithm for multiple objectives. While that work could detect strengths and weaknesses of different variants of the multi-objective Metropolis algorithm, compared to other MOEAs none of these variants could achieve better theoretical runtime guarantees for the bi-objective DLB benchmark.

2 Preliminaries

2.1 Multi-Objective Optimization

In multi-objective optimization, the goal is to find a set of optimal trade-off solutions that effectively balance multiple conflicting objectives, typically represented as a Pareto set. As common in the theory community, this paper considers the multi-objective pseudo-Boolean maximization problems, that is, to maximize $f(x) = (f_1(x), f_2(x), \ldots, f_m(x))$ subject to $x \in \{0,1\}^n$ where $m \in \mathbb{N}$ represents the number of objectives, and $n \in \mathbb{N}$ denotes the problem size. A key concept in multi-objective optimization is Pareto dominance as not all solutions are comparable. We say that a solution $x_1 \in \{0,1\}^n$ *weakly dominates* another solution $x_2 \in \{0,1\}^n$ (denoted as $x_1 \succeq x_2$) if $f_i(x_1) \geq f_i(x_2)$ for all $i \in \{1, \ldots, m\}$. In this case, x_1 is at least as good as x_2 in all objectives. If, in addition, x_1 is strictly better than x_2 in at least one objective (i.e., $f_i(x_1) > f_i(x_2)$ for some i), then we call x_1 *dominates* x_2 (denoted as $x_1 \succ x_2$). A solution x is called *Pareto optimal* if it is not dominated by any other solutions in $\{0,1\}^n$. The set of all Pareto optimal solutions is the *Pareto set*, and the corresponding function values form the *Pareto front*.

In the theory of MOEAs, the *runtime* is typically defined as the number of function evaluations required for the algorithm's population to fully cover the Pareto front [Auger and Doerr, 2011; Zhou *et al.*, 2019; Doerr and Neumann, 2020]. Additionally, in this paper, we will use $|x|_1$ to denote the exact number of ones in x, $|x|_0$ to denote the exact number of zeros in x, and $[a..b]$ to represent the set $\{a, a+1, ..., b\}$ for $a \leq b$ and $a, b \in \mathbb{Z}$. We also use $\overline{M}$ to denote the maximum size of a set of incomparable solutions for a given optimization problem.

2.2 The SMS-EMOA and Stochastic Population Update

The SMS-EMOA is a steady-state variant of the NSGA-II, which replaces the crowding distance by the hypervolume contribution as the secondary selection criterion. This algorithm operates with a fixed population size μ. In each iteration, a single offspring is generated and added to the combined parent and offspring population R_t. To maintain the fixed population size, one individual from R_t will be removed. It first uses the non-dominated sorting procedure to partition R_t into several fronts $F_1, \ldots, F_{i^*}$, where F_i contains all non-dominated individuals in $R_t \setminus \cup_{j=1}^{i-1} F_j$. Within the critical front F_{i^*}, the individual that contributes the least to the hypervolume will be removed (broken tie randomly). The hypervolume of a set S of individuals w.r.t. reference point r is calculated as $\mathrm{HV}_r(S) = \mathcal{L}\big(\bigcup_{u \in S}\{h \in \mathbb{R}^m \mid r \leq h \leq f(u)\}\big)$ where $\mathcal{L}$ represents the Lebesgue measure. The hypervolume contribution of an individual $x \in F_{i^*}$ is defined by $\Delta_r(x, F_i^*) := \mathrm{HV}_r(F_i^*) - \mathrm{HV}_r(F_i^* \setminus \{x\})$. See Algorithm 1 for the whole procedure of the classic SMS-EMOA.

The stochastic population update of the SMS-EMOA, proposed in [Bian *et al.*, 2023], only has one change compared to the classic SMS-EMOA. It uniformly at random selects half of the combined population for the survival selection, and the

Algorithm 1: SMS-EMOA

Generate P_0 by selecting μ solutions uniformly and randomly from $\{0,1\}^n$ with replacement;
for $t = 0, 1, 2, \ldots,$ **do**
  Select a solution x uniformly at random from P_t;
  Generate x' by flipping each bit of x independently with probability $1/n$;
  Use fast-non-dominated-sort() [Deb *et al.*, 2002] to divide $R_t = P_t \cup \{x'\}$ into $F_1, \ldots, F_{i^*}$;
  Calculate $\Delta_r(z, F_{i^*})$ for all $z \in F_{i^*}$ and find $D = \arg\min_{z \in F_{i^*}} \Delta_r(z, F_{i^*})$;
  Uniformly at random pick $z' \in D$ and $P_{t+1} = R_t \setminus \{z'\}$;

Algorithm 2: SMS-EMOA with the Stochastic Population Update

Generate P_0 by selecting μ solutions uniformly and randomly from $\{0,1\}^n$ with replacement;
for $t = 0, 1, 2, \ldots,$ **do**
  Select a solution x uniformly at random from P_t;
  Generate x' by flipping each bit of x independently with probability $1/n$;
  $R_t = \lfloor P_t \cup \{x'\}/2 \rfloor$ solutions uniformly and randomly selected from $P_t \cup \{x'\}$ without replacement;
  Use fast-non-dominated-sort() [Deb *et al.*, 2002] to divide R_t into $F_1, \ldots, F_{i^*}$;
  Calculate $\Delta_r(z, F_{i^*})$ for all $z \in F_{i^*}$ and find $D = \arg\min_{z \in F_{i^*}} \Delta_r(z, F_{i^*})$;
  Uniformly at random pick $z' \in D$ and $P_{t+1} = P_t \cup \{x'\} \setminus \{z'\}$;

other half directly enters into the next generation. With this strategy, the inferior solutions have the chance (like not being chosen to the survival selection) to the next generation. For a clear comparison with our proposed strategy, we also state its procedure in Algorithm 2.

