

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Aghaei Pour, Pouya; Bandaru, Sunith; Afsar, Bekir; Miettinen, Kaisa

Title: Desirable properties of performance indicators for assessing interactive evolutionary multiobjective optimization methods

Year: 2022

Version: Accepted version (Final draft)

Copyright: © 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM

Rights: In Copyright

Rights url: <http://rightsstatements.org/page/InC/1.0/?language=en>

Please cite the original version:

Aghaei Pour, P., Bandaru, S., Afsar, B., & Miettinen, K. (2022). Desirable properties of performance indicators for assessing interactive evolutionary multiobjective optimization methods. In J. E. Fieldsend (Ed.), *GECCO '22 : Proceedings of the Genetic and Evolutionary Computation Conference Companion* (pp. 1803-1811). ACM.
<https://doi.org/10.1145/3520304.3533955>

Desirable Properties of Performance Indicators for Assessing Interactive Evolutionary Multiobjective Optimization Methods

Pouya Aghaei Pour¹, Sunith Bandaru², Bekir Afsar¹, and Kaisa Miettinen¹

¹the University of Jyväskylä, Faculty of Information Technology
²School of Engineering Science, University of Skövde

Abstract

Interactive methods support decision makers in finding the most preferred solution in multiobjective optimization problems. They iteratively incorporate the decision maker's preference information to find the best balance among conflicting objectives. Several interactive methods have been developed in the literature. However, choosing the most suitable interactive method for a given problem can prove challenging and appropriate indicators are needed to compare interactive methods. Some indicators exist for a priori methods, where preferences are provided at the beginning of the solution process. We present some numerical experiments that illustrate why these indicators are not suitable for interactive methods. As the main contribution of this paper, we propose a set of desirable properties of indicators for assessing interactive methods as the first step of filling a gap in the literature. We discuss each property in detail and provide simple examples to illustrate their behavior.

Multiple criteria optimization, Performance evaluation, Performance assessment, Interactive methods

1 Introduction

In multiobjective optimization problems, we usually optimize several conflicting objectives simultaneously. This leads to multiple optimal solutions (known as *Pareto optimal solutions*) that are mathematically incomparable [1]. The set of Pareto optimal solutions is referred to as the Pareto front in the objective space.

Multiobjective evolutionary algorithms (MOEAs) are well-known methods for solving multiobjective optimization problems due to their ability to provide an approximation of the Pareto front. In addition, they can handle problems

without analytical functions, different types of decision variables, and so on [2, 3]. On the other hand, they cannot guarantee Pareto optimality, but generate an approximation of the Pareto front. However, in most real-world problems, only a single Pareto optimal solution needs to be selected for implementation. Typically, we use the knowledge of a domain expert, also known as a *decision maker* (DM), to provide some kind of preference information. Then, based on the DM's preference information, the most preferred solution is selected.

We can incorporate the DM's preference information for MOEAs in three main ways [1, 4]: 1) *a posteriori* methods, where the DM first sees an approximation of the Pareto front, and then chooses one or more solutions based on her/his preferences, 2) *a priori* methods, where the DM provides the preference information before the solution process, and then, a suitable MOEA tries to generate solutions that reflect the DM's preferences as well as possible and 3) *interactive* methods, where the DM provides her/his preferences iteratively during the solution process and guides the search to find one's most preferred solution in the approximated Pareto front.

A posteriori methods enable the DM to better understand existing trade-offs before expressing preferences. However, generating an approximation of the entire Pareto front is computationally expensive. In addition, it may be overwhelming for the DM to compare many solutions, especially, if we deal with a high number of objectives. *A priori* methods are usually computationally less expensive than *a posteriori* methods. However, it may be hard to provide preference information without knowing what kind of trade-offs are feasible. Moreover, besides preference information, most *a priori* methods require a parameter to identify a parameterized *region of interest* (ROI) which is a part of the approximated Pareto front that the DM is interested in. It is worth mentioning that according to [5], the definition of parameterized ROI is vague, and it can be identified in many different ways.

In interactive methods, the DM has the chance to learn about the trade-offs between objectives during the solution process and identify her/his most preferred solution in the ROI. The ROI is a part of the approximated Pareto front, where the DM likes to fine-tune the preferences and refine solutions. Moreover, unlike *a posteriori* methods, the DM has to process only a limited amount of information based on her/his preference information, which reduces the cognitive load set on her/him. There are different ways to provide preferences [6, 7]. For instance, in [8], the DM is able to provide his/her preferences in four different ways. Specifying aspiration levels representing desirable objective function values (constituting a so-called *reference point*) is a well-known way of providing preference information. The reference point is a popular way to provide preference information since it has been proven to be understandable to the DM [6, 9].

Many *performance indicators* (or indicators for simplicity) have been developed for *a posteriori* methods to be able to compare them [10, 11]. They assess the performance in approximating the whole Pareto front. In addition, some indicators have been dedicated to *a priori* methods [12, 13, 14, 15, 16]. They assess the performance in representing specific parts of the Pareto front identified

by preference information provided by a DM.

However, comparing interactive methods has been studied less. Typically, before the DM uses an interactive method, an analyst, who knows the behavior of interactive methods, should choose the most appropriate one. However, there are many aspects that the analyst should consider to be able to choose. To the best of our knowledge, no indicators have specifically been designed for assessing interactive methods.

