Finding ϵ-locally Optimal Solutions for Multi-objective Multimodal Optimization

Angel E. Rodriguez-Fernandez*, Lennart Schäpermeier*, Carlos Hernández*, Pascal Kerschke, *Member, IEEE*, Heike Trautmann, *Member, IEEE*, and Oliver Schütze

***Abstract*—In this paper, we address the problem of computing all locally optimal solutions of a given multi-objective problem whose images are sufficiently close to the Pareto front. Such *ϵ-locally optimal solutions* are particularly interesting in the context of multi-objective multimodal optimization (MMO). To accomplish this task, we first define a new set of interest, $L_{Q,\epsilon}$, that is strongly related to the recently proposed set of ϵ-acceptable solutions. Next, we propose a new unbounded archiver, $ArchiveUpdateL_{Q,\epsilon}$, aiming to capture $L_{Q,\epsilon}$ in the limit. This archiver can in principle be used in combination with any multi-objective evolutionary algorithm (MOEA). Further, we equip numerous MOEAs with $ArchiveUpdateL_{Q,\epsilon}$, investigate their performances across several benchmark functions, and compare the enhanced MOEAs with their archive-free counterparts. For our experiments, we utilize the well-established metrics HV, IGDX, and Δ_p. Additionally, we propose and use a new performance indicator, I_{EDR}, which results in comparable performances but which is applicable to problems defined in higher dimensions (in particular in decision variable space).**

***Index Terms*—multi-objective optimization, evolutionary computation, multimodal optimization, local solutions.**

I. Introduction

MULTI-OBJECTIVE optimization problems (MOPs), i.e., problems where several objectives have to be optimized concurrently, arise in many applications. One important characteristic of such problems is that the solution set – the so-called Pareto set – forms in the continuous case and under certain (mild) assumptions on the model, at least locally, a manifold of dimension $k-1$, where k is the number of objectives considered in the MOP.

For the numerical treatment of MOPs, multi-objective evolutionary algorithms (MOEAs) have caught the interest of many researchers and practitioners during the last three decades. Reasons for this include that MOEAs are applicable to a wide range of problems, require minimal assumptions on the model, and are of a global nature. Further, their set-based approach allows them to compute a finite-size representation of the entire solution set in one algorithm run. Hereby, however, most MOEAs focus on the approximation of the Pareto front, which is the image of the Pareto set.

On the one hand, this makes sense since the Pareto front represents the set of optimal options for the decision maker (DM) measured by the individual objectives of the given MOP. On the other hand, the possible realizations of these options are, in certain cases, not fully considered: It may happen that for an element y of the Pareto front, there exist two (or more) solutions x_1 and x_2 that map both to y (i.e., $F(x_1) = F(x_2) = y$, where F denotes the multi-objective map). For such problems, most MOEAs only store one of these solutions. This could represent a shortcoming since x_1 and x_2 could be very distinct and represent essentially different realizations, which is useful information for the DM.

The situation is essentially the same if $F(x_1) \approx F(x_2)$, even if, say, x_2 is dominated by x_1 (but by small amounts in terms of the objective function values). Analogously, *all* of such ϵ-efficient or nearly optimal solutions are of potential interest for the DM which has been studied in [1–5]. One major issue with these sets is that they are of dimension n (n being the number of decision variables) while the dimension of the Pareto set/front is $k-1$, and n is in most applications much larger than $k-1$. Recently, Li at al. [6] proposed the set of ϵ-acceptable solutions in the context of multimodal optimization (MMO) which can be seen as a particular discretization of the set of nearly optimal solutions. More precisely, the authors consider a feasible point ϵ-acceptable if it is locally optimal and ϵ-efficient. One can hence expect that these points form a set of dimension $k-1$ (i.e., the same as the Pareto set/front).

In this work, we address the computation of such locally optimal ϵ-efficient solutions. To this end, we will first define a set of interest, $L_{Q,\epsilon}$, that is highly related to the set of ϵ-acceptable solutions (but which uses a slightly different way to define ϵ-efficiency). In the next step, we propose a new unbounded archiver that aims to approximate this set. Further, we will discuss how to use this archiver within existing MOEAs and propose a new performance indicator. This indicator, I_{EDR}, is considering the "essentially different realizations" for points y on the Pareto front and, hence, is highly related to $L_{Q,\epsilon}$. Opposed to existing indicators in MMO, this one is applicable to problems defined in higher dimensions (in particular in decision variable space).

The remainder of this work is organized as follows: In Section II, we briefly present the required background to understand this work and discuss the related work. In Section III,

*Angel E. Rodriguez-Fernandez, Lennart Schäpermeier, and Carlos Hernández contributed equally to this work.

Angel E. Rodriguez-Fernandez and Oliver Schütze are with the CINVESTAV-IPN, Mexico. Lennart Schäpermeier and Pascal Kerschke are with the TU Dresden, Germany. Carlos Hernández is with IIMAS-UNAM, Mexico. Heike Trautmann is with the University of Paderborn, Germany, and the University of Twente, Netherlands.

This work was supported by the *European Research Center for Information Systems (ERCIS)*. The work of Lennart Schäpermeier and Pascal Kerschke was supported by the *Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig*. Angel E. Rodriguez-Fernandez acknowledges support from the CONAHCYT to pursue his postdoc fellowship at the CINVESTAV-IPN. Carlos Hernández also acknowledges the funding for this research provided by UNAM's project PAPIIT no. IA102923. Oliver Schütze acknowledges funding from CONAHCYT project CBF2023-2024-1463.

we define and discuss the set of interest, $L_{Q,\epsilon}$. In Section IV, we present the archiver $ArchiveUpdate L_{Q,\epsilon}$ that aims to capture all promising solutions. In Section V, we shortly discuss how this archiver can be integrated into MOEAs, and propose in Section VI a new performance indicator, I_{EDR}, related to $L_{Q,\epsilon}$. In Section VII, we present some numerical results, and finally, we draw our conclusions and discuss possible paths of future work in Section VIII.

II. Background and Related Work

Throughout this work, we consider continuous multi-objective optimization problems that can be expressed as

$$\min_{x \in Q} F(x). \tag{MOP}$$

Hereby, F denotes the map consisting of k individual objectives f_i, $i = 1, \ldots, k$,

$$F : Q \to \mathbb{R}^k, \qquad F(x) = (f_1(x), \ldots, f_k(x))^T, \tag{1}$$

which is defined on a domain $Q \subset \mathbb{R}^n$. Given $v, w \in \mathbb{R}^k$, we say that v is less than w ($v <_p w$), if $v_i < w_i$ for all $i \in \{1, \ldots, k\}$ (analog for the relation $\leq_p$). $y \in Q$ is dominated by a point $x \in Q$ ($x \prec y$) with respect to (MOP) if $F(x) \leq_p F(y)$ and $F(x) \neq F(y)$. $x \in Q$ is called a Pareto optimal point if there exists no $y \in Q$ that dominates x. The set of all Pareto optimal points is called the Pareto set (PS), denoted by P_Q, and its image $F(P_Q)$ is called the Pareto front (PF).

Given a vector $\epsilon = (\epsilon_1, \ldots, \epsilon_k)^T \in \mathbb{R}^k_+$ and $x, y \in Q$, we say that x ϵ-dominates y ($x \prec_\epsilon y$) with respect to (MOP) if $F(x) - \epsilon \leq_p F(y)$ and $F(x) - \epsilon \neq F(y)$.

For the performance evaluations we will use several performance indicators. One of them is the averaged Hausdorff distance ([7, 8]).

Definition 1: Let $A, B \subset \mathbb{R}^m$ be finite sets. The value

$$\Delta_p(A, B) = \max(D_p(A, B), ID_p(A, B)), \tag{2}$$

where

$$\begin{aligned} D_p(A, B) &= \left(\frac{1}{|A|} \sum_{a \in A} d(a, B)^p \right)^{1/p} \\ ID_p(A, B) &= \left(\frac{1}{|B|} \sum_{b \in B} d(b, A)^p \right)^{1/p}, \end{aligned} \tag{3}$$

$p \in \mathbb{N}$ and $d(a, B) = \min_{b \in B}(\|a - b\|_2)$, is called the *averaged Hausdorff distance* between A and B.

For $p = \infty$, the indicator coincides with the Hausdorff distance d_H. As opposed to d_H, the values of Δ_p are for finite values of p not entirely determined by single outliers in the candidate set. Δ_p is well understood in the context of Pareto front approximations (i.e., $m = k$ and B an approximation of the Pareto front). In that case, Δ_p prefers evenly spread solutions of the images of the candidate solutions along the Pareto front ([9]). Based on the the definition of Δ_p and our prior observations we expect Δ_p to behave very similarily for the sets we consider in this work: we will measure the approximation quality of a given archive A toward $L_{Q,\epsilon}$ (decision variable space), respectively $F(A)$ toward $F(L_{Q,\epsilon})$ (objective space). Since the latter does not have to be identical to the Pareto front of the given MOP, we are hence not dealing with (classical) Pareto front approximations. We will further use the indicators IGDX ([10], in decision variable space) and HV ([11], in objective space).

Over the last three decades, many different multi-objective evolutionary algorithms (MOEAs) have been proposed [12], which can be roughly divided into three main classes: (i) MOEAs based on the dominance relation [13, 14], (ii) on decompositions [15, 16] and (iii) on indicator functions [17, 18]. Since the sets of interest – mainly the Pareto set/front, but also other ones such as all locally optimal or nearly optimal solutions – typically consist of infinitely many elements, all of these MOEAs are equipped with particular selection strategies (i.e., strategies that decide which of the promising candidate solutions to keep during the run of the algorithm, and which ones to discard). Next to these selection strategies, several archiving strategies have been proposed during the last years, mainly intended to be used as external archivers to the base MOEA in order to obtain a more complete representation of the given set of interest. There exist, for instance, several bounded archivers, i.e., selection strategies that maintain archives whose magnitudes do not exceed a pre-described threshold. For instance, in [19, 20], several such archiving strategies are presented and discussed that are based on adaptive grid selections. In [21], bounded archivers are presented and analyzed, aiming for hypervolume approximations of the Pareto front. While in [22], the authors analyze the properties of several archiving techniques widely used in the literature.

Next, several archiving strategies have been proposed that utilize the concept of ϵ-dominance [23–26]. In these works, this concept is used to gather finite-size Pareto front approximations. All of these archivers yield finite-size approximations of the Pareto front with certain approximation qualities in the limit. However, the final magnitudes can not be adjusted a priori (since they mainly depend on the size of the Pareto front, which is not known beforehand). In [27], the authors propose an algorithm that converges to an ϵ-Pareto set of size k with smallest possible value of ϵ. In [4], the authors propose an archiver that aims to represent the entire set of approximate solutions of a given MOP.