As introduced before, we note here once again that a speed-up by a factor of $\max\{1, 2^{k/4}/n\}$ is proven in [Bian *et al.*, 2023] (which is subsequently proven to be $\max\{1, 2^k/n\}$ in [Zheng and Doerr, 2024b]) for the bi-objective OJZJ benchmark. We also note here, as introduced before as well, that a speed-up factor of only $\Theta(1)$ for large m (say $m > k$) is proven for the m-objective mOJZJ benchmark [Zheng and Doerr, 2024b].

3 The SMS-EMOA with the Aging Strategy

As mentioned before, the non-elitist stochastic population update cannot scale well from two objectives to more objectives. This section will introduce our non-elitist aging strategy and employ it into the SMS-EMOA for the easy comparison against the stochastic population update [Bian *et al.*, 2023] in the later sections.

Algorithm 3: SMS-EMOA with the Aging Strategy

Initialize P_0 with μ individuals chosen independently and uniformly at random from $\{0,1\}^n$ with replacement and set their age to τ;
for $t = 0, 1, 2, \ldots,$ **do**
  Select $x \in P_t$ uniformly at random;
  Generate y from x via standard bit-wise mutation and set its age $y.age = 0$;
  $R_t = \{x \mid x \in P_t \cup \{y\} \text{ and } x.\text{age} \geq \tau\}$;
  Use fast-non-dominated-sort() to divide R_t into $F_1, \ldots, F_{i^*}$;
  Let $D = \arg\min_{z \in F_i^*} \Delta_r(z, F_{i^*})$;
  Uniformly at random pick $z' \in D$;
  $P_{t+1} = P_t \cup \{y\} \setminus \{z'\}$;
  Raise the ages of solutions in P_{t+1} by 1;

3.1 Aging Strategy

Aging has been used in different kinds of randomized search heuristics, such as in evolutionary algorithms and artificial immune systems, e.g., [Horoba *et al.*, 2009; Jansen and Zarges, 2009; Jansen and Zarges, 2010a; Jansen and Zarges, 2010b; Jansen and Zarges, 2011b; Jansen and Zarges, 2011a; Oliveto and Sudholt, 2014; Corus *et al.*, 2019; Corus *et al.*, 2020]. Aging operators assign each solution in the population an individual age, which increases by 1 with each generation. Typically a maximum lifespan τ is defined, and any solution exceeding this age is removed from the population [Zarges, 2020]. We resort to this strategy but with some modifications. Instead of the immediate removal, only individuals with the age at least τ will join the survival selection and have the chance to be removed. This is the key simple aging strategy we use to let the inferior solutions have the chance to enter into the next generations, making the algorithm not shortsighted.

In this work, we choose the SMS-EMOA to equip with this strategy, since the efficiency of the SMS-EMOA with the non-elitist stochastic population update is proven for bi-objective OJZJ [Bian *et al.*, 2023]. In detail, we first initialize μ individuals randomly, and assign all individuals with the age limit τ. In each generation, one offspring individual is generated and assigned with the age of 0. Among the $\mu + 1$ combined parent and offspring individuals, the individuals with the age at least τ will be chosen to form the multiset R_t, to which the original survival selection will be applied. After the survival selection, the ages of all remaining individuals in the population are incremented by 1. All other procedures in the algorithm remain consistent with the original SMS-EMOA. The details are shown in Algorithm 3. Obviously, this strategy ensures that any newly generated individual, including the inferiors, can survive for at least τ generations.

3.2 Basic Behavior

The following lemma gives the basic behavior of the SMS-EMOA with the aging strategy. That is, if the size of the population is at least the same as the size of largest set of incomparable solutions plus τ, that is, $\overline{M} + \tau$, then any individual in the current population will have future individuals that weakly dominate it. Due to the space limitation, all proofs are omitted but will be available in an arXiv preprint [Li *et al.*, 2025].

Lemma 1. *Consider any m-objective optimization problem. Consider using the SMS-EMOA with the aging strategy and with the size of the population $\mu \geq \overline{M} + \tau$ to solve this problem. If P_t contains a solution x, then at any later time $t' > t$, the population $P_{t'}$ will contain a solution y such that $y \succeq x$. In particular, any Pareto front point once reached will be maintained in all future generations.*

This important behavior of not being worse will be frequently used for the runtime analysis for the Pareto front coverage.

4 Runtime for Bi-Objective OJZJ

Recall that the stochastic population update shows a speed-up by a factor of $\max\{1, 2^{\Theta(k)}/n\}$ for the SMS-EMOA optimizing the bi-objective OJZJ benchmark [Bian *et al.*, 2023]. In this section, we will also show that our proposed aging strategy can speed up the SMS-EMOA by a factor of $\max\{1, \Theta(k)^{k-1}\}$, which is also faster than the one with the stochastic population update by a factor of $\max\{1, \Theta(k)^k n/k\}$.

4.1 OJZJ

JUMP functions [Droste *et al.*, 2002] are the most extensively studied multimodal benchmark in the theory of randomized search heuristics, particularly in evolutionary algorithms [Bambury *et al.*, 2024]. They represent a class of problems with a tunable gap parameter k, which reflects different depths from the local optima to the global optimum. The OJZJ benchmark proposed by Doerr and Zheng [2021] is a bi-objective counterpart of the JUMP functions with problem size n and jump size k, which has been widely used to understand the theoretical behavior of the MOEAs for multimodal problems [Doerr and Qu, 2023a; Doerr and Qu, 2023b; Doerr and Qu, 2023c; Bian *et al.*, 2023; Doerr *et al.*, 2025]. The OJZJ benchmark involves two objectives, that is, one is the JUMP benchmark, and the other is the JUMP function applied to $\bar{x} = 1 - x$. See the formal definition in the following.