As the first step towards developing indicators for interactive methods, we must identify desirable properties for such indicators. As the main contribution of this paper, we identify such desirable properties. It is important to note that a single indicator is unlikely to possess all desirable properties. In fact, we suggest that several indicators should be developed for assessing different aspects of interactive methods. To support our motivation, we show that the indicators designed for a priori methods are not suitable to assess interactive methods. However, we do not claim that the list of properties presented in this paper is exhaustive. Our objective is to initiate research in this direction.

In this paper, we first briefly review existing indicators in Section 2. Then in Section 3, we propose the desirable properties that indicators designed for interactive methods should possess, and describe each property in detail. In Section 4, we assess existing indicators against our proposal. Section 5 includes numerical experiments to support our arguments. Finally, we conclude the paper and mention future research directions in Section 6.

2 Background

Different indicators have been developed for assessing a priori methods. Their desirable properties are discussed in [16] and the indicators are stated to possess most of the desirable properties. However, some of the indicators require the knowledge of the Pareto front [17, 18] but, according to [16], an indicator should not rely on the knowledge of the Pareto front. In what follows, we briefly describe some recent indicators designed for a priori methods.

In *R-metric* [12], a reference point is incorporated to identify the parameterized ROI. Then, based on an achievement scalarizing function [9], one of the solutions is selected as a pivot point. Next, all solutions are transferred into a virtual position using the pivot point. Finally, the hypervolume [19], or the IGD [20] of the solutions inside the parameterized ROI is calculated as the assessment of an a priori method. In this paper, we use R-metric by calculating the hypervolume (we refer to it as R-HV) and higher values of R-HV represent better performance.

PMOD [14] is a distance-based indicator. The main idea is to map solutions onto a hyperplane generated based on the DM's reference point. Then, three different distances are calculated. First, the distance between the solutions and the reference point, second, the standard deviation of each mapped point to the nearest point (for measuring diversity), and third, the distance of each solution and the origin point is calculated, but if the solutions are outside of parameter-

ized ROI, this value is multiplied by a penalty coefficient. Finally, the PMOD value is calculated by using these three distances (for more details see [14]). For PMOD, lower values represent better performance both in convergence and diversity.

The *preference metric based on distances* [15] (PMDA) indicator is based on light beam search [21] and decomposition-based multiobjective evolutionary algorithm [22]. PMDA has four main steps. First, the reference point is decomposed into $k + 1$ light beams, where k is the number of objectives. Then a preference-based hyperplane is constructed by means of the light beams. Next, the Euclidean distances of solutions to the ideal point are calculated as the main assessment. Following this, angles between solutions outside the parameterized ROI and the reference point are calculated to form a penalty function by multiplying them by a constant coefficient. Finally, the mean of distances and angles for all solutions generate the PMDA assessment of a set of solutions. The lower the value of PMDA, the better it is.

In the *user-preference composite front* (UPCF) indicator [16], first, all the solution sets are merged. Then, all of the nondominated solutions are selected. Next, the closest solution to the reference point is identified, and a parameterized ROI is formed around it by acquiring a parameter that determines the size of the parameterized ROI. Finally, the hypervolume or IGD values for the solutions inside the parameterized ROI are calculated as the final assessment. In this paper, we use only the hypervolume version of UPCF (we refer to it as UPCF-HV). Higher values in UPCF-HV indicate better performance.

The *EH-metric* [13] is a parameterless indicator designed to eliminate the problem of defining parameterized ROI required by the indicators above. Instead of asking the user to define the size of the parameterized ROI through a parameter, this indicator uses the concept of an *expanding hypercube*, which starts as a point at the reference point and expands (with the reference point at its center) until it envelops all solutions. The EH-metric value for an a priori method is calculated as the area under the curve generated by plotting the fraction of solutions enveloped by the hypercube as it expands versus the size of the hypercube. The former is considered to be a measure of diversity around the reference point, while the latter is a measure of convergence to the reference point. Thus, higher EH-metric values indicate good convergence and diversity of preferred solutions.

As mentioned earlier, to the best of our knowledge, there are no indicators for assessing interactive methods. So far, researchers have applied indicators developed for a priori methods (with some adjustments) as the best viable option. For example, in [23, 24] the R-metric has been used in this way.

According to [7], we can often observe two phases in interactive solution processes: a *learning phase* and a *decision phase*. In both of these phases, the provided preferences direct the search to a *desired region*, where interactive methods try to generate solutions. In the learning phase, the DM studies different parts of the Pareto front to increase her/his knowledge about the problem, how well different preferences can be reflected, and learn more about the achievable values for objectives. At the end of the learning phase, the DM is more

confident about which part of the Pareto front she/he is interested in and has identified her/his ROI. Here, the DM enters the decision phase, where she/he fine-tunes the search within the ROI until she/he is satisfied with one of the solutions.

One should note that we use the concept ROI with different meanings in different contexts. The ROI in a priori indicators is based on the preferences that the DM provides before the optimization process, whereas in interactive methods, the ROI is identified at the end of the learning phase to be further studied in the decision phase. In addition, we refer here to the act of providing new preference information by the DM as an *interaction*. It happens after every method-specific number of generations.

3 Desirable Properties

In this section, first, we provide a list of desirable properties for designing indicators suitable for interactive methods. Then, we discuss and describe them in detail. Since these properties are meant to assess different aspects of interactive methods, a good starting point is the list of desirable properties identified for interactive methods, provided in [25]. In that study, the authors divided the desirable properties of interactive methods into three categories. The first category consists of properties that should be considered during the whole solution process, that is, both in the learning and decision phases. These properties are referred to as general properties (GPs). The second set of desirable properties, referred to as LPs, relates to the learning phase. In this phase, the method is supposed to assist the DM in studying the objective space and learning about the different trade-offs to identify a ROI. The third and final set of desirable properties, DPs, relates to the decision phase, where the interactive method is intended to assist the DM in identifying the most preferred solution in the ROI. For more details about the three phases, see [25].

