

# Towards Running Time Analysis of Interactive Multi-Objective Evolutionary Algorithms

Tianhao Lu, Chao Bian, Chao Qian

National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China  
School of Artificial Intelligence, Nanjing University, Nanjing 210023, China  
{luth, bianc, qianc}@lamda.nju.edu.cn

## Abstract

Evolutionary algorithms (EAs) are widely used for multi-objective optimization due to their population-based nature. Traditional multi-objective EAs (MOEAs) generate a large set of solutions to approximate the Pareto front, leaving a decision maker (DM) with the task of selecting a preferred solution. However, this process can be inefficient and time-consuming, especially when there are many objectives or the subjective preference of DM is known. To address this issue, interactive MOEAs (iMOEAs) integrate decision making into the optimization process, i.e., update the population with the help of the DM. In contrast to their wide applications, there have existed only two pieces of theoretical works on iMOEAs, which only considered interactive variants of the two simple single-objective algorithms, RLS and (1+1)-EA. This paper provides the first running time analysis (the essential theoretical aspect of EAs) for practical iMOEAs. Specifically, we prove that the expected running time of the well-developed interactive NSGA-II (called R-NSGA-II) for solving the OneMinMax and OneJumpZeroJump problems are all asymptotically faster than the traditional NSGA-II. Meanwhile, we present a variant of OneMinMax, and prove that R-NSGA-II can be exponentially slower than NSGA-II. These results provide theoretical justification for the effectiveness of iMOEAs while identifying situations where they may fail. Experiments are also conducted to validate the theoretical results.

## Introduction

Multi-objective optimization (MOP) (Steuer 1986) requires to optimize several objective functions simultaneously, which appears in various real-world applications. Since these functions usually conflict with each other, there does not exist a solution which can outperform any other solution in every objective function. Therefore, the goal of MOP is to identify a set of solutions that represent diverse optimal trade-offs between objective functions, referred to as Pareto optimal solutions.

Evolutionary algorithms (EAs) (Bäck 1996) are a large class of randomized heuristic optimization algorithms, inspired by natural evolution. They keep a set of solutions (called population) and iteratively improve it by generating offspring solutions and selecting better ones. The

population-based nature makes EAs highly appropriate for solving MOPs, which have become a popular tool in various real-world MOP situations (Deb et al. 2002; Zhang and Li 2007; Beume, Naujoks, and Emmerich 2007).

After the optimization of multi-objective EAs (MOEAs), a decision maker (DM) needs to select one solution from the generated set of non-dominated solutions based on her (or his) preference. However, the paradigm of finding a large set of solutions at first and then selecting an appropriate one may be inefficient and time-consuming under some circumstances, especially when the number of objectives is large or DM has some subjective preference information. Specifically, when the number of objectives increases, the objective space gets increasingly complex, making it hard to find solutions with good convergence and spread; when the subjective preference of DM is known, such information can actually give a directional bias for the search, while MOEAs usually ignore it and simply search the whole solution space equally, resulting in the waste of computing resources in inessential regions.

Interactive MOEAs (iMOEAs) (Takagi 2001) is thus proposed, which integrates decision making into the optimization process. That is, in the population update procedure of each generation, the DM is used to select the preferred solution out of two incomparable solutions. The preference information provided by the DM can be a utility function (Gong et al. 2011; Köksalan and Karahan 2010), which outputs a function value to indicate the quality of a solution based on the solution's objective vector; a reference point (Huber, Geiger, and Sevaux 2015; Thiele et al. 2009), which represents the DM's favorite objective vector and prefers a solution with objective vector closer to it; and a surrogate model (Li 2019), which directly selects a preferred solution by using a machine learning model, e.g., neural network. Due to their fast convergence and the ability to reflect the DM's preference, iMOEAs have achieved many successful real-world applications, e.g., aircraft design (Bandt and Malinchik 2004), software design (Ramírez, Romero, and Ventura 2018), and land use planning (Janssen et al. 2008).

In contrast to wide applications, the theoretical foundation of MOEAs as well as iMOEAs is still underdeveloped, which is mainly due to the sophisticated behaviors of EAs and the hardness of MOPs. The running time analysis, the essential theoretical aspect, of MOEAs starts from analyzing

the expected running time of simple MOEAs (e.g., SEMO (or GSEMO) which employs the one-bit (or bit-wise) mutation operator to generate an offspring solution in each iteration and keeps the non-dominated solutions generated-so-far in the population) for solving multi-objective synthetic and combinatorial optimization problems (Giel 2003; Laumanns, Thiele, and Zitzler 2004a,b; Neumann 2007; Horoba 2009; Giel and Lehre 2010; Neumann and Theile 2010; Doerr, Kodric, and Voigt 2013; Qian, Yu, and Zhou 2013; Qian, Tang, and Zhou 2016; Dinot et al. 2023). There are also a series of theoretical studies showing that GSEMO or its variants can achieve good approximation guarantees for solving single-objective submodular optimization problems with constraints (Qian, Yu, and Zhou 2015; Friedrich and Neumann 2015; Qian et al. 2017, 2019; Zhou, Yu, and Qian 2019; Qian 2020; Roostapour et al. 2022).

Recently, researchers are shifting their focus to analyzing practical MOEAs. Zheng, Liu and Doerr (2022) analyzed the expected running time of the popular NSGA-II (Deb et al. 2002) solving the OneMinMax and LOTZ problems for the first time. Later, Zheng and Doerr (2022) considered a modified crowding distance method which recalculates the crowding distance each time the solution with the smallest crowding distance is deleted, and proved that it can lead to better approximation to the Pareto front. Bian and Qian (2022) proved that replacing binary tournament selection with stochastic tournament selection in parent selection can significantly decrease the running time of NSGA-II. More works include the analysis of NSGA-II solving the many-objective problem mOneMinMax (Zheng and Doerr 2023) and the bi-objective minimum spanning tree problem (Cerf et al. 2023), NSGA-II solving the LOTZ problem in the presence of noise (Dang et al. 2023b), the effectiveness of the crossover operator (Dang et al. 2023a; Doerr and Qu 2023b), and the lower bound of NSGA-II solving the OneMinMax and OneJumpZeroJump problems (Doerr and Qu 2023a). Besides, the expected running time of other popular MOEAs has also been considered, e.g., MOEA/D (Huang et al. 2019; Huang and Zhou 2020; Huang et al. 2021), SMS-EMOA (Bian et al. 2023), and NSGA-III (Doerr and Wietheger 2023).