Definition 2 ([Doerr and Zheng, 2021]). *Let $n \in \mathbb{N}$ and $k = [1..n]$. The function $\mathrm{OJZJ}_{n,k} = (f_1, f_2) : \{0,1\}^n \to \mathbb{R}^2$ is defined by*

$$f_1(x) = \begin{cases} k + |x|_1, & \text{if } |x|_1 \leq n - k \text{ or } x = 1^n, \\ n - |x|_1, & \text{else}; \end{cases}$$

$$f_2(x) = \begin{cases} k + |x|_0, & \text{if } |x|_0 \leq n - k \text{ or } x = 0^n, \\ n - |x|_0, & \text{else}. \end{cases}$$

Figure 1 illustrates the characteristics of the OJZJ benchmark. As proven in [Doerr and Zheng, 2021], the Pareto set is $\{x \mid |x|_1 \in [k..n-k] \cup \{0, n\}\}$ and the Pareto front is $\{(a, 2k + n - a) \mid a \in [2k..n] \cup \{k, n+k\}\}$. If $k \leq n/2$, the size of the largest set of incomparable individuals, $\overline{M}$, is $n - 2k + 3$. Since the objective values of OJZJ are all greater than 0, we set the reference point for the SMS-EMOA to $(-1, -1)$. As mentioned before, we will calculate the runtime to cover the full Pareto front.

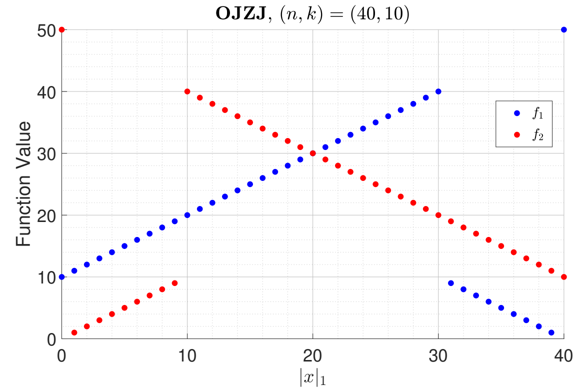

Figure 1: OJZJ with $\{n, k\} = \{40, 10\}$.

4.2 Runtime

The optimization process can be divided into two phases. The first phase ends when all inner Pareto front points $\{(a, 2k + n - a) \mid a \in [2k..n]\}$ are covered. Then the second phase starts, and ends when the full Pareto front is covered for the first time. Lemma 3 considers the runtime of the first phase, showing that after $O(\mu n \log n) + k\tau \left(\frac{2e}{1-\exp(-\frac{\tau}{\mu})}\right)^k$ iterations in expectation, the population will cover the inner Pareto front. The proof idea is that once the population includes at least one solution in the inner Pareto optimal set $\{x \mid |x|_1 \in [k..n-k]\}$, then its Hamming neighbor belonging to the inner Pareto optima can be generated by flipping one of the 0-bits or one of the 1-bits. According to Lemma 1, any element of Pareto front will be retained. By repeating this process and selecting the appropriate individuals, the entire inner Pareto front can be covered.

Lemma 3. *Let $k \leq n/2$. Consider using the SMS-EMOA with the aging strategy and with the size of the population $\mu \geq \overline{M} + \tau$ to optimize the* OJZJ *problem, then the expected number of iterations is at most $O(\mu n \log n) + k\tau \left(\frac{2e}{1-\exp(-\frac{\tau}{\mu})}\right)^k$ to cover all inner Pareto front points.*

The following lemma rigorously provides the runtime analysis of the algorithm for the critical second phase, which starts with a population already covering the inner Pareto front and ends when the full Pareto front is covered. In other words, during this phase, we need to find the two extreme solutions 1^n and 0^n. The proof uses a waiting time argument. That is, we first calculate the probability p of generating the desired extreme solution 1^n (or symmetrically 0^n) from the current boundary of the inner Pareto front, within at most $k\tau$ iterations. Subsequently we bound the expected waiting time by $\frac{k\tau}{p}$. More precisely, we divide $k\tau$ iterations into k steps and within each step of at most τ generation, we calculate the probability of making an improvement. Note that the aging strategy ensures that any newly generated solution survives for at least τ generations. Then we easily obtain the overall probability p via multiplying the probabilities in all steps for their independence, and thus the runtime argument is obtained.

Lemma 4. *Let $k \le n/2$. Consider using the SMS-EMOA with the aging strategy and with the size of the population $\mu \ge \overline{M} + \tau$ to optimize* OJZJ. *Assume that the current population covers all inner Pareto front points. Then the full Pareto front will be covered in at most $k\tau\left(\frac{e^2 n}{(1-\exp(-\frac{\tau}{\mu}))k}\right)^k$ iterations in expectation.*

Combining Lemmas 3 and 4, we derive the runtime of the SMS-EMOA with the aging strategy stated in the following theorem.