Finally, some researchers proposed to store all the non-dominated solutions found during the run of the algorithm (e.g., [28–32]). Since such archivers are unbounded by nature and since many DMs prefer to consider moderately-sized representations of the solution sets after the run of the algorithm, specialized subset selection techniques have been proposed for this purpose (e.g., [33, 34]). The interested reader is referred to [35] for a survey on archiving in multi-objective optimization.

The final concept that we need to introduce is multi-objective multimodal optimization (MMO). In single-objective optimization, multimodal optimization strives to identify local optima within a certain maximum distance ϵ (in objective space) to the global optimum, as illustrated in Figure 1. This approach allows to provide the decision maker with potential alternative or backup solutions. A good overview of (single-

objective) multimodal optimization is given in [36].

In contrast, multi-objective multimodal optimization (MMO) only recently started to attract increasing attention [37–39]. Still, there are several competing perspectives on this subject, which we will shortly outline here.

The most prevalent view presumes that MMO problems (MMOPs) feature a Pareto front with multiple distinct pre-images in decision space [38, 40–42]. It is then the task of an optimization algorithm to find and approximate all pre-images, or in other words, the entire Pareto set. To accomplish this task, the (temporary) consideration of locally optimal solution is a helpful tool (e.g., [5]), since all global solutions are also locally optimal. Test problems that follow this characterization are, for example, the suite of MMF problems [41], SYM-PART problems [40], Deterministic Distortion and Rotation Benchmark (DDRB) [43], and Omni-test [44]. However, this specification of MMOPs is rather limited: It presumes that all solutions of interest are Pareto-optimal and is thus only equivalent to the single-objective scenario where $\epsilon = 0$. This is evidenced in studies of bi-objective problems created from the well-known black-box optimization benchmark (BBOB) [45], specifically the bi-objective BBOB [46], as well as test problems created with a multiple peaks model [37, 47, 48], where the Pareto front has only one pre-image. Consequently, this view is also referred to as *multiglobal* multi-objective optimization [39].

A secondary perspective of MMO focuses on identifying locally efficient points [39, 49–51]. Landscape analysis and visualization studies based on this perspective highlight the importance of locally efficient sets and their induced search dynamics in a variety of test problems [52–56]. This perspective emphasizes the ability of locally efficient solutions to filter out irrelevant decision alternatives for which slight modifications already give a dominating performance.

Finally, the set of nearly optimal solutions is also sometimes referred to as MMO [1–5, 38, 57]. However, this definition does away with all notions of local optimality and may deliver solution alternatives that are all contained in the same basin of attraction. Based on these characteristics, we do not consider this to be an essentially MMO approach. However, it will turn out to be an important ingredient for the set of interest that we will consider for the rest of this study.

III. The Set of Interest

Here, we define and shortly discuss the set of interest, $L_{Q,\epsilon}$. As motivated above, we are interested in all locally optimal solutions x, whose images $F(x)$ are "close" to the Pareto front. Thus, given $\epsilon \in \mathbb{R}^k_+$ and Pareto set P_Q, let the set of nearly optimal solutions and the set of locally optimal solutions be

$$\begin{aligned} N_{Q,\epsilon} &:= \{x \in Q | \exists p \in P_Q : x \prec_\epsilon p\}, \\ L_Q &:= \{x \in Q | \exists \text{ neighborhood } N \text{ of } x \\ &\qquad \text{s.t. } \nexists y \in (N \cap Q) \ : \ y \prec x\}. \end{aligned} \tag{4}$$

Then, we define the set $L_{Q,\epsilon}$ as their intersection.

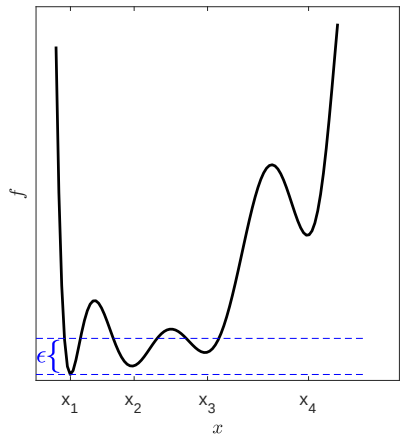

Fig. 1. Example of a scalar optimization problem with one global solution, x_1, and three locally optimal ones, x_2 to x_4.

Definition 2: $L_{Q,\epsilon}$ defines the set of all the nearly optimal solutions of (MOP) that are also locally optimal, i.e.,

$$\begin{aligned} L_{Q,\epsilon} &:= N_{Q,\epsilon} \cap L_Q \\ &= \{x \in Q | \exists p \in P_Q : x \prec_\epsilon p \text{ and } \exists \text{ neighborhood } N \\ &\qquad \text{of } x \text{ s.t. } \nexists y \in (N \cap Q) \ : \ y \prec x\}. \end{aligned} \tag{5}$$

Remark 1:
$L_{Q,\epsilon}$ is strongly related to the set of ϵ-acceptable solutions defined in [6]. The main difference between the two sets is given by the use of different definitions of ϵ-efficiency. While $L_{Q,\epsilon}$ uses the additive form defined above, the other set is based on a multiplicative definition of ϵ-dominance [23]. That is, both sets accomplish the same task, the difference is in the choice of the relaxation parameter ϵ: absolute values for $L_{Q,\epsilon}$ and relative ones for the set of ϵ-acceptable solutions.

Depending on the choice of $\epsilon \in \mathbb{R}^k_+$, the set $L_{Q,\epsilon}$ can, in principle, be used to deal with the two other existing MMO definitions: (i) It is $P_Q \subset L_{Q,\epsilon}$ for all $\epsilon \in \mathbb{R}^k_+$. That is, for "very small" values of ϵ, the set $L_{Q,\epsilon}$ will essentially only consist of all the pre-images of the Pareto front. (ii) For "large enough" values of ϵ, $L_{Q,\epsilon}$ is identical to the set of all locally optimal solutions L_Q, regardless of the distance to the Pareto front. However, as discussed above, we suggest here to use "small" values of ϵ so that $L_{Q,\epsilon}$ contains all locally optimal solutions (in particular all pre-images of the Pareto front) whose images are still close to the Pareto front. Needless to say, "small" and "close" are problem-dependent, so the values of ϵ must be adjusted by the DM to the given problem accordingly.

To familiarize with $L_{Q,\epsilon}$, we consider three examples:

(a) The easiest way to get familiar with $L_{Q,\epsilon}$ is to consider the special case $k = 1$, i.e., a scalar or single-objective optimization problem (SOP). For this case, we obtain

$$\begin{aligned} N^{(1)}_{Q,\epsilon} = \{x \in Q | f(x) - \epsilon < f(x^*) \text{ and } \exists \text{ neighborhood} \\ N \text{ of } x \text{ s.t. } \nexists y \in N \cap Q : \ f(y) < f(x)\}, \end{aligned} \tag{6}$$

where f is the (only) objective function, $\epsilon > 0$, and x^* a global optimum of the SOP. This definition reflects the

"good-subset" approach from [36] and related quality indicators proposed in [58]. Figure 1 shows the graph of a function f with one global minimum and three local minima. For the ϵ-value indicated in the figure, we obtain $L_{Q,\epsilon} = \{x_1, x_2, x_3\}$. In particular, x_4 is not considered as its image is "too far" from $f(x_1)$.

(b) Next, we consider a modification of SYM-PART [59], which we call SYM-PART9to9 in the sequel:

$$F : [-8, 8] \times [-6, 6] \to \mathbb{R}^2$$
$$F(x) = (f_1(x), f_2(x))^T \text{ where :}$$
$$f_i(x) = \begin{cases} r_i(x) & \text{if } t_1 = 0, \ t_2 = 0 \\ r_i(x) + \delta & \text{if } t_1 = -1, \ t_2 = 0 \\ r_i(x) + 2\delta & \text{if } t_1 = 1, \ t_2 = 0 \\ r_i(x) + 3\delta & \text{if } t_1 = 0, \ t_2 = -1 \\ r_i(x) + 4\delta & \text{if } t_1 = -1, \ t_2 = -1 \\ r_i(x) + 5\delta & \text{if } t_1 = 1, \ t_2 = -1 \\ r_i(x) + 6\delta & \text{if } t_1 = 0, \ t_2 = 1 \\ r_i(x) + 7\delta & \text{if } t_1 = -1, \ t_2 = 1 \\ r_i(x) + 8\delta & \text{if } t_1 = 1, \ t_2 = 1 \end{cases}$$
$$r_1(x) = (x_1 - t_1(c + 2a) + a)^2 + (x_2 - t_2 b)^2$$
$$r_2(x) = (x_1 - t_1(c + 2a) - a)^2 + (x_2 - t_2 b)^2$$
$$t_1 = sign(x_1) \min(\lceil (|x_1| - a - c/2)/(2a + c) \rceil, 1),$$
$$t_2 = sign(x_2) \min(\lceil (|x_2| - b/2)/b \rceil, 1),$$
$$a = 0.5, \quad b = 5, \quad c = 5, \quad \delta = 0.12.$$

In SYM-PART, each point on the Pareto front has nine different pre-images. In SYM-PART9to9, in eight out of these nine segments, the function values are increased by multiples of $\delta > 0$ so that the problem has one global Pareto set/front, and another eight local sets/fronts. If $8\delta \leq \epsilon_i$, $i = 1, 2$, $L_{Q,\epsilon}$ is hence given by all of these nine components. Figures 2 (a) and (b) show this scenario. Note that in this example $F(L_{Q,\epsilon})$ is not identical to the Pareto front. Since this is the case for all MOPs that have ϵ-locally optimal solutions that are not globally optimal, we are in this context not dealing with (classical) Pareto front approximations.

(c) At last, we consider the first instance of the 55th bi-objective BBOB problem in 2D [46], which we abbreviate as BBOB55. Figures 2 (c) and (d) show the corresponding sets $L_{Q,\epsilon}$ and $F(L_{Q,\epsilon})$ for $\epsilon = (2.5, 2.5)^T$. Noticeably, $L_{Q,\epsilon}$ falls into multiple connected components all over the search space.

IV. An Archiver for the Computation of $L_{Q,\epsilon}$

Here we propose $ArchiveUpdateL_{Q,\epsilon}$, an unbounded archiver that aims to capture $L_{Q,\epsilon}$ in the limit. Since the archiver tends to generate outliers, particularly for more complex problems, we also propose a cleaning step to be applied to the final archive.