Though much effort has been devoted to analyzing MOEAs, the theoretical analysis of iMOEAs is quite rare. To the best of our knowledge, there have been only two pieces of work. Brockhoff et al. (2012) analyzed the expected running time of iRLS and (1+1)-iEA solving two bi-objective problems LOTZ and COCZ, where iRLS and (1+1)-iEA maintain exactly one solution in the population and iteratively improve the solution by generating one offspring solution and keeping the better one. Specifically, the parent solution and offspring solution are first compared according to the Pareto dominance relationship, i.e., the dominating one will be picked; if they are incomparable, the DM will pick the preferred one, i.e., the one with a larger utility function value will be picked. They proved that the performance of (1+1)-iEA and iRLS may degrade dramatically if the utility function is non-linear instead of linear. Neumann and Nguyen (2014) analyzed the expected running time of iRLS and (1+1)-iEA solving the multi-objective

knapsack problem, and proved that using a linear utility function as DM can make iRLS and (1+1)-iEA perform as same as their single-objective counterparts, i.e., RLS and (1+1)-EA (Droste, Jansen, and Wegener 2002), solving the weighted objective sum of knapsack. Though the above two works analyzed (1+1)-iEA and iRLS solving multi-objective problems, the two algorithms are actually the variants of the simple single-objective algorithms RLS and (1+1)-EA, which cannot reflect the common structure of iMOEAs.

In this paper, we make the first attempt towards analyzing the expected running time of practical iMOEAs. We consider R-NSGA-II (Deb and Sundar 2006), which is an interactive variant of the well-known NSGA-II algorithm (Deb et al. 2002), and prefers a solution with a smaller distance to a reference point based on the DM's preference, instead of selecting from incomparable solutions based on their crowding distance. We prove that the expected running time of R-NSGA-II solving OneMinMax and OneJumpZeroJump is  $O(Nn \log n)$  (for any population size  $N$ ) and  $O(n^k)$  (for population size  $N \in O(n^{k/2})$ ), respectively, where  $n$  is the problem size, and  $k$  is the parameter of OneJumpZeroJump. Note that OneMinMax and OneJumpZeroJump are two commonly used bi-objective optimization problems in theoretical analyses of MOEAs (Zheng, Liu, and Doerr 2022; Zheng and Doerr 2023; Doerr and Qu 2023a; Bian et al. 2023). Here solving a problem means finding an optimal solution corresponding to the reference point, which is what the DM really wants. We also prove that the expected running time of NSGA-II solving these two problems is  $O(Nn \log n)$  and  $O(Nn^k)$ , respectively, where the population size  $N$  is required to be  $\Omega(n)$ . Thus, R-NSGA-II using a constant population size can be asymptotically faster than NSGA-II. The main reason is that R-NSGA-II can always keep the solution closest to the reference point in the population, while NSGA-II requires a large  $N$  to ensure this. These results provide theoretical justification for the effectiveness of iMOEAs, and are also verified by experiments.

Next, we present a variant of OneMinMax, and prove that the expected running time of NSGA-II with  $N \geq 4(n+1)$  is  $O(Nn \log n)$ , while R-NSGA-II using any polynomial population size requires  $\Omega(n^{n/4})$  expected running time. Thus, R-NSGA-II can be exponentially slower than NSGA-II. This is because the search direction of R-NSGA-II is guided towards the reference point in the objective space, but the solutions are actually getting far away from the solution corresponding to the reference point in the decision space. This finding may generally hold in complex real-world problems with rugged fitness landscape, e.g., we have verified it by the experiments on the combinatorial problem multi-objective NK-landscape (Aguirre and Tanaka 2007). Thus, our theoretical results suggest that when designing practical iMOEAs, relying solely on DM may be risky, while maintaining some diversity in the population may be helpful.

## Preliminaries

In this section, we first introduce multi-objective optimization, and then present the procedures of NSGA-II and its interactive variant, R-NSGA-II.

## Multi-Objective Optimization

Multi-objective optimization entails the simultaneous optimization of two or more objective functions, as illustrated in Definition 1. In this paper, we focus on maximization, and minimization can be defined similarly. The objectives often exhibit conflicting interests, resulting in the absence of a complete order within the solution space  $S$ . Consequently, comparing solutions often relies on the domination relationship, as presented in Definition 2. A solution is deemed Pareto optimal if no other solution in  $S$  dominates it. The collection of objective vectors of all the Pareto optimal solutions is called the Pareto front. The primary aim of multi-objective optimization is to identify the Pareto front, i.e., find at least one corresponding solution for each objective vector within this front. In this paper, we will study pseudo-Boolean problems, i.e., the solution space  $S = \{0, 1\}^n$ , where  $n$  denotes the problem size.

**Definition 1** (Multi-objective Optimization). *Given a feasible solution space  $S$  and objective functions  $f_1, f_2, \dots, f_m$ , multi-objective optimization can be formulated as*

$$\max_{\mathbf{x} \in S} (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_m(\mathbf{x})).$$

**Definition 2** (Dominance). *Let  $\mathbf{f} = (f_1, f_2, \dots, f_m) : S \rightarrow \mathbb{R}^m$  be the objective vector. For two solutions  $\mathbf{x}$  and  $\mathbf{y} \in S$ :*

- $\mathbf{x}$  weakly dominates  $\mathbf{y}$  (denoted as  $\mathbf{x} \succeq \mathbf{y}$ ) if  $\forall 1 \leq i \leq m, f_i(\mathbf{x}) \geq f_i(\mathbf{y})$ ;
- $\mathbf{x}$  dominates  $\mathbf{y}$  (denoted as  $\mathbf{x} \succ \mathbf{y}$ ) if  $\mathbf{x} \succeq \mathbf{y}$  and  $f_i(\mathbf{x}) > f_i(\mathbf{y})$  for some  $i$ ;
- $\mathbf{x}$  and  $\mathbf{y}$  are incomparable if neither  $\mathbf{x} \succeq \mathbf{y}$  nor  $\mathbf{y} \succeq \mathbf{x}$ .

## NSGA-II

The NSGA-II algorithm (Deb et al. 2002) as presented in Algorithm 1 is a widely-used MOEA known for its incorporation of two substantial features: non-dominated sorting and crowding distance. NSGA-II commences by initializing a population of  $N$  random solutions (line 1). In each generation, it applies bit-wise mutation on each solution in the current population  $P$  to generate  $N$  offspring solutions (lines 4–7). The bit-wise mutation operator flips each bit of an individual’s binary representation independently with a predefined probability, which is  $1/n$  in this paper.