Theorem 5. *Let $k \le n/2$. Consider using the SMS-EMOA with the aging strategy and with the size of the population $\mu \ge \overline{M} + \tau$ to optimize* OJZJ*, then the expected iterations to cover the full Pareto front is at most*

$$O(\mu n \log n) + \frac{k\tau\left((2e)^k + \left(\frac{e^2 n}{k}\right)^k\right)}{\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)^k}$$

$$= O\left(k\tau\left(\frac{e^2 n}{\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)k}\right)^k\right).$$

Note that the expected runtime of the original SMS-EMOA is $O(n^{k+1})$ for $\mu = \Theta(n)$ [Bian *et al.*, 2023]. We set $\tau = \Theta(n)$ and $\mu = \Theta(n - 2k + 4 + \tau)$ to also ensure $\mu = \Theta(n)$. In this setting, from Theorem 5, we know that the expected iterations for the SMS-EMOA with the aging strategy to cover the full Pareto front is $O(n^{k+1}/(\Theta(k))^{k-1})$. Hence, we see a speed-up by a factor of $\max\{1, \Theta(k)^{k-1}\}$. Also noting $O(n^{k+2}/2^{\Theta(k)})$ expected number of iterations for the stochastic population update [Bian *et al.*, 2023; Zheng and Doerr, 2024b], we see our aging strategy surpasses this by a factor of $\max\{1, \Theta(k)^k n/k\}$.

5 Runtime for m-Objective mOJZJ

In the previous section, we demonstrated that the SMS-EMOA with the aging strategy outperforms the original SMS-EMOA and the one with the stochastic population update strategy on the bi-objective OJZJ problem. As pointed out before, Zheng and Doerr [2024b] proved that the speed-up of the stochastic population update mechanism cannot scale well from two objectives to more objectives. In this section, we will show that our simple aging strategy performs well for many objectives, and achieves the speed-up by a factor of $\max\{1, \Theta(k)^{k-1}\}$ against the original SMS-EMOA as well as the one with the stochastic population update strategy.

5.1 mOJZJ

The mOJZJ benchmark proposed by [Zheng and Doerr, 2024b] is a m-objective counterpart of the bi-objective OJZJ benchmark [Doerr and Zheng, 2021], and is the first multimodal many-objective benchmark proposed for theoretical analysis. Zheng and Doerr [2024b] proved that the SMS-EMOA can compute the full Pareto front of this benchmark in an expected number of $O(\mu M n^k)$ iterations, where M denotes the size of the Pareto front. Later, Wietheger and Doerr [2024] proved near-tight runtime guarantees for the SEMO, GSEMO, SMS-EMOA and NSGA-III algorithms on this multimodal benchmark class and other popular unimodal many-objective variants of ONEMINMAX, COCZ, and LOTZ. Zheng, Gao and Doerr [2024a] conducted the mathematical runtime analysis for the NSGA-II with their newly proposed variant of the crowding distance, showing that this algorithm can solve the mOJZJ benchmark efficiently compared with the original NSGA-II. The mOJZJ benchmark involves m objectives. The bit string of length n is divided into $m/2$ blocks, each of length $2n/m$. For each block, a bi-objective OJZJ problem is defined. See the following definition.

Definition 6 ([Zheng and Doerr, 2024b])**.** *Let m be the even number of objectives. Let the problem size n be a multiple of $m/2$. Let $n' = \frac{2n}{m} \in \mathbb{N}$ and $k \in [1..n']$. Let $B_i := [(i-1)n' + 1..in']$ for $i \in [1..m/2]$ denote the i-th block of the n bit positions. Then for any $x = (x_1, \dots, x_n)$, the $m\text{OJZJ}_k = (f_1(x), \dots, f_m(x)) : \{0,1\}^n \to \mathbb{R}^m$ is defined by*

$$(f_{2i}, f_{2i-1}(x)) = \text{OJZJ}(x_{B_i}), i \in [1..m/2],$$

where the OJZJ *function is defined in Definition 2.*

The Pareto set of mOJZJ is $\{x \in \{0,1\}^n \mid \forall i \in [1..m/2], |x_{B_i}|_1 \in [k..n'-k] \cup \{0, n'\}\}$ and the Pareto front is $\{(a_1, n' + 2k - a_1, \dots, a_{m/2}, n' + 2k - a_{m/2}) \mid a_1, \dots, a_{m/2} \in [2k..n'] \cup \{k, n' + k\}\}$. Hence, the size of the Pareto front is $(n' - 2k + 3)^{\frac{m}{2}}$. Different from OJZJ, the size of the largest set of incomparable solutions is not equal to the size of the Pareto front, and a trivial upper bound of $(n'+1)^{\frac{m}{2}}$ is given in [Zheng and Doerr, 2023b] when $k \le n'/2$, $\overline{M} \le (n'+1)^{\frac{m}{2}}$. Here we set the reference point to $(-1, \dots, -1)$ with m dimensions.

5.2 Runtime

We now analyze the runtime of the SMS-EMOA with the aging strategy optimizing mOJZJ. To facilitate the proof, we resort to the technique used in [Wietheger and Doerr, 2024] which uses a union bound to establish a tail bound of the runtime. Before the analysis, we first recall several definitions from [Wietheger and Doerr, 2024]. Let $m' = m/2$. The set $K_{m,k}$ is the set of function values of individuals in which each block has either exactly k 0-bits or k 1-bits. Formally, $K_{m,k} = \left\{(a_1, \dots, a_{m'}) \mid a_i \in \left\{k, \frac{n}{m'} - k\right\} \text{ for all } i \in [m']\right\}$. Similarly, the set $C_{m,k}$ is the set of function values of individuals in which each block consists either all bits being 1, all bits being 0, exactly k 0-bits, or k 1-bits. Formally, $C_{m,k} = \left\{(a_1, \dots, a_{m'}) \mid a_i \in \left\{0, k, \frac{n}{m'} - k, \frac{n}{m'}\right\} \text{ for all } i \in [m']\right\}$.

We divide the optimization process into three phases. The first phase starts with the initialization of the algorithm and concludes when all points in $K_{m,k}$ are covered. The second phase begins thereafter and ends when all points in $C_{m,k}$

are covered. The final phase ends when the remaining Pareto front is covered.