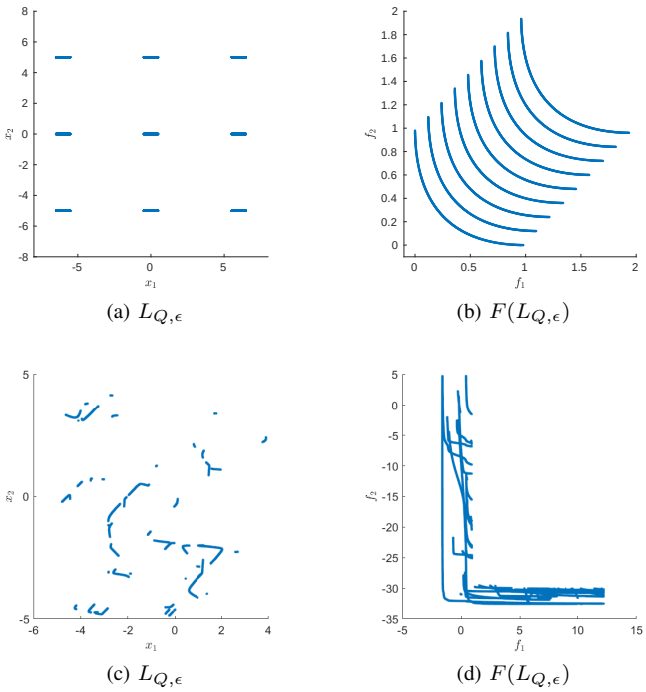

(a) $L_{Q,\epsilon}$ (b) $F(L_{Q,\epsilon})$ (c) $L_{Q,\epsilon}$ (d) $F(L_{Q,\epsilon})$

Fig. 2. The sets $L_{Q,\epsilon}$ and $F(L_{Q,\epsilon})$ for the two-decision variable problems SYM-PART9to9 (top row) and BBOB55 (bottom row).

A. Archiver

Algorithm 1 below shows the pseudo-code of $ArchiveUpdateL_{Q,\epsilon}$. Before we describe the algorithm in detail, we first discuss its acceptance strategy (lines 5 and 6), which is the core of the archiver. $L_{Q,\epsilon}$ is the intersection of $N_{Q,\epsilon}$ and L_Q, and the membership for each set cannot be decided directly during the algorithm's run. The Pareto set (for $N_{Q,\epsilon}$) is apparently not known beforehand, and local optimality (for L_Q) is hard to detect – at least if no gradient information is used (for which we opted for this study). Alternatively, we proceed as follows: We add a candidate solution d (at least temporarily) to the archive either if (i) it is non-dominated out of the considered candidate solutions from the entire run of the algorithm (this set is denoted by ND) or (ii) it ϵ-dominates one of the elements from ND and there exists no element b, considered candidate solution in the neighborhood of d, such that b dominates d. More precisely, we compare d to the current candidate solutions (current archive plus new population) within the δ_x-ball around d,

$$B_{\delta_x}(d) := \{x \in \mathbb{R}^n \ : \ \|x - d\|_2 \leq \delta_x\}. \tag{7}$$

The complete update process of the archiver is as follows. The archiver's inputs are the current archive A_0 and a population P, a new set of candidate solutions by which A_0 should be updated, and the thresholds $\epsilon \in \mathbb{R}_+^k$ (in objective space) and $\delta_x \in \mathbb{R}_+$ (in decision space). The set C is then defined as the union of A_0 and P. Next, compute the set of non-dominated points out of C, defined as ND, and the set of dominated solutions D. The new archive A is initialized with ND. Finally, all elements $d \in D$, which satisfy the condition (ii) described above, will be added to A.

Algorithm 1 $ArchiveUpdateL_{Q,\epsilon}$

Require: (objective-space-normalized) population P , archive A_0, thresholds $\epsilon \in \mathbb{R}^k_+$, $\delta_x \in \mathbb{R}_+$
Ensure: updated archive A
$C \leftarrow A_0 \cup P$
$ND \leftarrow nondom(C)$ ▷ non-dominated points
$D \leftarrow C \backslash ND$ ▷ dominated points
$A \leftarrow ND$
for all $d \in D$ **do**
  if $(\exists n \in ND : d \prec_\epsilon n)$ and $(\nexists b \in (B_{\delta_x}(d) \cap C) : b \prec d)$ **then**
    $A \leftarrow A \cup \{d\}$
  end if
end for
return A

It remains to determine the value of δ_x which is problem-dependent. Here, we have used five percent of the diagonal of the bounding box. That is, given the lower and upper bounds $l, u \in \mathbb{R}^n$ of the considered MOP, we have chosen

$$\delta_x := \frac{1}{20} \cdot \|u - l\|. \tag{8}$$

If the decision variables are in different ranges, one can of course use different values of δ_x in each coordinate (i.e., $\delta_x \in \mathbb{R}^n$), e.g., via using the neighborhood

$$\bar{B}_{\delta_x}(d) := \{x \in \mathbb{R}^n \ : \ |x_i - d_i| \leq \delta_{x_i}, \ i = 1, \ldots, n\}. \tag{9}$$

B. Cleaning

The archiver described above generates some outliers which is owed to the fact that a neighborhood of constant size is used to decide if a candidate solution could be locally optimal. We hence suggest applying a post-processing step that we discuss in the following on the final archive generated by $ArchiveUpdateL_{Q,\epsilon}$ (which we will refer to as *cleaning*). This post-processing step will be performed in decision space and on the dominated set of points, since we know the range of the decision variables in advance and all the outliers are dominated points. We want to remove points that are likely outliers, i.e., points that are isolated and far away from the rest (taking advantage of the fact that local and global solutions of MOPs typically form – at least locally – manifolds of a certain dimension). The clustering algorithm DBSCAN [60] that groups points based on density can identify outliers and does not need the number of clusters as input, making it an adequate choice for cleaning.

The only remaining step is to select appropriate values of $minpts$ and r, the two parameters of DBSCAN, which may be problem-dependent. Therefore, we propose to do a grid search on a small set of values for the parameters. For $minpts$, we use the values 2 and 3, and for r we first compute the average distance (in decision space) between dominated points $\bar{d}$, and then set $r = \{0.01\bar{d}, 0.02\bar{d}\}$, i.e. 1% and 2% of the average distance. This way we have an automated and problem-independent way of selecting the parameters. These values of r work in particular well for two and three decision variables. For higher dimensional problems (DTLZ and WFG) we used the same values, however, more adequate values might be selected which we leave for future work.

To select the *best* cluster out of the grid search, we need to use a function that measures cluster quality. For this we used the *weakest link* function defined in [61], which measures the smallest out of the longest path of points within a cluster. A summary of the cleaning process can be seen in Algorithm 2.

Algorithm 2 Cleaning

Input: Archive A
Output: Clean Archive A_{clean}
$ND \leftarrow nondom(A)$ ▷ non-dominated points
$D = \{d_1, d_2, \ldots\} \leftarrow A \backslash ND$ ▷ dominated points
Set $\bar{d} \leftarrow \frac{2}{|D|(|D|-1)} \sum_{d_i \neq d_j} \|d_i - d_j\|$ ▷ in decision space
Set $wlmin \leftarrow \infty$
for $minpts$ **in** $\{2, 3\}$ **do**
  for r **in** $\{0.01\bar{d}, 0.02\bar{d}\}$ **do**
    $C \leftarrow$ DBSCAN$(D, r, minpts)$ ▷ in decision space
    Set D_c as the union of clusters in C.
    $wl \leftarrow$ WeakestLink(D_c)
    if $wl \leq wlmin$ **then**
      $A_{clean} \leftarrow D_c \cup ND$
      $wlmin \leftarrow wl$
    end if
  end for
end for
return A_{clean}

C. Examples

We finally show some numerical results to show the effect of the archiver and the cleaning step. To this end, we feed the archiver with N candidate solutions that were chosen uniformly at random from the domain of the given MOP.

Figure 3 shows a numerical result for problem MMF1 [62]. The Pareto set of this problem consists of two connected components that both map to the entire Pareto front. Figures 3 (a) and (c) show the final archive A_f and its image $F(A_f)$ using $N = 100{,}000$. A_f contains 2,875 solutions which nicely cover both connected components, but also come with around 40 outliers (most visible in objective space). Figures 3 (b) and (d) show the result of the cleaning step on A_f and $F(A_f)$. The resulting set, called A_c, contains 2,630 elements, where in particular all of the outliers have been removed.

Figure 4 shows analog results for the more complex problem BBOB55 using $N = 2{,}000{,}000$. The final archive A_f (shown in black) consists of 9,495 points and captures almost the entire set of interest $L_{Q,\epsilon}$ (shown in blue), coming with around 400 outliers. The cleaned set A_c contains 7,289 elements and has a better overall approximation quality since all outliers (except for two) have been removed by the cleaning step.

Finally, we consider problem SYM-PART9to9 using the thresholds $\epsilon^{(1)} = (0.1, 0.1)^T$, $\epsilon^{(2)} = (0.4, 0.4)^T$ and $\epsilon^{(3)} = (1, 1)^T$ using $N = 100{,}000$ test points. Figure 5 shows numerical results (A_f and $F(A_f)$) of the archiver for $\epsilon^{(1)}$, $\epsilon^{(2)}$ and $\epsilon^{(3)}$ yielding approximations of one, four and nine out of

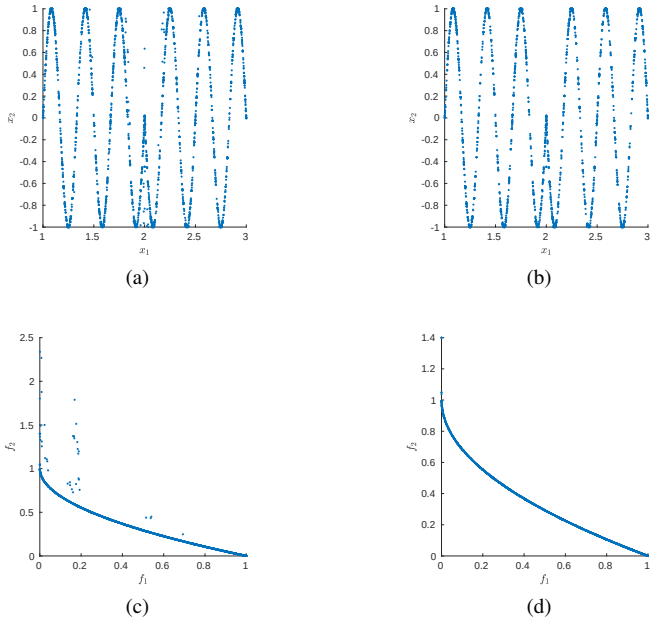

Fig. 3. Final archive for the MMF1 problem in decision (a) and objective space (c), respectively. Cleaned archive in the decision (b) and objective space (d), respectively.