After generating  $N$  offspring solutions, the next generation’s population is determined by selecting the best  $N$  solutions from the current population  $P$  and the offspring population  $Q$  (lines 8–14). These solutions are partitioned into non-dominated sets, denoted as  $F_1, F_2$ , and so on.  $F_1$  contains all the non-dominated solutions in the combined set  $P \cup Q$ , while  $F_i$  ( $i \geq 2$ ) contains all the non-dominated solutions in the set  $(P \cup Q) \setminus \bigcup_{j=1}^{i-1} F_j$ . The solutions from sets  $F_1, F_2$ , and so forth, are sequentially added to the next population (lines 10–12) until the population size reaches or exceeds  $N$ . When the inclusion of a critical set  $F_i$  threatens to exceed the population size limit  $N$ , the crowding distance for each solution in that set is computed. Intuitively, the crowding distance of a solution represents the average distance to its neighbouring solutions in the objective space, and a larger crowding distance implies a larger diversity of the solution. The solutions in the critical set  $F_i$  with larger

---

Algorithm 1: NSGA-II Algorithm (Deb et al. 2002) (without crossover, with fair selection)

---

**Input:** objective functions  $f_1, f_2, \dots, f_m$ , population size  $N$   
**Output:**  $N$  solutions from  $\{0, 1\}^n$

**Process:**

```

1:  $P \leftarrow N$  solutions randomly selected from  $\{0, 1\}^n$ ;
2: while criterion is not met do
3:    $Q = \emptyset$ ;
4:   for each solutions  $\mathbf{x}$  in  $P$  do
5:     apply bit-wise mutation on  $\mathbf{x}$  to generate  $\mathbf{x}'$ ;
6:     add  $\mathbf{x}'$  into  $Q$ 
7:   end for
8:   partition  $P \cup Q$  into non-dominated sets  $F_1, F_2, \dots$ ;
9:   let  $P = \emptyset, i = 1$ ;
10:  while  $|P \cup F_i| < N$  do
11:     $P = P \cup F_i, i = i + 1$ 
12:  end while
13:  assign each solution in  $F_i$  with a crowding distance;
14:  sort the solutions in  $F_i$  by crowding distance in descending order, and add the first  $N - |P|$  solutions into  $P$ 
15: end while
16: return  $P$ 

```

---

crowding distances are selectively included to fill the remaining population slots (line 14). The detailed procedures of non-dominated sorting and crowding distance are provided in the supplementary material due to space limitation.

## R-NSGA-II

NSGA-II, as the most widely applied MOEA, has been modified many times for designing its interactive variants, most of which replace the crowding distance with the DM’s preference information, e.g., characterized by a reference point (Deb and Sundar 2006), reference direction combining preference thresholds (Deb and Kumar 2007), or Gaussian functions (Narukawa, Tanigaki, and Ishibuchi 2014). Because most interactive variants of NSGA-II adopt a reference point as the DM, we consider R-NSGA-II (Deb and Sundar 2006), i.e., Reference-point-based NSGA-II, in this paper, which is widely applied in practice.

R-NSGA-II is similar to NSGA-II, except that the solution selection from the critical set  $F_i$  depends on the distance to the reference point instead of the traditional crowding distance. A solution having a smaller distance to the reference point is favored. Specifically, “a crowding distance” in line 13 of Algorithm 1 changes to “the distance to the reference point”, and “crowding distance in descending order” in line 14 changes to “the distance to the reference point in ascending order”. Note that the distance to the reference point denotes the Euclidean distance in the objective space.

The selection of reference point depends on the specific problem to be solved, and reflects the DM’s preference. For ease of theoretical analysis, we use an objective vector in the Pareto front as a reference point, following previous theoretical analysis of iMOEAs (Brockhoff et al. 2012; Neumann and Nguyen 2014). As the reference point reflects the DM’s

preference, the stopping criterion of an algorithm is set as finding the reference point, i.e., finding at least one solution whose objective vector equals to the reference point. Such a natural setting can avoid any reliance on subjective human factors and ensure objectivity for theoretical analysis.

## R-NSGA-II Can Do Better than NSGA-II

In this section, we analyze the expected running time of NSGA-II and R-NSGA-II for solving two bi-objective problems, i.e., OneMinMax and OneJumpZeroJump. We prove that for these two problems, R-NSGA-II can find the reference point asymptotically faster than NSGA-II.

### Analysis on OneMinMax

The OneMinMax problem as presented in Definition 3 aims to maximize the number of 0-bits and the number of 1-bits of a binary bit string simultaneously. The Pareto front of OneMinMax is

$$F^* = \{(k, n - k) \mid k \in [0..n]\},$$

and any solution  $\mathbf{x} \in \{0, 1\}^n$  is Pareto optimal, corresponding to the objective vector  $(n - |\mathbf{x}|_1, |\mathbf{x}|_1)$  in the Pareto front, where  $|\cdot|_1$  denotes the number of 1-bits of a solution. The size of the Pareto front  $|F^*| = n + 1$ . We let the reference point  $\mathbf{f}(\mathbf{z})$  be any vector  $(k, n - k), k \in [0..n]$ , since the selection of the reference point will not change our results.

**Definition 3** (OneMinMax (Giel and Lehre 2010)). *The OneMinMax problem of size  $n$  is to find  $n$ -bits binary strings which maximize*

$$\mathbf{f}(\mathbf{x}) = \left( n - \sum_{i=1}^n x_i, \sum_{i=1}^n x_i \right)$$

where  $x_i$  denotes the  $i$ -th bit of  $\mathbf{x} \in \{0, 1\}^n$ .

We prove in Theorems 1 and 2 that NSGA-II and R-NSGA-II can both find the reference point in  $O(Nn \log n)$  expected number of fitness evaluations. However, NSGA-II requires the population size  $N$  to be at least  $4(n + 1)$ , while R-NSGA-II does not have any requirement on the population size. Thus, our results show that if using a constant population size, R-NSGA-II can be at least  $n$  times faster than NSGA-II.

The main proof idea of Theorem 1 is that the population of NSGA-II can continuously expand in the Pareto front, until the reference point is found. Since the proof procedure is similar to that of Theorem 2 in (Zheng, Liu, and Doerr 2022), and the derived upper bound is also the same, we omit the proof of Theorem 1.

**Theorem 1.** *For NSGA-II solving the OneMinMax problem, if using a population size  $N$  of at least  $4(n + 1)$ , then the expected number of fitness evaluations for finding the reference point  $(k, n - k), k \in [0..n]$ , is  $O(Nn \log n)$ .*

The proof of Theorem 2 is similar to that of Theorem 1. Using the distance to the reference point to replace the crowding distance mechanism will not influence the proof procedure, thus we omit the detailed proof here. However, there is still a major difference in the requirement on the problem size. In the proof of the two theorems, the key point

is to keep decreasing the closet distance to the reference point until the reference point is found. The population update mechanism of NSGA-II aims at maintaining solutions with good diversity and does not particularly prefer solutions closer to the reference point, thus the objective vector closest to the reference point can be maintained in the population only if the population size is sufficiently large, i.e.,  $N \geq 4(n + 1)$ . However, R-NSGA-II adopts preference information, i.e., the distance to the reference point, to sort the solutions after non-dominated sorting, thus the objective vector closest to the reference point can always be maintained regardless of the population size  $N$ .