For the first phase and the final phase, we apply Lemma A1 and Lemma A3 from [Wietheger and Doerr, 2024], which provide the runtime analyses of the GSEMO on mOJZJ. These results can be adapted to the SMS-EMOA with the aging strategy when the size of the population $\mu \geq \overline{M} + \tau$. The reason is that Section 5 of [Wietheger and Doerr, 2024] outlines three sufficient properties to prove such upper bounds. The survival guarantee provided by Lemma 1 satisfies the first one for not losing already made progress. As for the remaining two properties, the probability of selecting an individual from the population for mutation is $1/\mu$ and standard bit-wise mutation is employed. Therefore, we obtain the following lemmas for the runtime of the first and final phases.

Lemma 7. *Let $k \leq n'/2$. Consider the SMS-EMOA with the aging strategy and with the size of the population $\mu \geq \overline{M} + \tau$ to optimize mOJZJ. Let T denote the number of iterations until the population covers $K_{m,k}$ and let*

$$t = \left(\ln(2)\frac{m'}{\ln(n)} + 2\right) e\mu(n-k)\ln(n-k).$$

Then $T \leq t$ with probability at least $1 - \frac{1}{n}$.

Lemma 8. *Let $k \leq n'/2$. Consider the SMS-EMOA with the aging strategy and with population size $\mu \geq \overline{M} + \tau$ optimizing mOJZJ starting with a population that covers $C_{m,k}$. Let T denote the number of iterations until the population covers the entire Pareto front and let*

$$t = 2em'\mu \cdot \max\Big\{2\Big(\frac{n}{2m'} - k\Big), 8\ln(m') + 8m'\ln\Big(\frac{n}{m'} - 2k + 3\Big) + 8\ln(n)\Big\}.$$

Then $T \leq \lceil t \rceil$ with probability at least $1 - \frac{1}{n}$.

Now we consider the second phase that begins when all points in $K_{m,k}$ are covered and ends when all points in $C_{m,k}$ are covered. From the analysis of Lemma 4, we know that the aging strategy helps to remain the inferior individuals, thereby increasing the probability of generating individuals with all 1-bits or all 0-bits. The following lemma shows that it continues to provide advantages in traversing fitness valleys for many-objective optimization. We first use the same approach to Lemma 4, calculating the probability of changing a block with $n' - k$ 1-bits (or 0-bits) to the extreme block $1^{n'}$ (or $0^{n'}$) within at most $k\tau$ iterations. Then based on this, we calculate the tail bound for the time required to find the extreme solution 1^n which has m' extreme blocks. A union bound is further used over the tail bounds to bound the time required to cover $C_{m,k}$.

Lemma 9. *Let $k \leq n'/2$. Consider the SMS-EMOA with the aging strategy and with population size $\mu \geq \overline{M} + \tau$ optimizing mOJZJ starting with a population that covers $K_{m,k}$. Let T denote the number of iterations until the population covers $C_{m,k}$ and let*

$$t = \left(1 + \frac{\ln(4)m' + \ln(n)}{\ln(m')}\right)\frac{\ln(m')k\tau}{2} \cdot \left(\frac{e^2 n}{k\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)}\right)^k.$$

Then $T \leq t$ with probability at least $1 - \frac{1}{n}$. Further,

$$E[T] \leq \left(1 - \frac{1}{m'}\right)^{-1}\left(2 + \frac{\ln(4)m'}{\ln(m')}\right)\frac{\ln(m')k\tau}{2} \cdot \left(\frac{e^2 n}{k\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)}\right)^k.$$

Combining Lemmas 7 to 9, similarly using a simple restart argument to obtain the expected value for the first and final phase, we can easily derive the runtime of the SMS-EMOA with the aging strategy in the following theorem.

Theorem 10. *Let $k \leq n'/2$. Consider the SMS-EMOA with the aging strategy that $\mu \geq \overline{M} + \tau$ to optimize mOJZJ. Let T denote the number of iterations until the population covers the Pareto front and let*

$$t = \left(1 + \frac{\ln(4)m' + \ln(n)}{\ln(m')}\right)\frac{3\ln(m')k\tau}{2} \cdot \left(\frac{e^2 n}{k\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)}\right)^k.$$

Then $T \leq t$ with high probability. Further,

$$E[T] \leq \left(1 - \frac{1}{m'}\right)^{-1}\left(2 + \frac{\ln(4)m'}{\ln(m')}\right)\frac{3\ln(m')k\tau}{2} \cdot \left(\frac{e^2 n}{k\left(1 - \exp\left(-\frac{\tau}{\mu}\right)\right)}\right)^k.$$

With $\tau = \Theta(\overline{M})$ and population size $\mu = \Theta(\overline{M} + \tau)$, the SMS-EMOA with the aging strategy requires $O(\overline{M}km(n/\Theta(k))^k)$ iterations in expectation to cover the full Pareto front of mOJZJ. Compared to the runtime guarantee $O(\overline{M}mn^k)$ proved in [Wietheger and Doerr, 2024] for the original SMS-EMOA, the aging strategy achieves a speed-up by a factor of $\max\{1, \Theta(k)^{k-1}\}$. Together with Theorem 5 for bi-objective OJZJ, we see that our aging strategy achieves a scalable speed-up factor of $\max\{1, \Theta(k)^{k-1}\}$ for all number of objectives. It is a nice property that the stochastic population update strategy doesn't have (Note the speed-up reduces significantly for increasing the number of objectives, as mentioned before). Besides, a speed-up factor of $\max\{1, \Theta(k)^{k-1}\}$ for our aging strategy is also obtained compared to the stochastic population update.

6 Experiments

This section conducts experiments to intuitively see the efficiency of the aging strategy for two and more objectives. For comparison, we also include experiments for the original SMS-EMOA and the one with the stochastic population update.