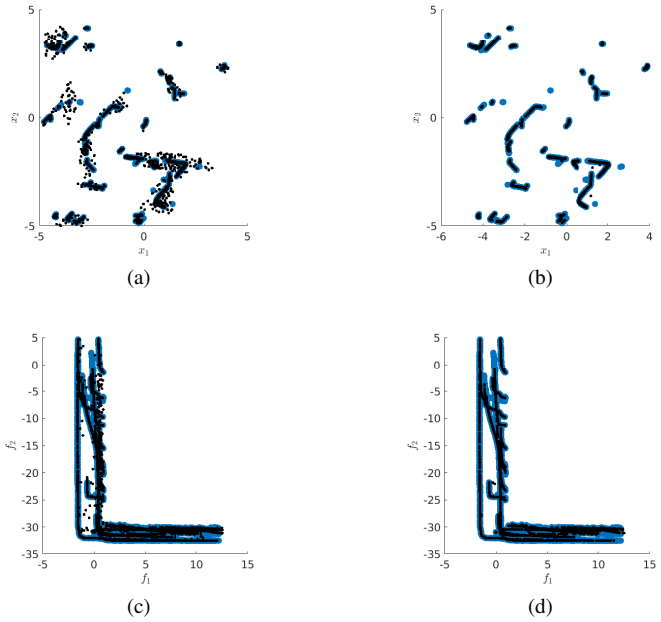

Fig. 4. (a), (c) Final archive for the BBOB55 problem in decision and objective space respectively. (b), (d) Cleaned archive in the decision and objective space respectively.

TABLE I
DISTANCES MEASURED BY Δ_2 OF THE IMAGES OF THE ARCHIVES $F(A_f)$ AND $F(A_c)$ TOWARD $F(L_{Q,\epsilon})$ FOR THE CONSIDERED TEST PROBLEMS.

Problem	$\Delta_2(F(A_f), F(L_{Q,\epsilon}))$	$\Delta_2(F(A_c), F(L_{Q,\epsilon}))$
MMF1	0.053945	0.013601
BBOB55	0.056124	0.050317
SYM-PART9to9 $\epsilon^{(1)}$	0.003018	0.003018
SYM-PART9to9 $\epsilon^{(2)}$	0.002850	0.002850
SYM-PART9to9 $\epsilon^{(3)}$	0.003184	0.003184

the nine components, respectively. We omitted the result of the cleaning step since it has no effect in these cases.

$L_{Q,\epsilon}$ is known analytically for SYM-PART9to9 and MMF1. In order to get a good approximation of $L_{Q,\epsilon}$ for BBOB55, we first fed $ArchiveUpdateL_{Q,\epsilon}$ with $1,000,000$ randomly chosen points. Next, we have used one point per connected component of the final approximation as starting point for the multi-objective continuation method Pareto Tracer [63].

Table I shows the approximation qualities of the image of the archives w.r.t. $F(L_{Q,\epsilon})$ using the averaged Hausdorff distance Δ_2 (i.e., $p = 2$) for all problems considered above which confirms the visual observations: Δ_2 goes down for MMF1 and BBOB55 because of all the outliers removed, and remains the same for SYM-PART9to9 where the cleaning had no effect.

V. INTEGRATION INTO MOEAS

So far, no specialized algorithm exists that aims for $L_{Q,\epsilon}$ approximations. However, it is, of course, possible to equip in principle any MOEA with $ArchiveUpdateL_{Q,\epsilon}$ as an external archive, as we will briefly describe in the following. Doing so, new $L_{Q,\epsilon}$ solvers are generated, which we will further investigate in the following sections.

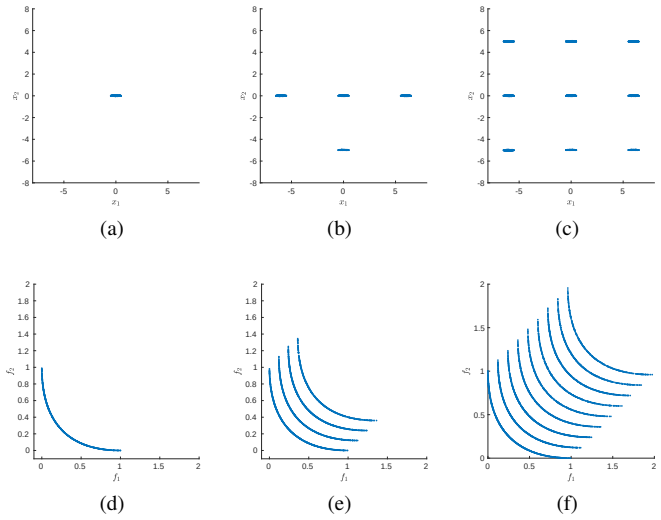

Fig. 5. Result of $ArchiveUpdateL_{Q,\epsilon}$ on SYM-PART9to9 using $\epsilon^{(1)} = (0.1, 0.1)^T$ ((a) and (d)), $\epsilon^{(2)} = (0.4, 0.4)^T$ ((b) and (e)), and $\epsilon^{(3)} = (1, 1)^T$ ((c) and (f)). Shown are the final archives and their images without cleaning. The cleaning step does not change the archives for all cases.

Algorithm 3 shows a framework that incorporates $ArchiveUpdateL_{Q,\epsilon}$ as external archiver to a given MOEA. P_i and A_i denote the i-th population and archive in generation i, respectively. At first, the initial population P_0 is initialized (depending on the chosen MOEA), and A_0 is built by the archiver based on this set. Next, the processes to generate new candidate solutions and the archiver are applied in a loop. Hereby, $Generate()$ symbolically refers to the generation process of the chosen MOEA. If the termination conditions are satisfied, the loop is terminated. The final population P_f

is simply the one that has been generated last by the MOEA. For the final archive A_f the findings are similar, however, first, the cleaning step is applied in order to remove possible outliers from the candidate set (only once at the end of the search).

Algorithm 3 MOEA equipped with $ArchiveUpdateL_{Q,\epsilon}$

Require: problem (MOP), chosen MOEA
Ensure: final population P_f, final archive A_f
1: $P_0 \leftarrow$ initialize()
2: $A_0 \leftarrow ArchiveUpdateL_{Q,\epsilon}(P_0, \emptyset)$
3: $t \leftarrow 0$
4: **while** Condition **do**
5: $\quad P_{t+1} = Generate(P_t)$
6: $\quad A_{t+1} \leftarrow ArchiveUpdateL_{Q,\epsilon}(P_{t+1}, A_t)$
7: $\quad t \leftarrow t+1$
8: **end while**
9: $P_f \leftarrow P_t$
10: $A_f \leftarrow$ Cleaning(A_t)
11: **return** $\{P_f, A_f\}$

For the efficient use of such an archive-equipped MOEA we suggest following the approach proposed in [64], which has a significant impact on the overall performance: Specifically, we suggest using the following modified objectives instead of the original objectives f_i, $i = 1, \ldots, k$, of the given MOP:

$$\bar{f}_i(x) = (1-\alpha) f_i(x) + \frac{\alpha}{k} \sum_{i=1}^{k} f_i(x), \quad i = 1, \ldots, k, \quad (10)$$

where $\alpha > 0$ is a "small" value. We have used $\alpha = 0.02$ for the computations presented here. The reason for this is that $ArchiveUpdateL_{Q,\epsilon}$ computes the non-dominated solutions out of the given candidate solutions (line 2 of Algorithm 1). If a MOP, however, contains weakly optimal solutions that are not globally optimal, stochastic search algorithms such as MOEAs may find and keep such "dominance-resistant solutions" which may have a negative impact on the solution quality. The use of the modified objectives helps to remove these dominance-resistant solutions. Finally, since the individual objectives' values may be in different ranges, we apply the normalization strategy in objective space as done in [16].

VI. A Performance Indicator for $L_{Q,\epsilon}$

Here, we propose a new performance indicator that is related to $L_{Q,\epsilon}$. A possible choice is certainly to compute the distance between the set of interest and the given archive A, e.g., to use $\Delta_2(A, L_{Q,\epsilon})$ and $\Delta_2(F(A), F(L_{Q,\epsilon}))$ in decision and objective space, respectively. Such distance calculations, however, are only possible if $L_{Q,\epsilon}$ is known completely and are therefore restricted to low-dimensional benchmark problems. The same discussion holds for any distance based indicators such as IGD and PSP [41]. Instead, we propose in the following to consider the "essentially different realizations" (in decision space) for certain target values (in objective space) enabled by A with respect to a reference set Y given in objective space. We will first consider the idea for one target value and then propose the indicator defined for a given MOP.

Assume we are given an archive $A \subset \mathbb{R}^n$ and the thresholds $\epsilon \in \mathbb{R}^k_+$ and $\delta_x > 0$. Further, assume we are given an element (target value) $y \in \mathbb{R}^k$, then

$$r(y, A, \epsilon) := \{a \in A \ : \ |f_i(a) - y_i| < \epsilon_i, \ \forall i = 1, \ldots, k\} \quad (11)$$

represents the set of possible realizations of y within A and using the threshold ϵ – assuming that any point $a \in A$ for which $F(a)$ is sufficiently close to y (defined by ϵ) is acceptable. Note that the magnitude of this set is not conclusive for our purpose: If $a_1 \in r(y, A, \epsilon)$, then any element $a_2 \in A$ that is sufficiently close to a_1 is certainly also included in $r(y, A, \epsilon)$, but may not represent a different realization for the decision maker. To get an estimation of the essentially different solutions, we suggest to proceed as follows: Let $C = C(\delta_x)$ be a collection of uniform boxes (hypercubes) that form a partition of (a superset of) the domain Q. Further, let the lengths of the boxes be given by δ_x. Then we say that two solutions $a_1, a_2 \in A$ that are included in two different boxes represent two essentially different realizations of y. Consequently, the number of *essentially different realizations* (edr) is given by the number of boxes $c \in C$ that intersect with $r(y, A, \epsilon)$. We hence define

$$edr(y, A, \epsilon, \delta_x) := |\{c \in C \ : \ r(y, A, \epsilon) \cap c \neq \emptyset\}|. \quad (12)$$

After having defined the value edr for one element $y \in \mathbb{R}^k$, we are now in the position to propose the indicator I_{EDR}: Given a finite-size representation $Y \subset \mathbb{R}^k$, we define the indicator value as the average of the edr values over all elements of Y,

$$I_{EDR}(Y) := \frac{1}{|Y|} \sum_{y \in Y} edr(y, A, \epsilon, \delta_x). \quad (13)$$

Hereby, the elements of Y should be uniformly distributed around the entire Pareto front, for each $y \in Y$ it ideally holds

$$\delta_y = dist(y, Y \backslash \{y\}) = \min_{\bar{y} \in Y \backslash \{y\}} \|y - \bar{y}\|_2 \approx \|\epsilon\|_\infty. \quad (14)$$

We stress that generating the entire box collection C that covers Q is not required. Instead, we can proceed as done in cell mapping techniques [65]: Given the lower and upper bounds of the problem, $l, u \in \mathbb{R}^n$, for a point $x \in r(y, A, \epsilon)$ we can compute the the box identifier z_i as follows:

$$z_i = \left\lceil \frac{x_i - l_i}{\delta_x} \right\rceil, \ i = 1, \ldots, n. \quad (15)$$

Their feasibility can easily be checked – $z_i \in \left[0, \left\lceil \frac{u_i - l_i}{\delta_x} \right\rceil\right]$ – and duplicate entries can be identified via considering the identifiers from $r(y, A, \epsilon)$ leading to $edr(y, A, \epsilon, \delta_x)$.