In the same way, we divide desirable properties of indicators for interactive methods into the corresponding categories. Ideally, indicators for interactive methods must be able to:

- GP1: Assess the convergence of solutions in those regions of the approximated Pareto front that reflect the DM's preferences the best (local convergence).
- GP2: Assess the diversity of solutions in those regions of the approximated Pareto front that reflect the DM's preferences the best (local diversity).
- GP3: Assess the performance irrespective of the number of objective functions (scalability).
- GP4: Assess the performance without knowledge of the Pareto front.
- GP5: Assess the performance by incorporating preferences that are provided in different ways.

- GP6: Assess the performance in a computationally inexpensive manner.
- GP7: Assess the performance in a manner that is independent of other interactive methods being compared.
- GP8: Assess the performance without introducing parameters that have an unclear effect on the performance or are unintuitive to set.
- GP9: Assess the performance as a whole process and not as a series of independent a priori steps.
- LP1: Assess how much of the Pareto front has been studied (expedition).
- LP2: Assess how well/fast the method can adapt to new (even very different) preferences (responsiveness).
- DP1: Assess the capability of fine-tuning solutions inside the ROI.
- DP2: Assess the decision phase by considering the amount of information shown to the DM at each interaction.

Next, we discuss each desirable property in more detail. Moreover, for some of the desirable properties, we provide hypothetical examples illustrating their role in designing indicators for assessing interactive methods. We consider two hypothetical interactive methods, I_1 , and I_2 , and visualize their solutions as red rectangles and orange circles, respectively.

In the provided examples, we use a reference point \hat{z} as preference information. Furthermore, we assume that the DM begins with a learning phase, which is the case in many practical scenarios. After having identified an ROI, the DM moves to the decision phase.

Moreover, the desired region and ROI may be identified differently in each indicator. For simplicity, in the provided examples, we use a cone (green dashed lines) to represent the desired region. In addition, the ROI is represented by a purple box, where we expect the DM would provide her/his reference points in the decision phase. Actually, the ROI is a subset of Pareto optimal solutions.

Figure 1 illustrates an example of how reference points are typically provided in learning and decision phases to reflect different needs. Reference points in the learning phase (denoted by \dagger) are often scattered as the DM goes on an expedition to learn more about the Pareto front. On the other hand, reference points in the decision phase (denoted by \oplus) have some conformance among them as the DM refines solutions in the ROI identified at the end of the learning phase.

3.1 General Properties

We have nine general desirable properties that should be considered when designing new indicators. These general properties are valid for both learning and decision phases. According to [16], indicators for a priori methods should have four desired properties. We can extend these desirable properties to be applicable in the case of interactive methods. Thus, the first four general properties





Figure 1: Example of reference points in the learning phase (denoted by \dagger) and the decision phase (denoted by \oplus)

correspond to those in [16]. In addition, we have formulated five more general properties for designing indicators for interactive methods. These general properties do not depend on specific desires of the DM regarding the learning or decision phases.

GP1

Convergence following the preferences, which we refer to as *local convergence*, is an important desirable property because with each interaction, whether in the learning phase or the decision phase, the DM expects to see solutions that reflect the preferences. According to [25], this gives the DM the feeling of being in control of the solution process.



(a) Example for showing local convergence.

(b) Example for showing local diversity.

Figure 2: Solutions generated with two hypothetical interactive methods I_1 (red rectangles) and I_2 (orange circles) for the reference point \hat{z} (denoted by \dagger). The desired region is shown by a green dashed cone.

Figure 2a illustrates local convergence. Here, both interactive methods I_1 and I_2 have generated four solutions. If we consider the union of the solutions and eliminate dominated ones, we can observe that I_2 retains all of its solutions, while I_1 loses one. However, since the remaining solutions of I_1 reflect \hat{z} better, an indicator should be able to identify I_1 as a better method.

GP2

An indicator should be able to measure the diversity of solutions reflecting the DM's preference, which we refer to as *local diversity* [16]. This is important because, at each interaction, the DM must have “discernibly distinct” solutions to choose from. A good balance between local convergence and local diversity is required so that the solutions are not too diverse to make the DM feel that the preferences are not being reflected by the interactive method (c.f. GP1).

Figure 2b illustrates local diversity. It is clear that the solutions generated by I_2 are more diverse than those of I_1 . We can observe that all solutions of I_1 are in the desired region, while I_2 has generated a solution outside it. We assume that interactive methods show all these solutions to the DM. Therefore, solutions that are outside of the desired region should not be disregarded but should influence the indicator's assessment in a negative way to reflect differences between methods compared.

GP3

Scalability is a desirable property of indicators. For example, when the number of objectives grows, and we cannot even visualize the solutions properly, it will be imperative for the analyst to be able to rely on the indicator when comparing interactive methods.

GP4

If an indicator needs knowledge of the Pareto front, the applicability of the indicator is limited. This is important to keep in mind since the main purpose of interactive methods is solving real-world problems, where we do not know the Pareto front in most cases. Therefore, it is essential that indicators do not depend on this information.

GP5

As mentioned in Section 1, different interactive methods assume preference information to be provided in different ways. Therefore, appropriate indicators are needed. This does not mean that one indicator should be able handle all different ways of providing preferences.