**Theorem 2.** *For R-NSGA-II with any population size  $N$  solving the OneMinMax problem, the expected number of fitness evaluations for finding the reference point  $(k, n - k), k \in [0..n]$ , is  $O(Nn \log n)$ .*

### Analysis on OneJumpZeroJump

The OneJumpZeroJump problem as presented in Definition 4 is a bi-objective variant of the single-objective problem  $\text{JUMP}_{n,k}$  (Droste, Jansen, and Wegener 2002). Its first objective is the same as  $\text{JUMP}_{n,k}$ , which aims at maximizing the number of 1-bits except a valley around  $1^n$ . This valley necessitates the flipping of precisely  $k$  correct bits to traverse it from the local optimum. The second objective shares isomorphism with the first one, wherein the roles of 0s and 1s are exchanged. The Pareto set is

$$S^* = \{\mathbf{x} \in \{0, 1\}^n \mid |\mathbf{x}|_1 \in [k..n - k] \cup \{0, n\}\},$$

and the Pareto front is

$$F^* = \{(i, n + 2k - i) \mid i \in [2k..n] \cup \{k, n + k\}\},$$

whose size is  $|F^*| = n - 2k + 3$ . We let the reference point  $\mathbf{f}(\mathbf{z}) = (n + k, k)$  to see if combining interactive decision making into NSGA-II can speed up the process of traversing the valley and finding the extreme objective vector in the Pareto front. Note that the solution corresponding to the reference point  $(n + k, k)$  is the solution with all 1s.

**Definition 4** (OneJumpZeroJump (Doerr and Zheng 2021)). *The OneJumpZeroJump problem of size  $n$  is to find  $n$ -bits binary strings which maximize*

$$\begin{aligned} f_1(\mathbf{x}) &= \begin{cases} k + |\mathbf{x}|_1, & \text{if } |\mathbf{x}|_1 \leq n - k \text{ or } \mathbf{x} = 1^n, \\ n - |\mathbf{x}|_1, & \text{else;} \end{cases} \\ f_2(\mathbf{x}) &= \begin{cases} k + |\mathbf{x}|_0, & \text{if } |\mathbf{x}|_0 \leq n - k \text{ or } \mathbf{x} = 0^n, \\ n - |\mathbf{x}|_0, & \text{else;} \end{cases} \end{aligned}$$

where  $n \in \mathbb{N}$ ,  $k \in [2..\sqrt{n}]$ , and  $|\mathbf{x}|_1$  and  $|\mathbf{x}|_0$  denote the number of 1-bits and 0-bits of  $\mathbf{x}$ , respectively.

For the convenience of proof, we refer to the notation in (Doerr and Qu 2022), i.e., the inner part of the Pareto set is denoted as

$$S_I^* = \{\mathbf{x} \mid |\mathbf{x}|_1 \in [k..n - k]\},$$

and the outer part of the Pareto set is denoted as

$$S_O^* = \{\mathbf{x} \mid |\mathbf{x}|_1 \in \{0, n\}\};$$

the inner part of the Pareto front is denoted as

$$F_I^* = \mathbf{f}(S_I^*) = \{(i, n + 2k - i) \mid i \in [2k..n]\},$$

and the outer part of the Pareto front is denoted as

$$F_O^* = \mathbf{f}(S_O^*) = \{(i, n + 2k - i) \mid i \in \{k, n + k\}\}.$$

We prove in Theorems 3 and 4 that NSGA-II using a population size  $N \geq 4|F^*| = 4(n - 2k + 3)$  can find the reference point in  $O(Nn^k)$  expected number of fitness evaluations, while R-NSGA-II using any population size  $N \in O(n^{k/2})$  can find the reference point in  $O(n^k)$  expected number of fitness evaluations. Thus, R-NSGA-II can be  $n$  times faster than NSGA-II.

The main proof idea of Theorem 3 is that the algorithm firstly finds the solution with  $(n - k)$  1-bits, and then the reference point can be found by flipping the remaining  $k$  0-bits through bit-wise mutation. Since the proof procedure is similar to that of Theorem 6 in (Doerr and Qu 2022) and the derived upper bound is also the same, we omit the proof of Theorem 3.

**Theorem 3.** *For NSGA-II solving the OneJumpZeroJump problem, if using a population size  $N$  of at least  $4(n + 2k - 3)$ , then the expected number of fitness evaluations for finding the reference point  $(n + k, k)$  is  $O(Nn^k)$ .*

Compared with finding the whole Pareto front, finding a reference point, i.e., a single objective vector in the Pareto front, can be easier. However, our setting  $(n + k, k)$  of the reference point is actually the most difficult one for OneJumpZeroJump, thus the expected running time for finding the reference point is the same as that for finding the whole Pareto front.

The basic proof idea of Theorem 4 is similar to that of Theorem 3. The main difference is that during the optimization procedure, the population of R-NSGA-II will first converge to the objective vector  $(n, 2k)$ , and then jump to the reference point  $(n + k, k)$ .

**Theorem 4.** *For R-NSGA-II with any population size  $N \in O(n^{k/2})$  solving the OneJumpZeroJump problem, the expected number of fitness evaluations for finding the reference point  $(n + k, k)$  is  $O(n^k)$ .*

*Proof.* We divide the optimization procedure into three stages, and we pessimistically assume that the reference point is not found in stages 1 and 2.

- Stage 1: this stage starts after initialization and ends when the solution with  $(n - k)$  1-bits is maintained in the population.
- Stage 2: this stage starts after stage 1 and ends when the population is full of the solutions with  $(n - k)$  1-bits.
- Stage 3: this stage starts after stage 2 and ends when the reference point is found, i.e., the solution  $1^n$  is maintained in the population.

**Stage 1.** We first analyze the expected number of generations until finding a solution in  $S_I^*$ . To this end, we pessimistically assume that any solution in the initial population has at most  $(k - 1)$  1-bits or at least  $(n - k + 1)$  1-bits. Without loss of generality, we can assume that one solution  $\mathbf{x}$  in

the population has at most  $(k - 1)$  1-bits. Then, selecting  $\mathbf{x}$  and flipping  $(k - |\mathbf{x}|_1)$  0-bits can generate a solution in  $S_I^*$ . Note that each solution in the population will be selected for mutation in each generation, thus the probability of generating a solution in  $S_I^*$  is at least

$$\frac{\binom{n-|\mathbf{x}|_1}{k-|\mathbf{x}|_1}}{n^{k-|\mathbf{x}|_1}} \cdot \left(1 - \frac{1}{n}\right)^{n-k+|\mathbf{x}|_1} \geq \frac{1}{e} \cdot \frac{\binom{n-|\mathbf{x}|_1}{k-|\mathbf{x}|_1}}{n^{k-|\mathbf{x}|_1}} \geq \frac{1}{ek^k},$$

where the last inequality is by the proof of Theorem 1 in (Bian et al. 2023) (i.e.,  $\binom{n-|\mathbf{x}|_1}{k-|\mathbf{x}|_1}/n^{k-|\mathbf{x}|_1} \geq \binom{n}{k}/n^k \geq (\frac{n}{k})^k/n^k$ ). Since the solution in  $S_I^*$  dominates any non-Pareto optimal solution, and has a smaller distance to the reference point compared with  $0^n$ , the newly generated solution must be maintained in the population. Thus, the expected number of generations for finding a solution in  $S_I^*$  is at most  $ek^k$ .