We finally stress that I_{EDR} is not Pareto compliant. This is due to the fact that an essentially different solution may be non-optimal. Figure 6 shows one example on MMF1 using the archive A shown in Figure 3 (c), the target value $y = (0.769, 0.123)^T$, and the thresholds $\epsilon = (0.1, 0.1)^T$ and $\delta_x = 0.4$. For this setting, we obtain 46 elements in $r(y, A, \epsilon)$, which fall into two relatively small clusters. These pre-images lie within two boxes, leading to $edr(y, A, \epsilon, \delta_x) = 2$ for the chosen thresholds. Smaller values of δ_x – and thus, smaller box sizes – may lead to edr values larger than two. However,

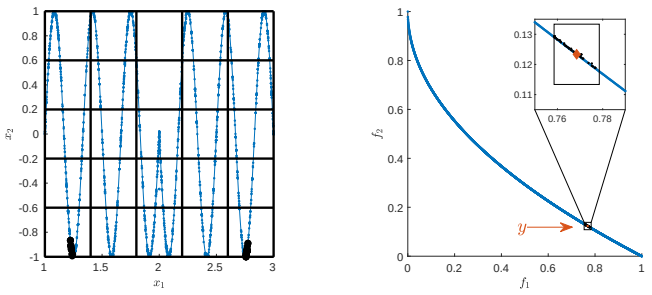

(a) decision space (b) objective space

Fig. 6. Hypothetical example on MMF1 using $\epsilon = (0.1, 0.1)^T$ and $\delta_x = 0.4$. Left: the Pareto set (blue line), the archive A (blue dots), and the realizations $r(y, A, \epsilon)$ (black dots). Right: the Pareto front (blue line), the target value y (red diamond) and the images of $r(y, A, \epsilon)$ (black dots). The inset plot in the top right corner provides a zoomed-in view of the neighborhood around y. For the chosen values of the thresholds we obtain $edr(y, A, \epsilon, \delta_x) = 2$.

this does not contradict the definition of edr since elements on the same connected component can indeed represent different realizations for the DM.

VII. Experimental Study

In this section, we empirically investigate the effect of integrating the archiver into different state-of-the-art MOEAs.

A. Experimental Setup

To assess the archiver's applicability, we performed an empirical analysis considering 13 algorithms and 43 problems from the literature. We selected NSGA-II [13], NSGA-III [16], MOEA/D [15] and SMS-EMOA [17] as representatives of classical MOEAs. Next, we selected CMMO [66], DN-NSGA-II [67], HHC-MMEA [68], HREA [6], MMEA-WI [69], MO_Ring_PSO_SCD [41], MPMMEA [70], TriMOEATAR [71] from multi-objective multimodal evolutionary algorithms, and a random search procedure as a baseline of the approach.

For the problems we selected MMF1-8 ($n = 2, k = 2$) [71], MMOP1-6 ($n = 3, k = 2$)[41], DTLZ1-7 ($n \in \{7, 12, 22\}, k = 3$)[72], WFG1-9 ($n = 12, k = 3$)[73], Aspar ($n = 2, k = 2$)[51], Bi-Rosenbrock ($n = 2, k = 2$)[55], bbob-biobj1, bbob-biobj2, bbob-biobj10, BBOB55 (cf. above) ($n = 2, k = 2$) [46], Deb99 ($n = 2, k = 2$) [74], Two-on-One ($n = 2, k = 2$) [59], SYM-PART and SYM-PART9to9 ($n = 2, k = 2$) [59], SSW ($n = 3, k = 2$)[75], LSS ($n = 5, k = 2$) [76], Omni ($n = 4, k = 2$) [77], and twelve problems from the imbalanced distance minimization benchmark problems with ϵ-efficient solutions (IDMP_e) [6] have been used ($n = \{2, 3, 4\}, k = \{2, 3, 4\}$). In all cases, n is the number of decision variables, k is the number of objectives and we normalize the objective space as done in [16]. All of these problems have already been used in the context of MMO and approximate solutions. All algorithms and problems were used from the PlatEMO platform version 4.2 [78].

We performed 20 independent executions of each (algorithm, problem)-combination and computed the corresponding indicators (hypervolume, IGDX (for $k = 2$), Δ_2 in decision variable space (for $k = 2$), and EDR indicator). We use $L_{Q,\epsilon}$ as the reference set for IGDX and Δ_2; for the hypervolume measure, we specify the reference point as the approximation of the nadir point plus 10% of the difference between the nadir and the ideal point; and for EDR we use discretizations Y of the PF.

Further, for computing $L_{Q,\epsilon}$, we used the reference Pareto sets/fronts provided in PlatEMO for MMF1-8 and MMOP1-6 since $L_{Q,\epsilon}$ coincides with the PS in those cases. For Deb99, Two-on-One, SYM-PART, SYM-PART9to9 and Omni, the PF and PS is known, so a uniform sampling was done. For the rest of the bi-objective problems (SSW, Aspar, Bi-Rosenbrock, bbob-biobj1, bbob-biobj2, bbob-biobj10, BBOB55, LSS), we performed a grid search of $1{,}000^n$ points, where n is the number of decision variables. For the three-objective problems (WFG and DTLZ), a grid search computation of $L_{Q,\epsilon}$ is unfeasible since all these problems have at least seven decision variables. We include only HV and EDR for these problems.

Finally, to obtain Y for the EDR computations, first we obtain the PF component out of the references described above, and then it was reduced so that each point is a distance of $\delta_y = 0.05$ from each other. This was made to simulate the decision making process, and obtain the finite size representation Y as required for I_{EDR}.

The parameters used for the experiments were as follows: Population size: 100; Number of generations: 100; $\epsilon = (0.2, \ldots, 0.2)^T$ (on normalized objectives); $\delta_x = \|u - l\|/20$; $\delta_y = 0.2$; Crossover: 1; η_c: 20; η_m: 20; Mutation: $1/n$. All other algorithm specific algorithms were set to PlatEMO's defaults. The locally efficient points are approximated using the visualization approach and software PLOT [51] and then filtered according to the ϵ-optimal solutions for each problem. We will publish the source code to replicate all experiments, intermediate results, and experimental analyses after the acceptance of this work.

B. Experimental Results

We view the experimental results from two different perspectives: First, we evaluate whether the utilization of the new archive (in comparison to the final population) is beneficial for any of the considered indicators (EDR, HV, Δ_2, and IGDX). Then, we will take a closer look at the overall performance of all algorithm configurations run in our study.

1) Archiver vs. Final Population: We start by illustrating the overall performance for each indicator, aggregated on all runs and problems, in Figure 7. Note that we normalize the indicator values per problem with min-max normalization to the $[0, 1]$ range to allow for better comparability between different indicator value ranges per problem instance. In this aggregated view, it becomes already clear that the archiver generally outperforms the final population, regardless of the indicator. Δ_2 and IGDX give qualitatively similar results to the EDR indicator, validating EDR as useful measure for MMO.

We then perform a Wilcoxon rank-sum test ($p = 0.05$) for each combination of algorithm and test problem for all indicators. Aggregated results for each optimizer are reported in Table II for EDR and in Table III for HV. Results on the IGDX and Δ_2 indicators (not shown) reflect the results

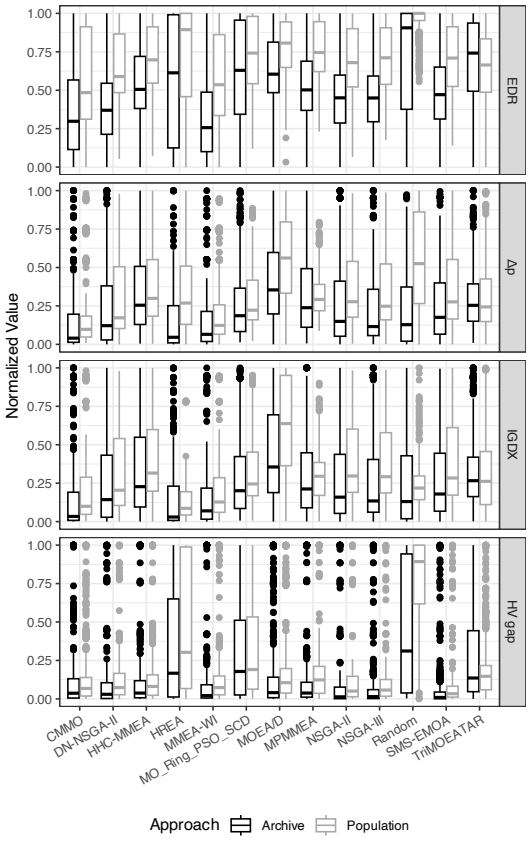

Fig. 7. Boxplots of the algorithms' performances for all indicators with and without archiver. Indicator values for each problem are normalized to $[0, 1]$ (lower is better) to aggregate across problem instances.

TABLE II
WINS, TIES, AND LOSSES ACROSS TEST PROBLEMS BASED ON WILCOXON RANK-SUM TESTS FOR THE ARCHIVER AND POPULATION-BASED RESULTS FOR EDR.

Optimizer	Archive better	Ties	Final population better
CMMO	**39**	12	4
DN-NSGA-II	**43**	12	0
HHC-MMEA	**45**	8	2
HREA	**41**	13	1
MMEA-WI	**40**	13	2
MO_Ring_PSO_SCD	**31**	21	3
MOEA/D	**42**	13	0
MPMMEA	**51**	3	1
NSGA-II	**50**	4	1
NSGA-III	**51**	4	0
Random	**46**	9	0
SMS-EMOA	**46**	9	0
TriMOEATAR	**33**	6	16

TABLE III
WINS, TIES, AND LOSSES ACROSS TEST PROBLEMS BASED ON WILCOXON RANK-SUM TESTS FOR THE ARCHIVER AND POPULATION-BASED RESULTS ON THE HV INDICATOR.