GP6

In general, the calculation of indicator values should be computationally inexpensive. This enables their more versatile usage. For example, one may want to calculate them at regular intervals during the solution process to monitor the progress of an interactive method or compare progress of different methods. For example, if an indicator is based on an inherently expensive computation, such as the hypervolume, the computation time increases exponentially as the number of objectives grows. It is impractical to use such indicators often (e.g., in regular intervals).

GP7

It is desirable that the indicator value for a given interactive method be independent of the other interactive methods being compared. This avoids the problem of recomputing the indicator when a new interactive method needs to be included in the comparison.

GP8

An indicator should be easy to use, not having parameters whose effect on assessing the performance is unclear. For example, many interactive methods have a parameter that identifies the desired region based on the DM's preferences (see e.g., [8]). Here, if the indicator asks for a new parameter to redefine the desired region, the analyst can get confused since she/he has to provide this information twice in different ways.

GP9

In some studies, the performance of interactive methods has been assessed by considering each interaction as a distinct a priori step [23, 24, 26], and indicators for a priori methods have been used to assess the median performance of interactions of each phase. This allows the use of existing (a priori) indicators in the absence of those designed for interactive methods. However, this can mislead assessments since the solution process as a whole and different roles of learning and decision phases are not supported.

Figure 3 illustrates why interactive methods should not be assessed as a series of a priori steps. Figures 3a, 3b, and 3c show three interactions of the decision phase with the two interactive methods I_1 and I_2 in a biobjective minimization problem. It is shown that the solutions generated by I_1 in Figures 3a and 3b are better than those of I_2 . However, in the third interaction, I_2 manages to find a solution that the DM prefers (orange circle inside a black rectangle). Here, if we calculate the mean of the performances of I_1 and I_2 with most of the indicators developed for a priori methods, I_1 will have a better performance than I_2 . However, the most preferred solution was generated by I_2 . Therefore, it is important for indicators to consider the interactive methods as more than a series of a priori steps.




Figure 3: Three interactions of interactive methods I_1 and I_2 with the reference points $\hat{z}_1, \hat{z}_2, \hat{z}_3$. The desired region is shown by a green dashed cone, and the ROI is shown by a purple box.

Before moving to desirable properties specific for decision and learning phases, it is worth pointing out that whether the DM is in the learning phase or the decision phase is not always obvious. Ideally, the indicator should have a mechanism to detect this transition based on the sequence of DM's preferences and modify its calculations to either suit the learning phase or the decision phase. Alternatively, one can design separate indicators for the two phases.

3.2 Learning Phase

In addition to the general properties, indicators should have specific desirable properties for the learning phase. In this phase, the DM wants to study the objective space to finally identify her/his ROI. Therefore, when designing indicators, we should consider the unique characteristics of the learning phase.

LP1

Measuring the expedition of an interactive method in the learning phase can help the analyst to figure out whether the interactive method has covered the approximation of the Pareto front well enough. Typically in this phase, the DM is not aware of the shape of the Pareto front. So, it is difficult for the DM to say how much expedition she/he has done. This is particularly true in many-objective problems. Moreover, measuring the expedition does not need to be exact since we do not want to rely on the knowledge of the Pareto front (GP4) and it is enough if the indicator can identify different regions of the Pareto front and communicates this information to the analyst.

It is worth mentioning that expedition is not the same as local diversity. The solutions shown to the DM should be diverse within the desired region so that they still reflect the DM's preferences (GP2). On the other hand, expedition is more about an approximation of how much of the approximated Pareto front has been covered by the generated solutions through the learning phase.

LP2

In the learning phase, the DM is still studying the objective space, and therefore her/his preferences may change drastically. As mentioned in [25], responsiveness to these changes is a desirable property for interactive methods in the learning phase. Therefore, in the learning phase, it is desirable for indicators to assess how well an interactive method can adapt to the changes in the preferences.

Moreover, as mentioned in Section 2, one of the main advantages of interactive methods is that they do not need as many function evaluations as a posteriori methods. Besides the responsiveness of an interactive method, measuring how fast it can converge toward new preferences is essential as well because usually the DM has limited time to wait for new solutions to be generated. Therefore, it is important that the interactive method can respond to the new preferences as fast as possible to minimize the waiting time of the DM. For example, we can track how many function evaluations it takes to generate solutions that reflect the new preferences.




Figure 4: Two interactions of methods I_1 and I_2 . Here, I_2 responds better to the change of preferences (from \hat{z}_1 to \hat{z}_2).

Figure 4 illustrates an example of the importance of responsiveness. Here, we assume that both interactive methods I_1 and I_2 had the same budget for function evaluations for each interaction. We can observe that in the first interaction (Figure 4a) with the reference point \hat{z}_1 , both interactive methods had almost similar results. However, when the DM provided the second reference point (Figure 4b), \hat{z}_2 , the solutions generated by I_2 are closer to \hat{z}_2 than what I_1 has generated. In other words, we could say that I_2 is more responsive to the changes in reference points than I_1 .

3.3 Decision Phase

As the DM begins the decision phase, new desirable properties dedicated to this phase should be considered. Now, she/he is more interested in fine-tuning

solutions in the ROI identified at the end of the learning phase.

DP1

Since the DM refines solutions by providing her/his preferences within the ROI, preferences are likely to be concordant.