For the bi-objective problem, we chose OJZJ as characterized in Theorem 5. We set the size of the problem $n \in \{10, 15, 20, 25, 30\}$ and fix the gap parameter at $k = 4$ in order to see whether the proven asymptotic runtime advantages hold for small and medium problem sizes. We set the age limit $\tau = \mu/2$ and the size of the population to $2(n-2k+4)$, which yields the optimal asymptotic bounds established in Theorem 5.

Each algorithm was tested with 50 independent runs, and terminated when the full Pareto front was covered for the first time. Figure 2 illustrates the mean (with standard deviations) number of iterations for each algorithm to cover the full Pareto front for the first time. From Figure 2, we can easily see a notable speed-up of about 7 of the aging strategy compared to the original SMS-EMOA, as well as a speed-up of around 5 compared to the stochastic population update strategy. These results clearly indicate that the superiority of the aging strategy already appears for small problem sizes and small gap size .

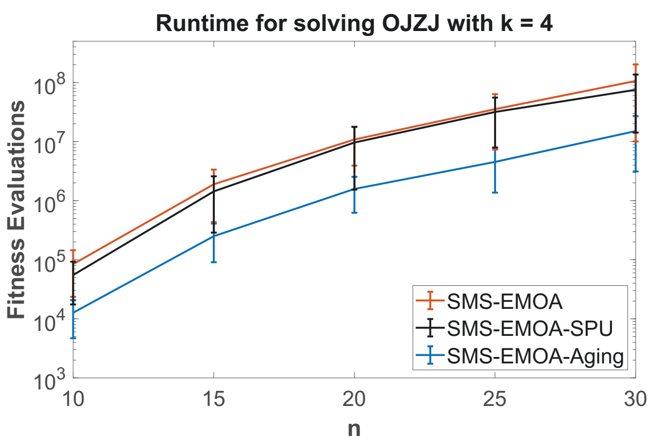

Figure 2: The mean (with standard deviations) number of fitness evaluations of the SMS-EMOA with different mechanisms for solving OJZJ with $k = 4$ and $n \in \{10, 15, 20, 25, 30\}$ in 50 independent runs.

For many-objective optimization, we chose mOJZJ as in Theorem 10, and we fixed the number of objectives to $m = 4$. We do not set larger numbers of objectives, since (i) experiments with $m = 4$ finished in a reasonable time and (ii) with $m = 4$ we intend to see whether our aging strategy can result in a good speed-up even for small number of objectives. We set the problem size $n \in \{12, 16, 20, 24, 28\}$ and gap size $k = 3$ to see whether the asymptotic results hold for small and medium problem sizes. The age limit τ was still set to $\mu/2$ and the population size was set to $2\Big((n'+1)^{m/2}+1\Big)$ for the best asymptotic runtime in Theorem 10. Due to the large number of fitness evaluations, each algorithm was tested with 20 independent runs, which we consider also sufficient for a reliable performance comparison, and terminated when the full Pareto front was covered for the first time. From Figure 3, we easily see the clear superiority of the SMS-EMOA with the aging strategy against the original one and the one with the stochastic population update. Note that the theoretical speed-up factor against the other two is $\Theta(k)^{k-1}$ from Theorem 10. For this setting, $k^{k-1} = 3^2 = 9$. From the plot, we know such speed-up factor is around 3 compared to the original SMS-EMOA and is around 2 compared to the one with the stochastic population update strategy, which means that the theoretical speed-up is already witnessed for small problem sizes, small gap size, and small number of objectives.

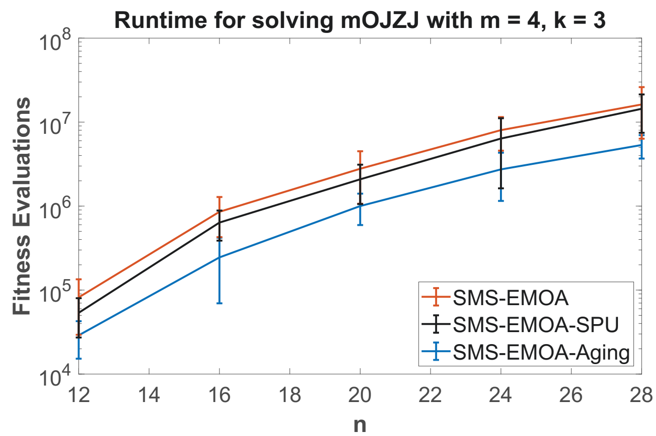

Figure 3: The mean (with standard deviations) number of fitness evaluations of the SMS-EMOA with different mechanisms for solving mOJZJ with $m = 4, k = 3$ and $n \in \{12, 16, 20, 24, 28\}$ in 20 independent runs.

7 Conclusion

In this paper, to tackle the quest for a scalably efficient non-elitist selection strategy for MOEAs, we proposed a simple aging strategy, essentially giving each new individual τ iterations to develop before being prone to removal, and added it to the SMS-EMOA. A speed-up of $\max\{1, \Theta(k)^{k-1}\}$ compared to the original SMS-EMOA (and also the one with the stochastic population update) is proven for the m-objective JUMP benchmark for all even m. This is the first non-elitist strategy for which a speed-up not vanishing for larger numbers of objectives is proven.

Also for the bi-objective JUMP benchmark, our speed-up of $\max\{1, \Theta(k)^{k-1}\}$ compares favorably with the previously shown speed-up of $\max\{1, 2^k/n\}$ by the stochastic population update strategy, and in particular means that a positive speed-up is seen for small values of k, which are most relevant (only constant k give a polynomial runtime). Our experiments support the above findings.