Optimizer	Archive better	Ties	Final population better
CMMO	**36**	11	8
DN-NSGA-II	**46**	5	4
HHC-MMEA	**43**	8	4
HREA	**45**	10	0
MMEA-WI	**35**	14	6
MO_Ring_PSO_SCD	26	**27**	2
MOEA/D	**42**	9	4
MPMMEA	**48**	5	2
NSGA-II	**43**	9	3
NSGA-III	**43**	8	4
Random	**46**	9	0
SMS-EMOA	**34**	15	6
TriMOEATAR	19	10	**26**

for EDR. It can clearly be seen that each optimizer performs statistically significantly better on most problem instances with the archiver than without. It is important to note that the cleaning step only removes dominated solutions, therefore the HV indicator will be unaffected by it.

In general, the performance of TriMOEATAR is not nearly as much improved by the archiver as for other algorithms, and its population even outscores the archiver in the HV indicator evaluation. However, it is not a particularly competitive optimizer overall, as is illustrated in Figure 7. We expect the underwhelming archiver performance here is rooted in problems with the TriMOEATAR optimizer.

2) Overall Evaluation: To evaluate the performance across all test problems, we rely on critical difference (CD) plots [79]. These show each optimizer's average rank and a CD stemming from a Nemenyi test ($\alpha = 0.05$) for pairwise comparisons. Classifiers are not significantly different according to this test when they are within a CD of each other, and such groups of classifiers are joined by a horizontal line. The CD plots for the EDR and HV indicators are shown in Figure 8.

A clear shift in preferences is visible between the indicators, where the MMO-based algorithms CMMO, MMEA-WI, and DN-NSGA-II perform better regarding EDR, while the top algorithms for HV include more globally-oriented ones such as SMS-EMOA, NSGA-II, and NSGA-III. Also, for EDR, some optimizers are tied with random search (MOEA/D and MO_Ring_PSO_SCD), and TriMOEATAR is even beaten by it. Note that random search is by far the weakest standalone algorithm. That is, its performance here can be fully attributed to the archiving strategy. Similar observations can be made for the rankings w.r.t. IGDX and Δ_2, but not for HV, where random search is tied for worst performance, again highlighting the complementarity of the approaches.

C. Sensitivity Analysis

The proposed archiver has two parameters: ϵ and δ_x. We assume in this work that ϵ is given by the DM (and hence, fixed). In the following we will perform a small sensitivity analysis for δ_x on SYM-PART9to9 and BBOB55. The behavior of ϵ for SYM-PART9to9 can be seen in Figure 5.

1) δ_x Experimental Setup: We performed a sensitivity analysis using the values:

$$\delta_x := \frac{1}{\delta_i} \cdot \|u - l\|, \tag{16}$$

for $\delta_i \in \{1, 5, 10, 15, 20, 25\}$ and applied $ArchiveUpdateL_{Q,\epsilon}(P, A_0, \epsilon, \delta_x)$. It is important to remember that δ_x is inversely proportional to δ_i. The parameters used were: one million points sampled uniformly at random (P); $A_0 = \emptyset$; $\epsilon = (1, 1)$ for SYM-PART9to9 and $\epsilon = (2.5, 2.5)$ for BBOB55; and δ_x as in Eq. (16). Then, we

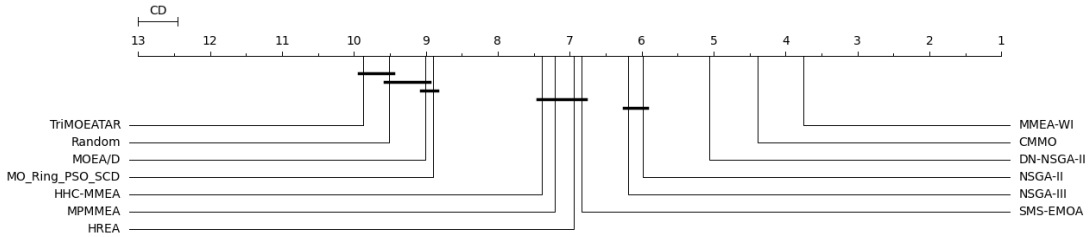

(a) EDR indicator

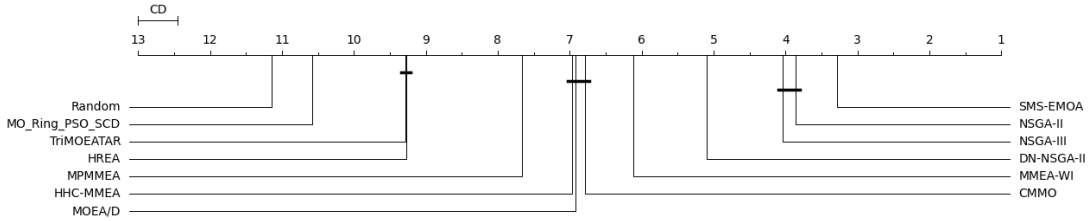

(b) HV indicator

Fig. 8. Critical difference plots for the EDR (top) and HV (bottom) indicators for all algorithms with archiver. The scale is given in ranks, i.e., lower is better.

measured the following performance indicators: IGD; EDR; $IGDX$; HV; computation time; and archive size for all the δ_x values.

2) δ_x Results: We show the plots for all the indicators as a function of δ_x in Figure 9 and the plots of the final archives in Figures 10 and 11.

a) SYM-PART9to9 Results: Before discussing the indicator plots, we need to observe in Figure 10 that for the value of $\delta_i = 1$ we get only the component that corresponds to the PF (as expected, since $\delta_i = 1$ represents the entire decision space), and for the rest of the δ_i values we obtain all the local components. Next, looking at the indicator plots (Figure 9) for SYM-PART9to9, we can see the expected result: EDR has its lowest value for $\delta_i = 1$ because we only have one component, and for the remaining values we have an almost constant value for EDR, even though the archive size increases. Note that IGD and HV remain constant, since regardless of the δ_x value, the archiver will always capture the non-dominated points. Further, in the IGDX case, we use $L_{Q,\epsilon}$ as the reference. This explains why the highest value for IGDX is obtained for $\delta_i = 1$, where we have only one out of the nine components. Finally, we observe a linear increase in the size of the archive and a nonlinear increase in the elapsed time.

b) BBOB55 Results: Before discussing the indicator plots, we need to observe in Figure 11 that, as we increase the value of δ_i, we get more local components as expected. Looking at the indicator plots (Figure 9), we can see the expected results: EDR is increasing with δ_i starting at $\delta_i = 15$. This is due to the parameters δ_x and δ_y of EDR, for which we used the same values as in our experiments ($\delta_x = \frac{1}{20}\|u - l\|$ and $\delta_y = 0.2$). HV is almost constant as expected. IGD remains constant and goes down for $\delta_i = 20$ and $\delta_i = 25$. In these two cases we observe (Figure 11f for $\delta_i = 25$) that a local front near the PF (near the reference points $(-2, -5)$ and $(-2, -1)$) is kept by the archiver due to the smaller value of δ_x, resulting in extra (non-dominated) points closer to the PF and therefore improving IGD. IGDX has the same behavior as in SYM-PART9to9 since the reference used is an approximation of $L_{Q,\epsilon}$ and not only the PS, and therefore the more components we find (with increasing δ_i), the higher IGDX will be. Finally, we observe an almost linear increase in the size of the archive and a nonlinear increase in the elapsed time.

VIII. Conclusions and Future Work

In this paper, we have addressed the problem of computing all locally optimal solutions that are nearly optimal which can be used in the context of multimodal optimization (MMO). To this end, we have first proposed and discussed a set of interest, $L_{Q,\epsilon}$, which is closely related to the previously proposed set of ϵ-acceptable solutions [6].

To capture $L_{Q,\epsilon}$ via evolutionary multi-objective optimization (EMO) techniques, we have proposed a novel archiver, $ArchiveUpdateL_{Q,\epsilon}$. This archiver is unbounded and aims to approximate the entire set of interest in the limit. To assess the performance of EMO strategies, we have proposed the performance indicator I_{EDR}. This indicator is based on the "essentially different realizations" for given points y on the Pareto front and strongly relates to $L_{Q,\epsilon}$. Finally, we have presented some numerical results, showing that the inclusion of the archiver into a wide range of MOEAs improves their performance w.r.t. I_{EDR}, as well as w.r.t. the hypervolume indicator, significantly for all but one of thirteen optimizers. Further analyses show that even a random search equipped with the archiver performs adequately, emphasizing the archiver's impact on the solution quality.

Though the results presented here are very promising, there are several possible paths of future work that are worth investigating. First of all, it would be desired to design a specialized MOEA that aims for $L_{Q,\epsilon}$ approximations. In this context, it would also be beneficial to have further bounded archivers aiming for this new set of interest. Another interesting aspect

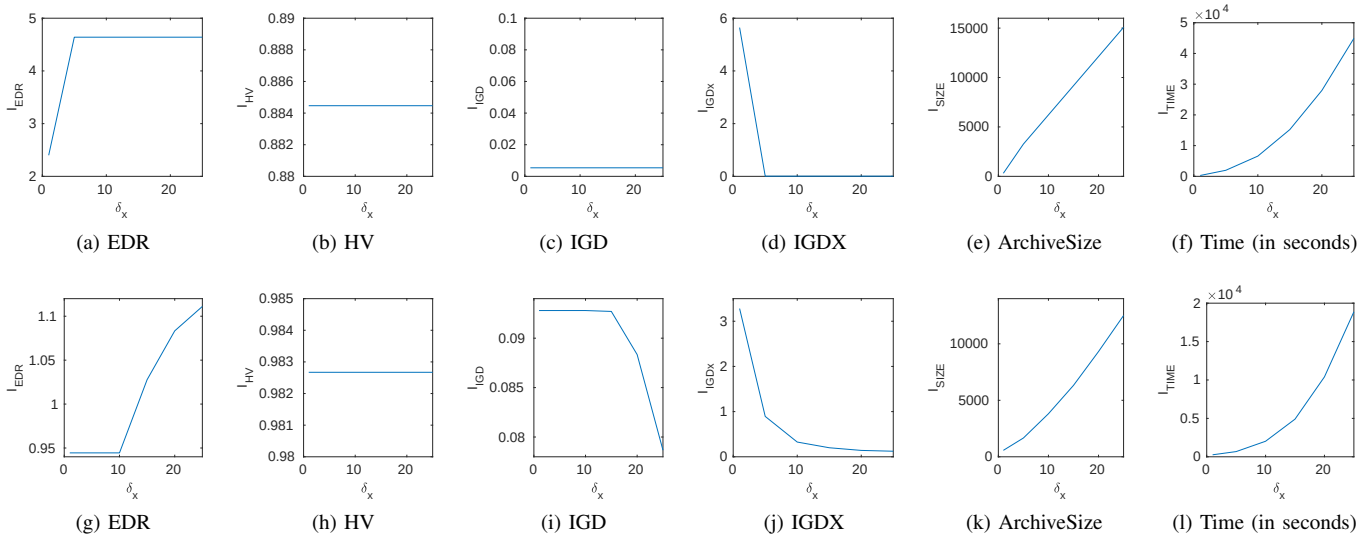

(a) EDR (b) HV (c) IGD (d) IGDX (e) ArchiveSize (f) Time (in seconds)

(g) EDR (h) HV (i) IGD (j) IGDX (k) ArchiveSize (l) Time (in seconds)

Fig. 9. Plots of the indicators for different values of δ_x and for the problems SYM-PART9to9 (first row) and BBOB55 (second row).