However, it is not easy to confirm whether the provided preferences have concordance with each other in problems with more objectives. For example, when we are using reference points in the decision phase, if the new reference point at each interaction in the ROI dominates the old reference point (which is also in the ROI), they are concordant. Otherwise, even if the new reference point dominates the previous one, they are not concordant if the new one is provided outside of the ROI. An indicator must identify the preferences with concordance and increase the role of corresponding solutions in assessing the decision phase.




Figure 5: Hypothetical fine-tuning in the decision phase during four interactions. The ROI is represented by a purple box.

Figure 5 shows a simple example of the fine-tuning capability of an interactive method. Here, \hat{z}_1 , \hat{z}_2 , \hat{z}_3 and \hat{z}_4 are the reference points that the DM has provided in successive decision phase interactions. Assume that the most preferred solution is chosen in the last interaction (from solutions corresponding to \hat{z}_4). In this example, we can easily observe that \hat{z}_1 , \hat{z}_2 and \hat{z}_4 are all inside the ROI (denoted by a purple box), but \hat{z}_3 is not. Here, \hat{z}_4 dominates \hat{z}_1 and therefore, they are concordant. However, there is no concordance between \hat{z}_4 , \hat{z}_2 , and \hat{z}_3 because \hat{z}_2 is not dominated by \hat{z}_4 , and \hat{z}_3 is outside the ROI.

Moreover, since the most preferred solution is chosen from the fourth interaction, \hat{z}_4 must play the most significant role in assessing the decision phase. An indicator should be able to measure the concordance between the DM's preferences in the decision phase. Here, each interaction does not influence the assessment equally, Thus, this approach is different from seeing interactive methods as a series of a priori steps where all interactions have an equal effect on the assessment of the performance.

Table 1: Indicator values for iRVEA and iNSGA for 3-objective DTLZ7 problem in the learning phase. Here, \uparrow means that higher values are better for the corresponding indicator, and \downarrow means that lower values are better. Bold values indicate that the corresponding interactive method has a better performance.

	iRVEA					iNSGA				
	R-HV \uparrow	EH-metric \uparrow	PMOD \downarrow	PMDA \downarrow	UPCF-HV \uparrow	R-HV \uparrow	EH-metric \uparrow	PMOD \downarrow	PMDA \downarrow	UPCF-HV \uparrow
interaction 1	5.927	0.276	6.543	4.632	3.437	6.669	0.561	7.888	5.050	4.310
interaction 2	5.859	0.569	5.713	4.248	3.985	5.401	0.419	6.626	5.092	2.546
interaction 3	7.301	0.688	5.418	3.226	4.015	6.489	0.121	5.833	4.940	2.135
interaction 4	5.909	0.482	5.690	4.245	3.874	2.930	0.196	6.392	4.893	2.497

DP2

Another important aspect of interactive methods is the amount of information shown to the DM at each interaction. In relation to GP5, the amount and nature of information depend on the way preference information is provided. In the learning phase, the amount of information required by the DM to learn about the shape of the Pareto front can vary as long as the cognitive load is acceptable to the DM. However, in the decision phase, a typical requirement from a DM is the number of solutions that she/he wishes to analyze within the ROI [1]. An interactive method that generates fewer solutions than what the DM desires may delay the solution process, whereas one that generates more solutions may increase the cognitive load on the DM. In essence, the desires of the DM should be respected. An indicator should be able to take both these aspects into account.

4 Applicability of Existing Indicators

As mentioned earlier, since there are no indicators in the literature designed specifically for comparing interactive methods, some studies [23, 24, 27] have resorted to using indicators developed for a priori methods. In this section, we assess the five a priori indicators presented in Section 2 with respect to the desirable properties discussed in the previous section.

Table 2 shows how well the five indicators stack against the desirable properties. All of the indicators satisfy the first four desirable properties concerning local convergence and local diversity, scalability and knowledge of the Pareto front (GP1, GP2, GP3 and GP4). However, for GP1 and GP2, some of the indicators like R-HV and UPCF-HV remove solutions outside the desired region before calculating GP1 and GP2. This may be misleading since the DM sees these solutions. In other words, if the interactive method presents some solutions outside the desired region (or the ROI), then these solutions should have a negative effect in assessing the method (instead of being deleted).

All the indicators consider reference point(s) as the preference information. Hence, none of the indicators satisfy the desirable property GP5. Moreover, R-HV and UPCF-HV do not satisfy GP6, since they are based on computationally expensive calculations (hypervolume). Here, as the number of objectives increases, the computation time of these indicators grows exponentially, which