Next, we analyze the expected number of generations until finding a solution with  $(n - k)$  1-bits. By flipping a 0-bit of the solution with the most 1-bits in  $S_I^*$ , a solution in  $S_I^*$  with smaller distance to the reference point can be generated, whose probability is at least  $\frac{1}{n} \cdot (1 - \frac{1}{n})^{n-1} \geq \frac{1}{en}$ . Note that the new solution is the most preferred and thus will be maintained in the population. Repeating the above process at most  $(n - 2k)$  times, the solution with  $(n - k)$  1-bits can be maintained in the population. Thus, the total expected number of generations of stage 1 is at most  $ek^k + (n - 2k) \cdot en$ .

**Stage 2.** Note that any non-Pareto optimal solution must be dominated by the solution with  $(n - k)$  1-bits, and any Pareto optimal solution (except  $1^n$ ) has smaller  $f_1$  value and larger  $f_2$  value than the solution with  $(n - k)$  1-bits, thus the solution with  $(n - k)$  1-bits is the most preferred. After a solution with  $(n - k)$  1-bits is found in stage 1, it can consecutively copy itself by flipping none of the bits, and occupy the whole population. The probability of the copying event is at least  $(1 - \frac{1}{n})^n \geq (1 - \frac{1}{n}) \cdot \frac{1}{e} \geq \frac{1}{6}$ . As the duplication procedure needs to repeat at most  $N$  times, the expected number of generations of stage 2 is at most  $6N$ .

**Stage 3.** For a solution  $\mathbf{x}$  with  $|\mathbf{x}|_1 = n - k$ , flipping the  $k$  0-bits can generate the reference point, whose probability is  $(\frac{1}{n})^k \cdot (1 - \frac{1}{n})^{n-k} \geq \frac{1}{en^k}$ . Since all the  $N$  solutions in the population have  $(n - k)$  1-bits, the probability of generating the reference point in each generation is at least  $1 - (1 - \frac{1}{en^k})^N$ . Thus, the expected number of generations is at most

$$\frac{1}{1 - (1 - 1/en^k)^N} \leq \frac{1}{1 - 1 + N \frac{1}{en^k} - \frac{N(N-1)}{2} (\frac{1}{en^k})^2} = O(n^k/N),$$

where the equality holds by  $N = o(n^k)$ .

Combining the three stages, the expected number of generations for finding the reference point is  $O(k^k + n^2 + N + n^k/N) = O(n^k/N)$ , where the equality is by  $N = O(n^{k/2})$ . Note that NSGA-II generates  $N$  solutions in each generation, thus the expected number of fitness evaluations is  $O(n^k)$ .  $\square$

From the proofs of Theorems 3 and 4, we can find the main reason for the acceleration brought by R-NSGA-II.

NSGA-II requires  $N \geq 4(n - 2k + 3)$  to ensure that every objective vector in the Pareto front obtained by the algorithm will not be lost in the future. However, R-NSGA-II can drive the whole population to converge towards the reference point and thus increase the probability of jumping across the valley.

The above comparison on OneMinMax and OneJumpZeroJump discloses two advantages of R-NSGA-II. First, due to the preference to the reference point, it can automatically maintain the solution closest to the reference point in the objective space, without requirement on the population size  $N$ . Second, it can drive the whole population to converge towards the reference point, bringing acceleration. For example, on OneJumpZeroJump, all solutions in the population will keep approaching the reference point, which can significantly increase the probability of leaping to the optima.

### Can R-NSGA-II Always Work?

In the previous section, we have shown the advantage of R-NSGA-II theoretically, which echoes what people said, i.e., combining decision making into MOEAs can accelerate finding the DM's favorite solution. However, we will show that R-NSGA-II may fail sometimes. For this purpose, we consider a variant of OneMinMax, called OneMinMax\* as presented in Definition 5. OneMinMax\* is the same as OneMinMax, except that the objective vector of the solution  $0^n$  is set to  $(-n, 2n)$ . The Pareto front of OneMinMax\* is

$$F^* = \{(k, n - k) \mid k \in [0..n - 1]\} \cup \{(-n, 2n)\},$$

and any solution is still Pareto optimal. We set the reference point to be  $f(z) = (-n, 2n)$ , i.e.,  $z = 0^n$ .

**Definition 5 (OneMinMax\*).** *The OneMinMax\* problem of size  $n$  is to find  $n$ -bits binary strings which maximize*

$$f(x) = \begin{cases} (\sum_{i=1}^n (1 - x_i), \sum_{i=1}^n x_i), & \text{if } x \in \{0, 1\}^n \setminus \{0^n\}, \\ (-n, 2n), & \text{if } x = 0^n. \end{cases}$$

We prove in Theorems 5 and 6 that NSGA-II with  $N \geq 4(n + 1)$  can find the reference point in  $O(Nn \log n)$  expected number of fitness evaluations, while R-NSGA-II using any polynomial population size requires  $\Omega(n^{n/4})$  expected number of fitness evaluations. Thus, R-NSGA-II can be exponentially slower than NSGA-II. We omit the proof of Theorem 5 because its procedure is the same as that of Theorem 1, i.e., NSGA-II can find any reference point by continuously expanding in the Pareto front.

**Theorem 5.** *For NSGA-II solving the OneMinMax\* problem, if using a population size  $N$  of at least  $4(n + 1)$ , then the expected number of fitness evaluations for finding the reference point  $(-n, 2n)$  is  $O(Nn \log n)$ .*

R-NSGA-II drives the whole population to converge towards the reference point in the objective space. However, on the OneMinMax\* problem, a solution closer to the reference point  $(-n, 2n)$  in the objective space will be farther from the corresponding solution  $0^n$  in the decision space. This leads to the inefficiency of R-NSGA-II, which is also the main proof idea of Theorem 6.

**Theorem 6.** *For R-NSGA-II with any polynomial population size  $N$  solving the OneMinMax\* problem, the expected number of fitness evaluations for finding the reference point  $(-n, 2n)$  is  $\Omega(n^{n/4})$ .*

*Proof.* By Chernoff bound, an initial solution randomly selected from  $\{0, 1\}^n$  has at least  $n/4$  1-bits with probability  $1 - e^{-\Omega(n)}$ . Thus, the probability that any individual in the initial population has at least  $n/4$  1-bits is at least

$$(1 - e^{-\Omega(n)})^N \geq 1 - Ne^{-\Omega(n)} = 1 - o(1),$$

where the equality holds because  $N$  is a polynomial of  $n$ .