From our understanding of the mathematical proofs, we conjecture that our simple non-elitist aging strategy can be effective in other MOEAs as well, and list this as our most interesting future work.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No. 62306086, 62350710797), Science, Technology and Innovation Commission of Shenzhen Municipality (Grant No. GXWD20220818191018001), and Guangdong Basic and Applied Basic Research Foundation (Grant No. 2025A1515011936). This research benefited from the support of the FMJH Program PGMO. This work has profited from many fruitful discussions at the Dagstuhl Seminars 23361 *Multiobjective Optimization on a Budget* and 24271 *Theory of Randomized Optimization Heuristics*.

References

[Alghouass *et al.*, 2025] Yasser Alghouass, Benjamin Doerr, Martin S. Krejca, and Mohammed Lagmah. Proven approximation guarantees in multi-objective optimization: SPEA2 beats NSGA-II. In *International Joint Conference on Artificial Intelligence, IJCAI 2025*. ijcai.org, 2025. Accepted.

[Auger and Doerr, 2011] Anne Auger and Benjamin Doerr, editors. *Theory of Randomized Search Heuristics*. World Scientific Publishing, 2011.

[Bambury *et al.*, 2024] Henry Bambury, Antoine Bultel, and Benjamin Doerr. An extended jump functions benchmark for the analysis of randomized search heuristics. *Algorithmica*, 86:1–32, 2024.

[Beume *et al.*, 2007] Nicola Beume, Boris Naujoks, and Michael Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181:1653–1669, 2007.

[Bian and Qian, 2022] Chao Bian and Chao Qian. Better running time of the non-dominated sorting genetic algorithm II (NSGA-II) by using stochastic tournament selection. In *Parallel Problem Solving From Nature, PPSN 2022*, pages 428–441. Springer, 2022.

[Bian *et al.*, 2023] Chao Bian, Yawen Zhou, Miqing Li, and Chao Qian. Stochastic population update can provably be helpful in multi-objective evolutionary algorithms. In *International Joint Conference on Artificial Intelligence, IJCAI 2023*, pages 5513–5521. ijcai.org, 2023.

[Corus *et al.*, 2019] Dogan Corus, Pietro S. Oliveto, and Donya Yazdani. Artificial immune systems can find arbitrarily good approximations for the NP-hard number partitioning problem. *Artificial Intelligence*, 274:180–196, 2019.

[Corus *et al.*, 2020] Dogan Corus, Pietro S. Oliveto, and Donya Yazdani. When hypermutations and ageing enable artificial immune systems to outperform evolutionary algorithms. *Theoretical Computer Science*, 832:166–185, 2020.

[Dang *et al.*, 2023] Duc-Cuong Dang, Andre Opris, Bahare Salehi, and Dirk Sudholt. A proof that using crossover can guarantee exponential speed-ups in evolutionary multi-objective optimisation. In *Conference on Artificial Intelligence, AAAI 2023*, pages 12390–12398. AAAI Press, 2023.

[Deb and Jain, 2014] Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, 18:577–601, 2014.

[Deb *et al.*, 2002] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6:182–197, 2002.

[Deng *et al.*, 2025] Renzhong Deng, Weijie Zheng, and Benjamin Doerr. The first theoretical approximation guarantees for the non-dominated sorting genetic algorithm III (NSGA-III). In *International Joint Conference on Artificial Intelligence, IJCAI 2025*. ijcai.org, 2025. Accepted.

[Dinot *et al.*, 2023] Matthieu Dinot, Benjamin Doerr, Ulysse Hennebelle, and Sebastian Will. Runtime analyses of multi-objective evolutionary algorithms in the presence of noise. In *International Joint Conference on Artificial Intelligence, IJCAI 2023*, pages 5549–5557. ijcai.org, 2023.

[Doerr and Neumann, 2020] Benjamin Doerr and Frank Neumann, editors. *Theory of Evolutionary Computation—Recent Developments in Discrete Optimization*. Springer, 2020. Also available at http://www.lix.polytechnique.fr/Labo/Benjamin.Doerr/doerr_neumann_book.html.

[Doerr and Qu, 2023a] Benjamin Doerr and Zhongdi Qu. A first runtime analysis of the NSGA-II on a multimodal problem. *IEEE Transactions on Evolutionary Computation*, 27:1288–1297, 2023.

[Doerr and Qu, 2023b] Benjamin Doerr and Zhongdi Qu. From understanding the population dynamics of the NSGA-II to the first proven lower bounds. In *Conference on Artificial Intelligence, AAAI 2023*, pages 12408–12416. AAAI Press, 2023.

[Doerr and Qu, 2023c] Benjamin Doerr and Zhongdi Qu. Runtime analysis for the NSGA-II: provable speed-ups from crossover. In *Conference on Artificial Intelligence, AAAI 2023*, pages 12399–12407. AAAI Press, 2023.

[Doerr and Zheng, 2021] Benjamin Doerr and Weijie Zheng. Theoretical analyses of multi-objective evolutionary algorithms on multi-modal objectives. In *Conference on Artificial Intelligence, AAAI 2021*, pages 12293–12301. AAAI Press, 2021.

[Doerr *et al.*, 2025] Benjamin Doerr, Tudor Ivan, and Martin S. Krejca. Speeding up the NSGA-II with a simple tie-breaking rule. In *Conference on Artificial Intelligence, AAAI 2025*, pages 26964–26972. AAAI Press, 2025.

[Droste *et al.*, 2002] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.

[Horoba *et al.*, 2009] Christian Horoba, Thomas Jansen, and Christine Zarges. Maximal age in randomized search heuristics with aging. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 803–810, 2009.

[Jansen and Zarges, 2009] Thomas Jansen and Christine Zarges. Comparing different aging operators. In *Artificial Immune Systems: 8th International Conference, ICARIS 2009, York, UK, August 9-12, 2009. Proceedings 8*, pages 95–108. Springer, 2009.