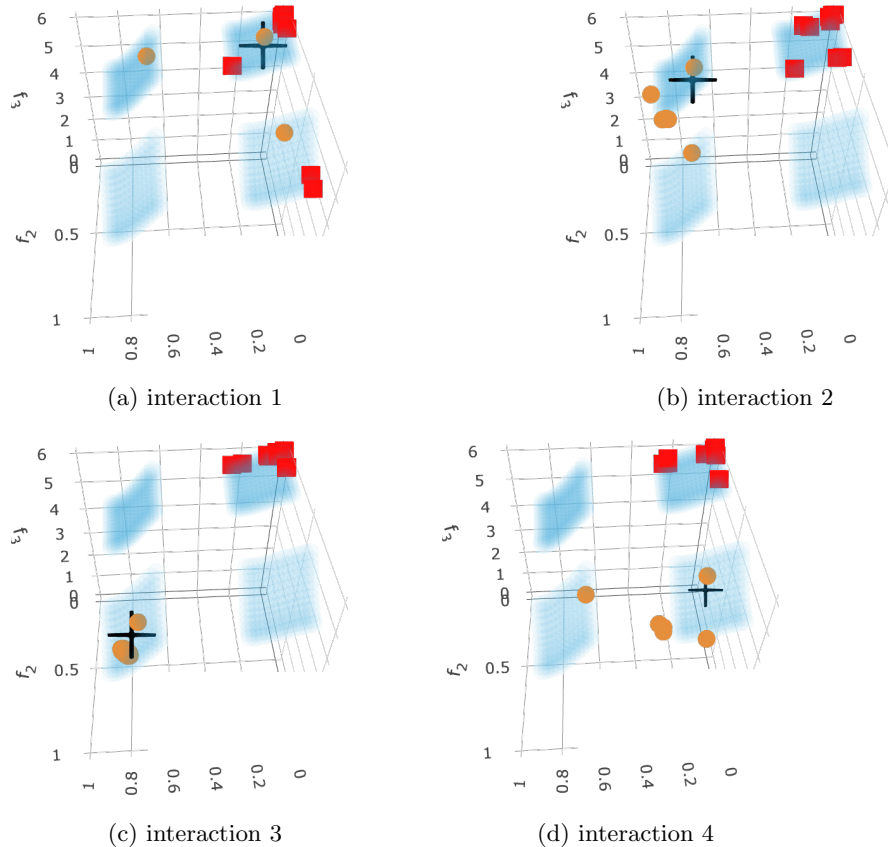


Figure 6: DTLZ7 in the learning phase where we provided a reference point in each part of the Pareto front manually.

is not desirable in interactive methods.

Among the indicators we listed in Table 2, R-HV, EH-metric and UPCF-HV do not satisfy GP7. They employ a prescreening step which combines and sorts the final solutions from all methods being compared. Thus, the values of these indicators depend on the methods being compared. If a new method is to be compared, these values may need to be recomputed (for more details see [16, 12, 13]).

Except for EH-metric, the rest of the indicators need an analyst to set at least one parameter. For example, R-HV, PMOD, PMDA and UPCF-H require the size of parameterized ROI. In addition, PMOD and PMDA require a penalty coefficient for the solutions outside the parameterized ROI. It may be confusing for the analyst to provide these parameters values. None of the studies have analyzed the effect of these parameters on their assessment of the performance.

The rest of the desirable properties are understandably not satisfied by any

Table 2: Proposed desirable properties and their presence in existing indicators for a priori methods.

Properties	R-HV	EH-metric	UPCF-HV	PMOD	PMDA
GP1	✓	✓	✓	✓	✓
GP2	✓	✓	✓	✓	✓
GP3	✓	✓	✓	✓	✓
GP4	✓	✓	✓	✓	✓
GP5	✗	✗	✗	✗	✗
GP6	✗	✓	✗	✓	✓
GP7	✗	✗	✗	✓	✓
GP8	✗	✓	✗	✗	✗
GP9	✗	✗	✗	✗	✗
LP1	✗	✗	✗	✗	✗
LP2	✗	✗	✗	✗	✗
DP1	✗	✗	✗	✗	✗
DP2	✗	✗	✗	✗	✗

Table 3: Mean indicator values of iRVEA and iNSGA for 5-objective DTLZ3 problem in the decision phase. As before, \uparrow means that higher values are better for the corresponding indicator, and \downarrow means that lower values are better. Bold values indicate that the corresponding interactive method has a better performance.

	iRVEA					iNSGA				
	R-HV \uparrow	EH-metric \uparrow	PMOD \downarrow	PMDA \downarrow	UPCF-HV \uparrow	R-HV \uparrow	EH-metric \uparrow	PMOD \downarrow	PMDA \downarrow	UPCF-HV \uparrow
interaction 1	14.145	0.475	5.835	2.888	0.185	32.212	0.590	8.533	3.329	0.195
interaction 2	24.263	0.619	6.405	1.430	0.203	32.253	0.643	9.265	3.329	0.262
interaction 3	26.126	0.766	13.813	0.704	0.311	31.956	0.744	12.160	3.329	0.347
interaction 4	28.699	0.768	14.950	0.687	0.412	31.957	0.741	12.160	3.329	0.458

of the indicators, because they were not designed to consider the learning and decision phases. We have already mentioned that there could be separate indicators dedicated to assess learning and decision phases. In addition, it may be even too difficult to design an indicator that satisfies all the desirable properties in the learning or decision phases. Hence, we may need several indicators for different purposes in each phase.

5 Numerical Experiments

In this section, we present two numerical examples to emphasize the need for developing new indicators specifically designed for interactive methods. We demonstrate that the current practice of assessing interactive methods as a series of a priori steps is inappropriate. The first example shows the importance of desirable properties for the learning phase (LP1 and LP2). The second example focuses on the importance of DP1 in the decision phase and GP9 of the general properties. We compare interactive NSGAIII (iNSGA) [23] and interactive RVEA (iRVEA) [8] using DTLZ benchmark problems [28]. For both examples, the number of generations is limited to 100 per interaction, while the number of function evaluations for both methods is 100000. These numbers are examples

as our goal here is not to find the best method but to study the behavior of the indicators.

Since using a priori indicators for comparing interactive methods is not the main focus of this study, we have provided the details of this part in the Supplementary Material.¹

5.1 An Example for the Learning Phase

Here, we use the 3-objective 11-variable (number of variables is based on [29]) DTLZ7 problem due to the unique shape of its Pareto front, which has four disconnected regions as shown by the blue areas in Figure 6. These distinct regions enable demonstrating the importance of measuring the expedition (LP1) and responsiveness (LP2) in the learning phase. The following reference points corresponding to each region were used in successive interactions to test the expedition capability of iRVEA and iNSGA (a) [0.11, 0.10, 5.4], (b) [0.70, 0.14, 4.50], (c) [0.76, 0.76, 3.5], and (d) [0.14, 0.70, 4.5].