Now, we assume that the above event happens, i.e., any solution in the initial population has at least  $n/4$  1-bits. Note that any solution in the solution space is Pareto optimal, thus it will belong to  $F_1$  after non-dominated sorting. Meanwhile, since the reference point is  $(-n, 2n)$ , the solution with smaller  $f_1$  and larger  $f_2$  is preferred. Thus, during the optimization procedure, the solution with more 1-bits must be preferred, implying that the population can maintain only solutions with at least  $n/4$  1-bits. To find the reference point, i.e., the solution  $0^n$ , the 1-bits of a parent solution must be flipped simultaneously. As each solution has at least  $n/4$  1-bits, the probability is at most  $1/n^{n/4}$ . Thus, in each generation, the probability of finding  $0^n$  is at most  $1 - (1 - 1/n^{n/4})^N \leq N/n^{n/4}$ , where the inequality is by Bernoulli's inequality.

Combining the above analyses, the expected number of fitness evaluations for finding the reference point is at least

$$(1 - o(1)) \cdot (n^{n/4}/N) \cdot N = \Omega(n^{n/4}),$$

where the factor  $N$  corresponds to  $N$  fitness evaluations in each generation.  $\square$

## Experiments

In the previous sections, we have proved that combining interactive decision making into NSGA-II can bring acceleration for solving the OneMinMax and OneJumpZeroJump problems, while leading to significant deceleration for solving OneMinMax\*. In this section, we first conduct experiments to verify these theoretical results, and then show that the finding from the analysis of OneMinMax\* also holds on the combinatorial problem multi-objective NK-landscape (Aguirre and Tanaka 2007). The data and code are provided in the supplementary material.

For each problem OneMinMax, OneJumpZeroJump or OneMinMax\*, we compare the number of fitness evaluations of NSGA-II and R-NSGA-II until finding the reference point, which is set to  $(0, n)$ ,  $(n + k, k)$  and  $(-n, 2n)$ , respectively, consistent with our theoretical analysis. The first two reference points correspond to the solution with all 1s, and the last one corresponds to the solution with all 0s. We set the problem size  $n$  from 10 to 50, with a step of 10. For each problem size  $n$  and each algorithm, we conduct 1000 independent runs, and record the average and standard deviation of the number of fitness evaluations required to find the reference point.

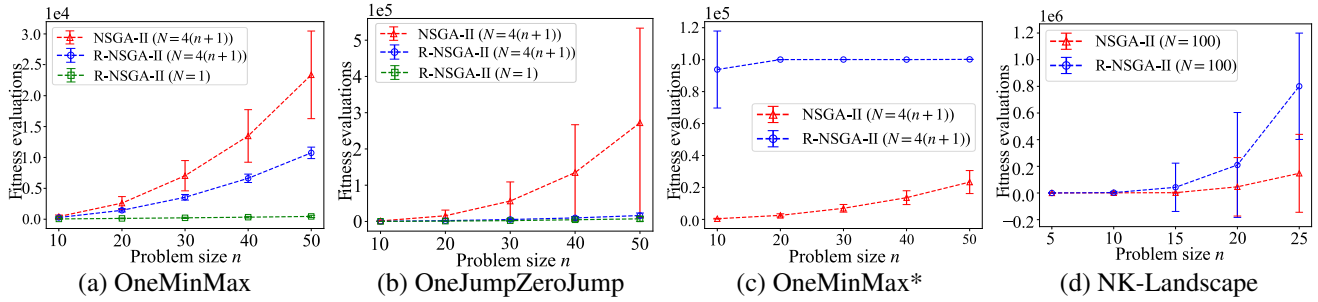


Figure 1: Average number of fitness evaluations of 1000 independent runs of NSGA-II and R-NSGA-II for solving the OneMinMax, OneJumpZeroJump, OneMinMax\* and bi-objective NK-landscape problems, respectively.

For the OneMinMax problem, we set the population size  $N$  of NSGA-II and R-NSGA-II to  $4(n+1)$  and 1, respectively, as suggested by Theorems 1 and 2. Figure 1(a) shows that R-NSGA-II with  $N=1$  can achieve a clear acceleration compared with NSGA-II. We also test the performance of R-NSGA-II with the population size  $N=4(n+1)$  to see whether the acceleration is only brought by the decrease of the population size. We can observe that even using the same population size, R-NSGA-II is still faster.

For the OneJumpZeroJump problem, the parameter  $k$  is set to 2. The population size  $N$  of NSGA-II and R-NSGA-II is set to  $4(n-2k+3)$  and 1, respectively, according to Theorems 3 and 4. We can see from Figure 1(b) that R-NSGA-II with  $N=1$  needs much less fitness evaluations than NSGA-II with  $4(n-2k+3)$ , to find the reference point. Moreover, by increasing the population size  $N$  to  $4(n-2k+3)$ , the average time of R-NSGA-II only increases slightly, which is consistent with Theorem 4 that the expected running time of R-NSGA-II does not depend on  $N$ .

For the OneMinMax\* problem, we set the population size  $N$  of NSGA-II to  $4(n+1)$  as suggested by Theorem 5, and use the same  $N$  for R-NSGA-II for fairness. As we find that it's very time-consuming for R-NSGA-II to find the reference point, we set a maximum number of fitness evaluation, i.e.,  $10^5$ . Figure 1(c) shows that NSGA-II is clearly better, and R-NSGA-II almost reaches the maximum time limit, which is consistent with Theorem 6 that R-NSGA-II requires at least exponential time to find the reference front.

From the proof of Theorem 6, we find that for the OneMinMax\* problem, a solution closer to the reference point in the objective space is farther from it in the decision space, leading to the inefficiency of R-NSGA-II. Such characteristic may hold in complex real-world problems, where R-NSGA-II can be inefficient. To validate this finding, we compare R-NSGA-II with NSGA-II on the bi-objective NK-Landscape problem, where the problem size  $n$  is set from 5 to 25, with a step of 5. For each  $n$ , we choose a set of randomly initialized parameters, i.e., fitness contribution  $c_{ij}$  and loci  $k_{ij}$ , to generate the problem, following (Li, Han, and Chu 2023). The parameter  $K$  is set to 3, and the population size  $N$  of NSGA-II and R-NSGA-II is both set to 100 for fairness. We acquire the Pareto front by exhaustive enumeration over the whole solution space  $\{0, 1\}^n$ , and then randomly choose one from the Pareto front as the reference

point, which is  $[0.821, 0.483]$ ,  $[0.643, 0.647]$ ,  $[0.770, 0.481]$ ,  $[0.741, 0.521]$  and  $[0.760, 0.541]$  for the problem size  $n \in \{5, 10, \dots, 25\}$ , respectively. For R-NSGA-II, we set a maximum number of fitness evaluations, i.e.,  $10^6$ . Figure 1(d) shows that R-NSGA-II is much worse than NSGA-II.