[Jansen and Zarges, 2010a] Thomas Jansen and Christine Zarges. Aging beyond restarts. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 705–712, 2010.

[Jansen and Zarges, 2010b] Thomas Jansen and Christine Zarges. On the benefits of aging and the importance of details. In *International Conference on Artificial Immune Systems*, pages 61–74. Springer, 2010.

[Jansen and Zarges, 2011a] Thomas Jansen and Christine Zarges. On benefits and drawbacks of aging strategies for randomized search heuristics. *Theoretical Computer Science*, 412:543–559, 2011.

[Jansen and Zarges, 2011b] Thomas Jansen and Christine Zarges. On the role of age diversity for effective aging operators. *Evolutionary Intelligence*, 4:99–125, 2011.

[Li *et al.*, 2025] Mingfeng Li, Weijie Zheng, and Benjamin Doerr. Scalable speed-ups for the SMS-EMOA from a simple aging strategy. In *International Joint Conference on Artificial Intelligence, IJCAI 2025*. ijcai.org, 2025. Accepted.

[Oliveto and Sudholt, 2014] Pietro S. Oliveto and Dirk Sudholt. On the runtime analysis of stochastic ageing mechanisms. In *Conference on Genetic and Evolutionary Computation, GECCO 2014*, pages 113–120, 2014.

[Opris *et al.*, 2024] Andre Opris, Duc Cuong Dang, Frank Neumann, and Dirk Sudholt. Runtime analyses of NSGA-III on many-objective problems. In *Genetic and Evolutionary Computation Conference, GECCO 2024*, pages 1596–1604. ACM, 2024.

[Opris, 2025a] Andre Opris. A first runtime analysis of NSGA-III on a many-objective multimodal problem: provable exponential speedup via stochastic population update. In *International Joint Conference on Artificial Intelligence, IJCAI 2025*. ijcai.org, 2025. Accepted.

[Opris, 2025b] Andre Opris. A many-objective problem where crossover is provably indispensable. In *Conference on Artificial Intelligence, AAAI 2025*, pages 27108–27116. AAAI Press, 2025.

[Ren *et al.*, 2024] Shengjie Ren, Chao Bian, Miqing Li, and Chao Qian. A first running time analysis of the Strength Pareto Evolutionary Algorithm 2 (SPEA2). In *Parallel Problem Solving from Nature, PPSN 2024, Part III*, pages 295–312. Springer, 2024.

[Wang *et al.*, 2024] Shouda Wang, Weijie Zheng, and Benjamin Doerr. Choosing the right algorithm with hints from complexity theory. *Information and Computation*, 296:105125, 2024.

[Wietheger and Doerr, 2023] Simon Wietheger and Benjamin Doerr. A mathematical runtime analysis of the non-dominated sorting genetic algorithm III (NSGA-III). In *International Joint Conference on Artificial Intelligence, IJCAI 2023*, pages 5657–5665. ijcai.org, 2023.

[Wietheger and Doerr, 2024] Simon Wietheger and Benjamin Doerr. Near-tight runtime guarantees for many-objective evolutionary algorithms. In *Parallel Problem Solving from Nature, PPSN 2024, Part IV*, pages 153–168. Springer, 2024.

[Zarges, 2020] Christine Zarges. Theoretical foundations of immune-inspired randomized search heuristics for optimization. *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, pages 443–474, 2020.

[Zheng and Doerr, 2023a] Weijie Zheng and Benjamin Doerr. Mathematical runtime analysis for the non-dominated sorting genetic algorithm II (NSGA-II). *Artificial Intelligence*, 325:104016, 2023.

[Zheng and Doerr, 2023b] Weijie Zheng and Benjamin Doerr. Runtime analysis of the SMS-EMOA for many-objective optimization. *CoRR*, abs/2312.10290, 2023.

[Zheng and Doerr, 2024a] Weijie Zheng and Benjamin Doerr. Runtime analysis for the NSGA-II: proving, quantifying, and explaining the inefficiency for many objectives. *IEEE Transactions on Evolutionary Computation*, 28:1442–1454, 2024.

[Zheng and Doerr, 2024b] Weijie Zheng and Benjamin Doerr. Runtime analysis of the SMS-EMOA for many-objective optimization. In *Conference on Artificial Intelligence, AAAI 2024*, pages 20874–20882. AAAI Press, 2024.

[Zheng *et al.*, 2022] Weijie Zheng, Yufei Liu, and Benjamin Doerr. A first mathematical runtime analysis of the non-dominated sorting genetic algorithm II (NSGA-II). In *Conference on Artificial Intelligence, AAAI 2022*, pages 10408–10416. AAAI Press, 2022.

[Zheng *et al.*, 2024a] Weijie Zheng, Yao Gao, and Benjamin Doerr. A crowding distance that provably solves the difficulties of the nsga-ii in many-objective optimization. *arXiv preprint arXiv:2407.17687*, 2024.

[Zheng *et al.*, 2024b] Weijie Zheng, Mingfeng Li, Renzhong Deng, and Benjamin Doerr. How to use the Metropolis algorithm for multi-objective optimization? In *Conference on Artificial Intelligence, AAAI 2024*, pages 20883–20891. AAAI Press, 2024.

[Zhou *et al.*, 2011] Aimin Zhou, Bo-Yang Qu, Hui Li, Shi-Zheng Zhao, Ponnuthurai Nagaratnam Suganthan, and Qingfu Zhang. Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation*, 1:32–49, 2011.

[Zhou *et al.*, 2019] Zhi-Hua Zhou, Yang Yu, and Chao Qian. *Evolutionary Learning: Advances in Theories and Algorithms*. Springer, 2019.