Figure 6 shows the solutions that iNSGA (red rectangles) and iRVEA (orange circles) have generated corresponding to different reference points. We can observe that iNSGA could not respond to the changes of the reference point and stayed in one region. On the other hand, iRVEA could provide solutions in the same region with the reference point. Therefore, for expedition (LP1) and responsiveness (LP2), iRVEA was better than iNSGA in this example.

The assessments of indicators for a priori methods for each interaction have been gathered in Table 1. We can observe that most of the indicators declare that iRVEA is better than iNSGA in the last three interactions. However, based on these values, we cannot get the essential information that iNSGA was stuck in the initial region and could not cover the Pareto front well. In fact, by looking at these values, the analyst may be misled to think that the performance of these interactive methods is not that different (e.g., see PMOD or PMDA values). Therefore, it is essential that the indicators can communicate such important insight to the analyst.

5.2 An Example for the Decision Phase

In this section, we compare iRVEA and iNSGA using the 5-objective DTLZ3 problem in the decision phase. Thus, we assume that the ROI has already been identified. To generate reference points, we used an artificial decision maker [23]. These reference points are: (a) [0.000, 0.000, 0.000, 3.072], (b) [0.000, 0.000, 0.000, 0.000, 1.951], (c) [0.000, 0.000, 0.000, 0.000, 1.010], and (d) [0.000, 0.000, 0.000, 0.000, 1.010]. We can observe that the artificial decision maker changed the reference point for the first three interactions, but at the last interaction used the same reference point as the previous one. After generating the reference points, we ran each interactive method ten independent times. Then, we calculated the average of results for each interaction (see Table 3).

¹Link to the implementation: <https://github.com/ppouyaa/desirable-properties-master>
<https://github.com/ppouyaa/desirable-properties-master>

According to Table 3, EH-metric indicates that for the first two interactions, iNSGA was better than iRVEA, while for the third and fourth interactions, iRVEA was better than iNSGA. Earlier, we mentioned that in the current literature, the mean of indicators values for each interaction is typically calculated to find the best interactive methods. Here, if we calculate the mean of the EH-metric values, iRVEA has the value of 0.656, and iNSGA has the value of 0.679. Therefore, based on this way of calculation, iNSGA is better than iRVEA.

However, earlier, we noted that one of the desirable properties in the decision phase is fine-tuning solutions in the ROI (DP1). This involves information about the concordance of reference points. Moreover, the reference points did not change in the last two interactions (where iRVEA had a better performance). If we consider the concordance of reference points and let solutions corresponding to the third and fourth interaction influence the results more, iRVEA could probably be regarded to have a better performance than iNSGA. This shows that it is important to have specifically designed indicators for the decision phase (or the learning phase). Besides, this example shows why it is important not to assess the interactive methods as a series of a priori steps (GP9).

Finally, we can observe in Table 3 that the indicators are not similar for each interaction. For example, based on R-HV, iNSGA was better than iRVEA at every interaction. However, based on PMDA, iRVEA was better than iNSGA at every interaction. In addition, EH-metric indicates that at the first two interactions, iNSGA was better, and for the third and fourth interactions, iRVEA was better. On the other hand, PMOD gave opposite results to EH-metric. This is interesting since most of these indicators were designed to calculate local convergence and local diversity, and still, the results are so different from one another. Thus, if these indicators are used to assess interactive methods, the results may be sensitive to the choice of the indicator. This supports the need of indicators designed specifically for assessing interactive methods.

6 Conclusions

In this paper, we identified the desirable properties for designing indicators suitable for interactive methods and discussed them in detail. There are three main categories for these desirable properties. The general properties that should be considered in both the learning and decision phase, the desirable properties regarding the learning phase, and the decision phase properties that focus on aspects of interactive methods that help the DM refine a solution. Together, we suggested 13 desirable properties that indicators designed for interactive methods should possess. However, one indicator cannot satisfy all the desirable properties, and there should be different indicators for different purposes.

We also demonstrated why indicators developed for a priori methods should not be applied for interactive methods. We showed that these indicators do not satisfy most of the desirable properties that we presented. Furthermore, we provided two numerical examples to support the claim that there is a need for indicators specifically designed for assessing interactive methods.

By assessing interactive methods, we can analyze their characteristics and choose the appropriate method for different real-world problems. Therefore, as a future research direction, we plan to develop indicators that satisfy at least most of the desirable properties we presented in this paper. Moreover, these desirable properties only consider algorithmic aspects of interactive methods. It is also important to study interactive methods from human perspectives such as cognitive load set on the DM.

Acknowledgments

This research was partly supported by the Academy of Finland (grant no 311877 and 322221) and is related to the thematic research area DEMO (Decision Analytics utilizing Causal Models and Multiobjective Optimization, jyu.fi/demo) of the University of Jyväskylä. The authors like to thank Dr. Michael Emmerich for providing valuable comments.