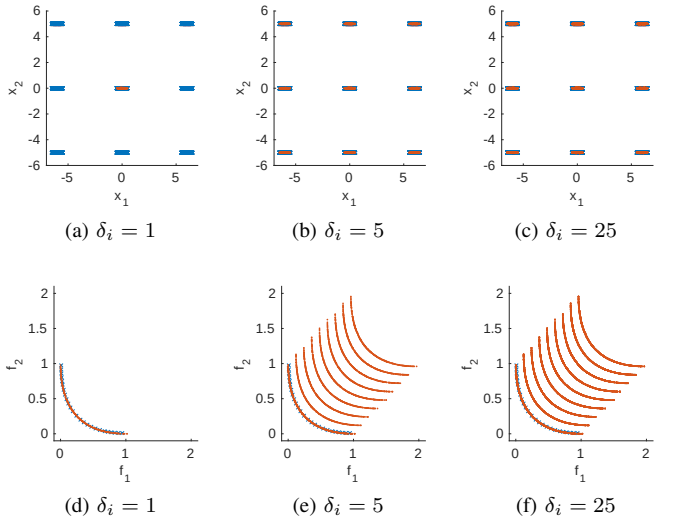

(a) $\delta_i = 1$ (b) $\delta_i = 5$ (c) $\delta_i = 25$

(d) $\delta_i = 1$ (e) $\delta_i = 5$ (f) $\delta_i = 25$

Fig. 10. Plots of the final archive for different values of δ_x and for the problem SYM-PART9to9. The first column is decision space, the second column is objective space. Blue crosses represent the reference set, and red dots represent the final archive.

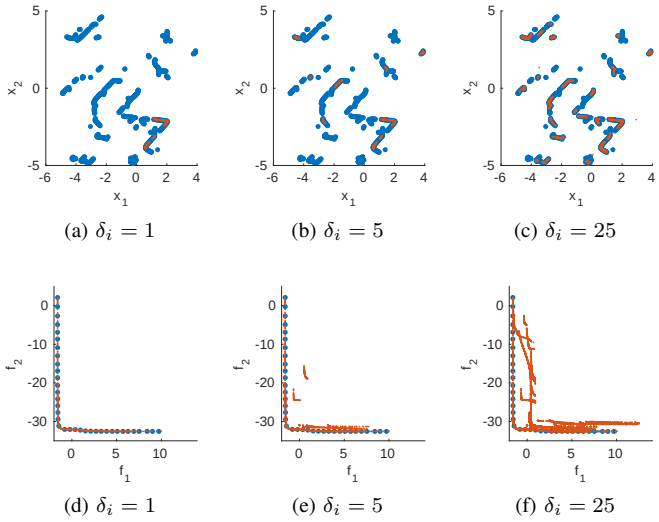

(a) $\delta_i = 1$ (b) $\delta_i = 5$ (c) $\delta_i = 25$

(d) $\delta_i = 1$ (e) $\delta_i = 5$ (f) $\delta_i = 25$

Fig. 11. Plots of the final archive for different values of δ_x and for the problem BBOB55. The first column is decision space, the second column is objective space. Blue-filled dots represent the reference set, and red dots represent the final archive.

would be to incorporate gradient information. This could help increase the security of the information if a candidate solution is locally optimal. As a by-product, this should reduce the number of outliers maintained by the archiver.

A third aspect will be the tuning of the cleaning parameters for problems with more than three decision variables. Finally, we would like to stress that the restriction to locally optimal solutions is *one* way to reduce the set of nearly optimal solutions. For further studies it is conceivable to discuss other possible discretization/reduction strategies.

REFERENCES

[1] O. Schütze, C. A. Coello Coello, and E. G. Talbi, "Approximating the ϵ-efficient set of an MOP with stochastic search algorithms," in *MICAI 2007: Advances in Artificial Intelligence, 6th Mexican International Conference on Artificial Intelligence* (A. F. Gelbukh and A. F. K. Morales, eds.), Lecture Notes in Computer Science, pp. 128–138, Springer, 2007.

[2] O. Schütze, C. A. Coello Coello, E. Tantar, and E.-G. Talbi, "Computing finite size representations of the set of approximate solutions of an MOP with stochastic search algorithms," in *Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pp. 713–720, 2008.

[3] O. Schütze, M. Vasile, and C. A. Coello Coello, "Approximate solutions in space mission design," in *Proceedings of the International Conference on Parallel*

Problem Solving from Nature (PPSN 2008) (G. Rudolph, T. Jansen, S. M. Lucas, C. Poloni, and N. Beume, eds.), pp. 805 – 814, Springer Berlin Heidelberg, 2008.

[4] O. Schütze, C. Hernández, E.-G. Talbi, J. Q. Sun, Y. Naranjani, and F. R. Xiong, "Archivers for the representation of the set of approximate solutions for MOPs," *Journal of Heuristics*, vol. 5, no. 1, pp. 71 – 105, 2019.

[5] O. Schütze, A. E. Rodriguez-Fernandez, C. Segura, and C. Hernández, "Finding the set of nearly optimal solutions of a multi-objective optimization problem," *IEEE Transactions on Evolutionary Computation*, pp. 1–1, 2024.

[6] W. Li, X. Yao, T. Zhang, R. Wang, and L. Wang, "Hierarchy ranking method for multimodal multiobjective optimization with local Pareto fronts," *IEEE Transactions on Evolutionary Computation*, vol. 27, no. 1, pp. 98–110, 2022.

[7] O. Schütze, X. Esquivel, A. Lara, and C. A. Coello Coello, "Using the averaged Hausdorff distance as a performance measure in evolutionary multi-objective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 16, no. 4, pp. 504–522, 2012.

[8] J. M. Bogoya, A. Vargas, and O. Schütze, "The averaged Hausdorff distances in multi-objective optimization: A review," *Mathematics*, vol. 7, no. 10, 2019.

[9] G. Rudolph, O. Schütze, C. Grimme, C. Domínguez-Medina, and H. Trautmann, "Optimal averaged Hausdorff archives for bi-objective problems: theoretical and numerical results," *Computational Optimization and Applications*, vol. 64, pp. 589–618, Jun 2016.

[10] A. Zhou, Q. Zhang, and Y. Jin, "Approximating the set of Pareto-optimal solutions in both the decision and objective spaces by an estimation of distribution algorithm," *IEEE Transactions on Evolutionary Computation*, vol. 13, no. 5, pp. 1167–1189, 2009.

[11] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach," *IEEE Transactions on Evolutionary Computation*, vol. 3, no. 4, pp. 257–271, 1999.

[12] C. A. Coello Coello, E. Goodman, K. Miettinen, D. Saxena, O. Schütze, and L. Thiele, "Interview: Kalyanmoy Deb talks about formation, development and challenges of the EMO community, important positions in his career, and issues faced getting his works published," *Mathematical and Computational Applications*, vol. 28, no. 2, 2023.

[13] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 182–197, 2002.

[14] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimization," in *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)* (K. Giannakoglou *et al.*, eds.), pp. 95–100, International Center for Numerical Methods in Engineering (CIMNE), 2002.

[15] Q. Zhang and H. Li, "MOEA/D: A multi-objective evolutionary algorithm based on decomposition," *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.

[16] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints.," *Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2014.

[17] N. Beume, B. Naujoks, and M. Emmerich, "SMS-EMOA: Multiobjective selection based on dominated hypervolume," *European Journal of Operational Research*, vol. 181, no. 3, pp. 1653–1669, 2007.

[18] S. Zapotecas-Martínez, A. López-Jaimes, and A. García-Nájera, "LIBEA: A Lebesgue indicator-based evolutionary algorithm for multi-objective optimization," *Swarm and Evolutionary Computation*, vol. 44, pp. 404 – 419, 2019.

[19] J. D. Knowles and D. W. Corne, "Approximating the nondominated front using the Pareto archived evolution strategy," *Evolutionary Computation*, vol. 8, no. 2, pp. 149 – 172, 2000.

[20] J. D. Knowles and D. W. Corne, "Bounded Pareto archiving: Theory and practice," in *Metaheuristics for Multiobjective Optimisation*, pp. 39 – 64, Springer, 2004.

[21] J. D. Knowles, D. W. Corne, and M. Fleischer, "Bounded archiving using the Lebesgue measure," in *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 2490 – 2497, IEEE Press, 2003.

[22] M. López-Ibáñez, J. D. Knowles, and M. Laumanns, "On sequential online archiving of objective vectors," in *Evolutionary Multi-Criterion Optimization (EMO 2011)*, pp. 46 – 60, Springer, Berlin, Heidelberg, 2011.

[23] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler, "Combining convergence and diversity in evolutionary multiobjective optimization," *Evolutionary computation*, vol. 10, pp. 263 – 282, 2002.

[24] O. Schütze, M. Laumanns, E. Tantar, C. A. Coello Coello, and E.-G. Talbi, "Computing gap free Pareto front approximations with stochastic search algorithms," *Evolutionary Computation*, vol. 18, no. 1, pp. 65–96, 2010.

[25] O. Schütze and C. Hernández, *Archiving Strategies for Evolutionary Multi-objective Optimization Algorithms*. Springer, 2021.

[26] C. I. H. Castellanos and O. Schütze, "A bounded archiver for Hausdorff approximations of the Pareto front for multi-objective evolutionary algorithms," *Mathematical and Computational Applications*, vol. 27, no. 3, 2022. Article Number: 48.

[27] M. Laumanns and R. Zenklusen, "Stochastic convergence of random search methods to fixed size Pareto front approximations," *European Journal of Operational Research*, vol. 213, pp. 414 – 421, 2011.

[28] J. Fieldsend, R. Everson, and S. Singh, "Using unconstrained elite archives for multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 3, pp. 305–323, 2003.

[29] O. Schütze, C. A. C. Coello, S. Mostaghim, M. Dellnitz,

and E.-G. Talbi, "Hybridizing evolutionary strategies with continuation methods for solving multi-objective problems," *IEEE Transactions on Evolutionary Computation*, vol. 19, no. 7, pp. 762 – 769, 2008.

[30] R. Wang, Z. Zhou, H. Ishibuchi, T. Liao, and T. Zhang, "Localized weighted sum method for many-objective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 1, pp. 3–18, 2016.

[31] L. M. Pang, H. Ishibuchi, and K. Shang, "Algorithm configurations of MOEA/D with an unbounded external archive," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1087–1094, 2020.