## Conclusion and Discussion

In this paper, we make a first step towards theoretically analyzing the expected running time of practical iMOEA. We consider the interactive variant, R-NSGA-II, of the popular NSGA-II. We prove that R-NSGA-II can be asymptotically faster than NSGA-II for solving the OneMinMax and OneJumpZeroJump problems, providing theoretical justification for iMOEA. Furthermore, we find that the integration of interactive decision-making within MOEAs does not always expedite convergence by proving that R-NSGA-II can be exponentially slower than NSGA-II for solving a variant of the OneMinMax problem. Though being proved in a special case, this finding may hold more generally, especially on complex real-world problems with rugged fitness landscape. We have empirically verified it on the popular practical problem NK-landscape. This finding suggests the importance of striking a balance between interactive decision-making and diversity preservation within the population, which might be helpful for designing efficient iMOEAs in practice.

Building upon the insights gained from our theoretical and empirical results, there are several avenues for future research. In this work, we mainly focus on scenarios where R-NSGA-II exhibits significant advantages or disadvantages compared to NSGA-II, while it would be interesting to identify the situations where R-NSGA-II and NSGA-II show similar performance. We believe such benchmark can be created by decreasing the gap in the decision space between the reference point and the point closest to it in the objective space, e.g., set the reference point with objective vector  $(-n, 2n)$  in OneMinMax\* to the solution closer to  $1^n$  in the decision space. That is, if setting the objective vector of a solution with  $k$  1-bits to  $(-n, 2n)$ , then R-NSGA-II will perform better with larger  $k$ . Moreover, considering NSGA-II's inefficiency in many-objective optimization (Zheng and Dorr 2023), analysis of R-NSGA-II's performance in many-objective optimization would be a meaningful direction.



## Acknowledgements

The authors want to thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the National Key R&D Program of China (2022ZD0116600) and National Science Foundation of China (62022039, 62276124). Chao Qian is the corresponding author.

## References

- Aguirre, H. E.; and Tanaka, K. 2007. Working principles, behavior, and performance of MOEAs on MNK-landscapes. *European Journal of Operational Research*, 181(3): 1670–1690.
- Bandte, O.; and Malinchik, S. 2004. A broad and narrow approach to interactive evolutionary design—an aircraft design example. In *Proceedings of the 6th GECCO*, 883–895. Seattle, WA.
- Beume, N.; Naujoks, B.; and Emmerich, M. 2007. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3): 1653–1669.
- Bian, C.; and Qian, C. 2022. Better running time of the non-dominated sorting genetic algorithm II (NSGA-II) by using stochastic tournament selection. In *Proceedings of the 17th International Conference on Parallel Problem Solving from Nature (PPSN)*, 428–441. Dortmund, Germany.
- Bian, C.; Zhou, Y.; Li, M.; and Qian, C. 2023. Stochastic population update can provably be helpful in multi-objective evolutionary algorithms. In *Proceedings of the 32nd IJCAI*, 5513–5521. Macao, SAR, China.
- Brockhoff, D.; López-Ibáñez, M.; Naujoks, B.; and Rudolph, G. 2012. Runtime analysis of simple interactive evolutionary biobjective optimization algorithms. In *Proceedings of the 12th International Conference on Parallel Problem Solving from Nature (PPSN)*, 123–132. Taormina, Italy.
- Bäck, T. 1996. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford, UK: Oxford University Press.
- Cerf, S.; Doerr, B.; Hebras, B.; Kahane, Y.; and Wietheger, S. 2023. The first proven performance guarantees for the non-dominated sorting genetic algorithm II (NSGA-II) on a combinatorial optimization problem. In *Proceedings of the 32nd IJCAI*, 5522–5530. Macao, SAR, China.
- Dang, D.; Opris, A.; Salehi, B.; and Sudholt, D. 2023a. A proof that using crossover can guarantee exponential speed-ups in evolutionary multi-objective optimisation. In *Proceedings of the 37th AAAI*, 12390–12398. Washington, DC.
- Dang, D.-C.; Opris, A.; Salehi, B.; and Sudholt, D. 2023b. Analysing the robustness of NSGA-II under noise. In *Proceedings of the 25th GECCO*, 642–651. Lisbon, Portugal.
- Deb, K.; and Kumar, A. 2007. Light beam search based multi-objective optimization using evolutionary algorithms. In *Proceedings of the 2007 IEEE Congress on Evolutionary Computation (CEC)*, 2125–2132. Singapore.
- Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2): 182–197.
- Deb, K.; and Sundar, J. 2006. Reference point based multi-objective optimization using evolutionary algorithms. In *Proceedings of the 8th GECCO*, 635–642. Seattle, WA.
- Dinot, M.; Doerr, B.; Hennebelle, U.; and Will, S. 2023. Runtime analyses of multi-objective evolutionary algorithms in the presence of noise. In *Proceedings of the 32nd IJCAI*, 5549–5557. Macao, SAR, China.
- Doerr, B.; Kodric, B.; and Voigt, M. 2013. Lower bounds for the runtime of a global multi-objective evolutionary algorithm. In *Proceedings of the 2013 IEEE Congress on Evolutionary Computation (CEC)*, 432–439. Cancun, Mexico.
- Doerr, B.; and Qu, Z. 2022. A first runtime analysis of the NSGA-II on a multimodal problem. In *Proceedings of the 17th International Conference on Parallel Problem Solving from Nature (PPSN)*, 399–412. Dortmund, Germany.
- Doerr, B.; and Qu, Z. 2023a. From understanding the population dynamics of the NSGA-II to the first proven lower bounds. In *Proceedings of the 37th AAAI*, 12408–12416. Washington, DC.
- Doerr, B.; and Qu, Z. 2023b. Runtime analysis for the NSGA-II: Provable speed-ups from crossover. In *Proceedings of the 37th AAAI*, 12399–12407. Washington, DC.
- Doerr, B.; and Wietheger, S. 2023. A mathematical runtime analysis of the non-dominated sorting genetic algorithm III (NSGA-III). In *Proceedings of the 32nd IJCAI*, 5657–5665. Macao, SAR, China.
- Doerr, B.; and Zheng, W. 2021. Theoretical analyses of multi-objective evolutionary algorithms on multi-modal objectives. In *Proceedings of the 35th AAAI*, 12293–12301. Virtual.
- Droste, S.; Jansen, T.; and Wegener, I. 2002. On the analysis of the  $(1+1)$  evolutionary algorithm. *Theoretical Computer Science*, 276(1-2): 51–81.
- Friedrich, T.; and Neumann, F. 2015. Maximizing submodular functions under matroid constraints by evolutionary algorithms. *Evolutionary Computation*, 23(4): 543–558.
- Giel, O. 2003. Expected runtimes of a simple multi-objective evolutionary algorithm. In *Proceedings of the 2003 IEEE Congress on Evolutionary Computation (CEC)*, 1918–1925. Canberra, Australia.
- Giel, O.; and Lehre, P. K. 2010. On the effect of populations in evolutionary multi-objective optimisation. *Evolutionary Computation*, 18(3): 335–356.
- Gong, M.; Liu, F.; Zhang, W.; Jiao, L.; and Zhang, Q. 2011. Interactive MOEA/D for multi-objective decision making. In *Proceedings of the 13th GECCO*, 721–728. Dublin, Ireland.
- Horoba, C. 2009. Analysis of a simple evolutionary algorithm for the multiobjective shortest path problem. In *Proceedings of the 10th International Workshop on Foundations of Genetic Algorithms (FOGA)*, 113–120. Orlando, FL.