References

- [1] K. Miettinen, *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, 1999.
- [2] K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, 2001.
- [3] Y. Jin, H. Wang, and S. Chaoli., *Data-Driven Evolutionary Optimization*. Springer, 2021.
- [4] C.-L. Hwang and A. S. M. Masud, *Multiple Objective Decision Making Methods and Applications: a State-of-the-Art Survey*. Springer, 1979.
- [5] K. Li, R. Chen, G. Min, and X. Yao, “Integration of preferences in decomposition multiobjective optimization,” *IEEE Transactions on Cybernetics*, vol. 48, no. 12, pp. 3359–3370, 2018.
- [6] S. Bechikh, M. Kessentini, L. Ben Said, and K. Ghédira, “Chapter four - preference incorporation in evolutionary multiobjective optimization: a survey of the state-of-the-art,” in *Advances in Computers* (A. R. Hurson, ed.), vol. 98, pp. 141–207, Elsevier, 2015.
- [7] K. Miettinen, F. Ruiz, and A. Wierzbicki, “Introduction to multiobjective optimization: Interactive approaches,” in *Multiobjective Optimization: Interactive and Evolutionary Approaches* (J. Branke, K. Deb, K. Miettinen, and R. Slowinski, eds.), pp. 27–57, Berlin, Heidelberg: Springer, 2008.
- [8] J. Hakanen, T. Chugh, K. Sindhya, Y. Jin, and K. Miettinen, “Connections of reference vectors and different types of preference information in interactive multiobjective evolutionary algorithms,” in *Proceedings of the 2016*

- IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1–8, 2016.
- [9] A. P. Wierzbicki, “The use of reference objectives in multiobjective optimization,” in *Multiple Criteria Decision Making Theory and Application* (G. Fandel and T. Gal, eds.), pp. 468–486, Springer, 1980.
 - [10] N. Riquelme, C. Von Lüken, and B. Baran, “Performance metrics in multiobjective optimization,” in *2015 Latin American Computing Conference, Proceedings*, pp. 1–11, IEEE, 2015.
 - [11] C. Audet, J. Bignon, D. Cartier, S. Le Digabel, and L. Salomon, “Performance indicators in multiobjective optimization,” *European Journal of Operational Research*, vol. 292, no. 2, pp. 397–422, 2021.
 - [12] K. Li, K. Deb, and X. Yao, “R-metric: Evaluating the performance of preference-based evolutionary multiobjective optimization using reference points,” *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 6, pp. 821–835, 2017.
 - [13] S. Bandaru and H. Smedberg, “A parameterless performance metric for reference-point based multi-objective evolutionary algorithms,” in *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 499–506, ACM, 2019.
 - [14] Z. Hou, S. Yang, J. Zou, J. Zheng, G. Yu, and G. Ruan, “A performance indicator for reference-point-based multiobjective evolutionary optimization,” in *2018 IEEE Symposium Series on Computational Intelligence (SSCI), Proceedings*, pp. 1571–1578, IEEE, 2018.
 - [15] G. Yu, J. Zheng, and X. Li, “An improved performance metric for multiobjective evolutionary algorithms with user preferences,” in *2015 IEEE Congress on Evolutionary Computation, Proceedings*, pp. 908–915, IEEE, 2015.
 - [16] A. Mohammadi, M. N. Omidvar, and X. Li, “A new performance metric for user-preference based multi-objective evolutionary algorithms,” in *2013 IEEE Congress on Evolutionary Computation, Proceedings*, pp. 2825–2832, IEEE, 2013.
 - [17] D. A. Van Veldhuizen, *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. Air Force Institute of Technology, 1999.
 - [18] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. Da Fonseca, “Performance assessment of multiobjective optimizers: An analysis and review,” *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 117–132, 2003.

- [19] 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.
- [20] P. A. Bosman and D. Thierens, “The balance between proximity and diversity in multiobjective evolutionary algorithms,” *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 174–188, 2003.
- [21] A. Jaskiewicz and R. Słowiński, “The ‘light beam search’ approach—an overview of methodology applications,” *European Journal of Operational Research*, vol. 113, no. 2, pp. 300–314, 1999.
- [22] Q. Zhang and H. Li, “MOEA/D: A multiobjective evolutionary algorithm based on decomposition,” *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.
- [23] B. Afsar, A. Ruiz, and K. Miettinen, “Comparing interactive evolutionary multiobjective optimization methods with an artificial decision maker,” *Complex & Intelligent Systems*, 2021. to appear.
- [24] B. Afsar, K. Miettinen, and A. Ruiz, “An artificial decision maker for comparing reference point based interactive evolutionary multiobjective optimization methods,” in *Evolutionary Multi-Criterion Optimization, 11th International Conference, EMO 2021, Proceedings* (H. Ishibuchi, Q. Zhang, R. Cheng, K. Li, H. Li, H. Wang, and A. Zhou, eds.), (Cham), pp. 619–631, Springer, 2021.
- [25] B. Afsar, K. Miettinen, and F. Ruiz, “Assessing the performance of interactive multiobjective optimization methods: A survey,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–27, 2021.
- [26] A. Mazumdar, *Novel Approaches for Offline Data-Driven Evolutionary Multiobjective Optimization*. JYU Dissertations 456, University of Jyväskylä, 2021.
- [27] P. Aghaei Pour, T. Rodemann, J. Hakanen, and K. Miettinen, “Surrogate assisted interactive multiobjective optimization in energy system design of buildings,” *Optimization and Engineering*, vol. 23, pp. 303–327, 2022.
- [28] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, “Scalable multi-objective optimization test problems,” in *Proceedings of the 2002 Congress on Evolutionary Computation, CEC’02*, vol. 1, pp. 825–830, IEEE, 2002.
- [29] J. Knowles, “ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems,” *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 50–66, 2006.