[32] H. Ishibuchi, L. M. Pang, and K. Shang, "A new framework of evolutionary multi-objective algorithms with an unbounded external archive," in *ECAI 2020*, pp. 283–290, IOS Press, 2020.

[33] Y. Peng and H. Ishibuchi, "A diversity-enhanced subset selection framework for multimodal multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 26, no. 5, pp. 886–900, 2021.

[34] Z. Wang, B. Mao, H. Hao, W. Hong, C. Xiao, and A. Zhou, "Enhancing diversity by local subset selection in evolutionary multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, 2022.

[35] M. Li, M. López-Ibáñez, and X. Yao, "Multi-objective archiving," *IEEE Transactions on Evolutionary Computation*, pp. 1–1, 2023.

[36] M. Preuss, *Multimodal optimization by means of evolutionary algorithms*. Springer, 2015.

[37] P. Kerschke, H. Wang, M. Preuss, C. Grimme, A. Deutz, H. Trautmann, and M. Emmerich, "Towards Analyzing Multimodality of Continuous Multiobjective Landscapes," in *Parallel Problem Solving from Nature–PPSN XIV: 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings 14*, pp. 962–972, Springer, 2016.

[38] R. Tanabe and H. Ishibuchi, "A review of evolutionary multimodal multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 24, no. 1, pp. 193–200, 2019.

[39] C. Grimme, P. Kerschke, P. Aspar, H. Trautmann, M. Preuss, A. H. Deutz, H. Wang, and M. Emmerich, "Peeking beyond peaks: Challenges and research potentials of continuous multimodal multi-objective optimization," *Computers & Operations Research*, vol. 136, p. 105489, 2021.

[40] G. Rudolph and M. Preuss, "A multiobjective approach for finding equivalent inverse images of Pareto-optimal objective vectors," in *2009 IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM)*, pp. 74–79, IEEE, 2009.

[41] C. Yue, B. Qu, and J. Liang, "A multiobjective particle swarm optimizer using ring topology for solving multimodal multiobjective problems," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 5, pp. 805–817, 2017.

[42] A. Ahrari, S. Elsayed, R. Sarker, D. Essam, and C. A. Coello Coello, "Adaptive multilevel prediction method for dynamic multimodal optimization," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 3, pp. 463–477, 2021.

[43] A. Ahrari, S. Elsayed, R. Sarker, D. Essam, and C. A. C. Coello, "A novel parametric benchmark generator for dynamic multimodal optimization," *Swarm and Evolutionary Computation*, vol. 65, p. 100924, 2021.

[44] T. Hiroyasu, S. Nakayama, and M. Miki, "Comparison study of SPEA2+, SPEA2, and NSGA-II in diesel engine emissions and fuel economy problem," in *2005 IEEE congress on evolutionary computation*, vol. 1, pp. 236–242, IEEE, 2005.

[45] N. Hansen, S. Finck, R. Ros, and A. Auger, *Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions*. PhD thesis, INRIA, 2009.

[46] D. Brockhoff, A. Auger, N. Hansen, and T. Tušar, "Using well-understood single-objective functions in multiobjective black-box optimization test suites," *Evolutionary computation*, vol. 30, no. 2, pp. 165–193, 2022.

[47] S. Wessing, *The multiple peaks model 2*. Technische Universität Dortmund, Faculty of Computer Science, Algorithm Engineering (Ls 11), 2015.

[48] L. Schäpermeier, P. Kerschke, C. Grimme, and H. Trautmann, "Peak-A-Boo! Generating Multi-objective Multiple Peaks Benchmark Problems with Precise Pareto Sets," in *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 291–304, Springer, 2023.

[49] J. E. Fieldsend, "Using an adaptive collection of local evolutionary algorithms for multi-modal problems," *Soft Computing*, vol. 19, 06 2014.

[50] J. E. Fieldsend and K. Alyahya, "Visualising the landscape of multi-objective problems using local optima networks," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1421–1429, 2019.

[51] L. Schäpermeier, C. Grimme, and P. Kerschke, "One PLOT to Show Them All: Visualization of Efficient Sets in Multi-objective Landscapes," in *International Conference on Parallel Problem Solving from Nature*, pp. 154–167, Springer, 2020.

[52] P. Kerschke, M. Preuss, C. Hernández, O. Schütze, J.-Q. Sun, C. Grimme, G. Rudolph, B. Bischl, and H. Trautmann, "Cell Mapping Techniques for Exploratory Landscape Analysis," in *EVOLVE - A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation V* (Tantar, A. A. et al., ed.), (Cham), pp. 115–131, Springer International Publishing, 2014.

[53] L. Schäpermeier, C. Grimme, and P. Kerschke, "To Boldly Show What No One Has Seen Before: A Dashboard for Visualizing Multi-objective Landscapes," in *Evolutionary Multi-Criterion Optimization: 11th International Conference, EMO 2021*, pp. 632–644, Springer, 2021.

[54] A. Liefooghe, S. Verel, B. Lacroix, A.-C. Zăvoianu, and J. McCall, "Landscape features and automated algorithm selection for multi-objective interpolated continuous optimisation problems," in *Proceedings of the Genetic*

and Evolutionary Computation Conference, GECCO '21, p. 421–429, 2021.

[55] L. Schäpermeier, C. Grimme, and P. Kerschke, "MOLE: Digging Tunnels Through Multimodal Multi-objective Landscapes," in *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 592–600, 2022.

[56] L. Schäpermeier, C. Grimme, and P. Kerschke, "Plotting Impossible? Surveying Visualization Methods for Continuous Multi-objective Benchmark Problems," *IEEE Transactions on Evolutionary Computation*, vol. 26, no. 6, pp. 1306–1320, 2022.

[57] C. I. Hernández, O. Schütze, J.-Q. Sun, and S. Ober-Blöbaum, "Non-epsilon dominated evolutionary algorithm for the set of approximate solutions," *Mathematical and Computational Applications*, vol. 25, no. 1, 2020.

[58] M. Preuss and S. Wessing, "Measuring multimodal optimization solution sets with a view to multiobjective techniques," in *EVOLVE-A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation IV*, pp. 123–137, Springer, 2013.

[59] G. Rudolph, B. Naujoks, and M. Preuss, "Capabilities of EMOA to detect and preserve equivalent Pareto subsets," in *EMO '07: Proceedings of the Evolutionary Multi-Criterion Optimization Conference*, pp. 36–50, 2007.

[60] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise.," in *KDD* (E. Simoudis, J. Han, and U. M. Fayyad, eds.), pp. 226–231, AAAI Press, 1996.

[61] S. Ben-David and M. Ackerman, "Measures of clustering quality: A working set of axioms for clustering," in *Advances in Neural Information Processing Systems* (D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds.), vol. 21, Curran Associates, Inc., 2008.

[62] C. Yue, B. Qu, K. Yu, J. Liang, and X. Li, "A Novel Scalable Test Problem Suite for Multimodal Multiobjective Optimization," *Swarm and Evolutionary Computation*, vol. 48, pp. 62 – 71, 2019.

[63] A. Martín and O. Schütze, "Pareto tracer: A predictor-corrector method for multi-objective optimization problems," *Engineering Optimization*, vol. 50, no. 3, pp. 516–536, 2018.

[64] H. Ishibuchi, T. Matsumoto, N. Masuyama, and Y. Nojima, "Effects of dominance resistant solutions on the performance of evolutionary multi-objective and many-objective algorithms," in *GECCO '20: Proceedings of the 22th annual Conference on Genetic and Evolutionary Computation*, pp. 507 – 515, 2020.

[65] J.-Q. Sun, F.-R. Xiong, O. Schütze, and C. Hernández, *Cell Mapping Methods - Algorithmic Approaches and Applications*. Springer, 2019.

[66] F. Ming, W. Gong, L. Wang, and L. Gao, "Balancing convergence and diversity in objective and decision spaces for multimodal multi-objective optimization," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 7, no. 2, pp. 474–486, 2022.

[67] J. J. Liang, C. Yue, and B.-Y. Qu, "Multimodal multi-objective optimization: A preliminary study," in *2016 IEEE Congress on Evolutionary Computation (CEC)*, pp. 2454–2461, IEEE, 2016.

[68] Z. Ding, L. Cao, L. Chen, D. Sun, X. Zhang, and Z. Tao, "Large-scale multimodal multiobjective evolutionary optimization based on hybrid hierarchical clustering," *Knowledge-Based Systems*, vol. 266, p. 110398, 2023.

[69] W. Li, T. Zhang, R. Wang, and H. Ishibuchi, "Weighted indicator-based evolutionary algorithm for multimodal multiobjective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 6, pp. 1064–1078, 2021.

[70] Y. Tian, R. Liu, X. Zhang, H. Ma, K. C. Tan, and Y. Jin, "A multipopulation evolutionary algorithm for solving large-scale multimodal multiobjective optimization problems," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 3, pp. 405–418, 2020.

[71] Y. Liu, G. G. Yen, and D. Gong, "A multimodal multiobjective evolutionary algorithm using two-archive and recombination strategies," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 4, pp. 660–674, 2018.

[72] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, *Scalable Test Problems for Evolutionary Multiobjective Optimization*, pp. 105–145. London: Springer London, 2005.

[73] S. Huband, P. Hingston, L. Barone, and L. While, "A review of multiobjective test problems and a scalable test problem toolkit," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 5, pp. 477–506, 2006.

[74] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*. Chichester, UK: John Wiley & Sons, 2001. ISBN 0-471-87339-X.

[75] S. Schaeffler, R. Schultz, and K. Weinzierl, "Stochastic method for the solution of unconstrained vector optimization problems," *Journal of Optimization Theory and Applications*, vol. 114, no. 1, pp. 209–222, 2002.

[76] M. Emmerich and A. Deutz, "Test problems based on Lamé superspheres," in *Evolutionary Multi-Criterion Optimization: 4th International Conference, EMO 2007, Matsushima, Japan, March 5-8, 2007. Proceedings* (S. O. et al., ed.), (Berlin, Heidelberg), pp. 922–936, Springer Berlin Heidelberg, 2007.

[77] K. Deb and S. Tiwari, "Omni-optimizer: A generic evolutionary algorithm for single and multi-objective optimization," *European Journal of Operational Research*, vol. 185, no. 3, pp. 1062 – 1087, 2008.

[78] Y. Tian, R. Cheng, X. Zhang, and Y. Jin, "PlatEMO: A MATLAB platform for evolutionary multi-objective optimization," *IEEE Computational Intelligence Magazine*, vol. 12, no. 4, pp. 73–87, 2017.

[79] J. Demšar, "Statistical comparisons of classifiers over multiple data sets," *The Journal of Machine learning research*, vol. 7, pp. 1–30, 2006.