- Huang, Z.; and Zhou, Y. 2020. Runtime analysis of somatic contiguous hypermutation operators in MOEA/D framework. In *Proceedings of the 34th AAAI*, 2359–2366. New York, NY.
- Huang, Z.; Zhou, Y.; Chen, Z.; and He, X. 2019. Running time analysis of MOEA/D with crossover on discrete optimization problem. In *Proceedings of the 33rd AAAI*, 2296–2303. Honolulu, HI.
- Huang, Z.; Zhou, Y.; Luo, C.; and Lin, Q. 2021. A runtime analysis of typical decomposition approaches in MOEA/D framework for many-objective optimization problems. In *Proceedings of the 30th IJCAI*, 1682–1688. Virtual.
- Huber, S.; Geiger, M. J.; and Sevaux, M. 2015. Simulation of preference information in an interactive reference point-based method for the bi-objective inventory routing problem. *Journal of Multi-Criteria Decision Analysis*, 22(1-2): 17–35.
- Janssen, R.; Van Herwijnen, M.; Stewart, T. J.; and Aerts, J. C. 2008. Multiobjective decision support for land-use planning. *Environment and planning B: Planning and design*, 35(4): 740–756.
- Köksalan, M.; and Karahan, I. 2010. An interactive territory defining evolutionary algorithm: iTDEA. *IEEE Transactions on Evolutionary Computation*, 14(5): 702–722.
- Laumanns, M.; Thiele, L.; and Zitzler, E. 2004a. Running time analysis of evolutionary algorithms on a simplified multiobjective knapsack problem. *Natural Computing*, 3: 37–51.
- Laumanns, M.; Thiele, L.; and Zitzler, E. 2004b. Running time analysis of multiobjective evolutionary algorithms on pseudo-Boolean functions. *IEEE Transactions on Evolutionary Computation*, 8(2): 170–182.
- Li, K. 2019. Progressive preference learning: Proof-of-principle results in MOEA/D. In *Proceedings of the 10th International Conference on Evolutionary Multi-Criterion Optimization (EMO)*, 631–643. East Lansing, MI.
- Li, M.; Han, X.; and Chu, X. 2023. MOEAs are stuck in a different area at a time. In *Proceedings of the 25th GECCO*, 303–311. Lisbon, Portugal.
- Narukawa, K.; Tanigaki, Y.; and Ishibuchi, H. 2014. Evolutionary many-objective optimization using preference on hyperplane. In *Proceedings of the 16th GECCO*, 91–92. Vancouver, Canada.
- Neumann, F. 2007. Expected runtimes of a simple evolutionary algorithm for the multi-objective minimum spanning tree problem. *European Journal of Operational Research*, 181(3): 1620–1629.
- Neumann, F.; and Nguyen, A. Q. 2014. On the impact of utility functions in interactive evolutionary multi-objective optimization. In *Proceedings of the 10th International Conference on Simulated Evolution and Learning (SEAL)*, 419–430. Dunedin, New Zealand.
- Neumann, F.; and Theile, M. 2010. How crossover speeds up evolutionary algorithms for the multi-criteria all-pairs-shortest-path problem. In *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature (PPSN)*, 667–676. Krakow, Poland.
- Qian, C. 2020. Distributed pareto optimization for large-scale noisy subset selection. *IEEE Transactions on Evolutionary Computation*, 24(4): 694–707.
- Qian, C.; Shi, J.-C.; Yu, Y.; Tang, K.; and Zhou, Z.-H. 2017. Subset selection under noise. In *Advances in Neural Information Processing Systems 30 (NIPS)*, 3562–3572. Long Beach, CA.
- Qian, C.; Tang, K.; and Zhou, Z.-H. 2016. Selection hyperheuristics can provably be helpful in evolutionary multi-objective optimization. In *Proceedings of the 14th International Conference on Parallel Problem Solving from Nature (PPSN)*, 835–846. Edinburgh, Scotland.
- Qian, C.; Yu, Y.; Tang, K.; Yao, X.; and Zhou, Z.-H. 2019. Maximizing submodular or monotone approximately submodular functions by multi-objective evolutionary algorithms. *Artificial Intelligence*, 275: 279–294.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2013. An analysis on recombination in multi-objective evolutionary optimization. *Artificial Intelligence*, 204: 99–119.
- Qian, C.; Yu, Y.; and Zhou, Z.-H. 2015. Subset selection by pareto optimization. In *Advances in Neural Information Processing Systems 28 (NIPS)*, 1765–1773. Montreal, Canada.
- Ramírez, A.; Romero, J. R.; and Ventura, S. 2018. Interactive multi-objective evolutionary optimization of software architectures. *Information Sciences*, 463–464: 92–109.
- Roostapour, V.; Neumann, A.; Neumann, F.; and Friedrich, T. 2022. Pareto optimization for subset selection with dynamic cost constraints. *Artificial Intelligence*, 302: 103597.
- Steuer, R. E. 1986. *Multiple Criteria Optimization: Theory, Computations, and Application*. New York, NY: John Wiley & Sons.
- Takagi, H. 2001. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9): 1275–1296.
- Thiele, L.; Miettinen, K.; Korhonen, P. J.; and Molina, J. 2009. A preference-based evolutionary algorithm for multi-objective optimization. *Evolutionary Computation*, 17(3): 411–436.
- Zhang, Q.; and Li, H. 2007. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6): 712–731.
- Zheng, W.; and Doerr, B. 2022. Better approximation guarantees for the NSGA-II by using the current crowding distance. In *Proceedings of the 24th GECCO*, 642–651. Boston, MA.
- Zheng, W.; and Doerr, B. 2023. Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives. *IEEE Transactions on Evolutionary Computation*, in press.
- Zheng, W.; Liu, Y.; and Doerr, B. 2022. A first mathematical runtime analysis of the non-dominated sorting genetic algorithm II (NSGA-II). In *Proceedings of the 36th AAAI*, 10408–10416. Virtual.
- Zhou, Z.-H.; Yu, Y.; and Qian, C. 2019. *Evolutionary learning: Advances in Theories and Algorithms*. Springer.