On the Diversity of Non–dominated Sets.

Giovanni Lizárraga Lizárraga[1,a], Salvador Botello Rionda[2,b]

Corporación Mexicana de Investigación en Materiales S.A. de C.V. Ciencia y Tecnología 790, Saltillo México[1]

Center of Research in Mathematics. Jalisco S/N, Valenciana, Guanajuato, México[2]

glizarraga@comimsa.com[a], botello@cimat.mx[b]

Abstract. When evaluating the performance of Multiobjective Evolutionary Algorithms (MOEA), the most important factors to consider are the convergence and the diversity of the non–dominated sets generated. Even when convergence is the most important property, diversity has been considered very relevant when approximating the Pareto Front. Unfortunately, diversity is very difficult to define and evaluate, to the point that some studies on performance measures simply ignore diversity. In this work we explore the concept of diversity with the goal to find a better interpretation of that property. Also, we propose a model that can be helpful to evaluate the diversity of a subset of the Pareto Front.

1 Introduction

The performance of a multi–objective evolutionary algorithm (MOEA), is very difficult to evaluate. The output of a MOEA is a set of vectors in R^m, usually known as a non–dominated set (NS). In order to evaluate the performance of MOEAs, we need to evaluate the quality of non–dominated sets. One of the most important properties of a non–dominated set is diversity. Diversity is related with how diverse is the information presented by the set, and it is very difficult to evaluate. The goal of this work is to study the concept of diversity in order to clarify this concept. A better understanding of diversity allow us to define more accurate performance measures. When evaluating and comparing the performance of MOEAs, there are many factors that are important. In this paper we focus on the diversity.

The rest of the article is organized as follows: in Section 2 we introduce some concepts related to multi–objective optimization. In Section 3 we discuss the concept of diversity of an approximation set. In Section 4 we describe some of models of diversity. Finally, we state our conclusions in Section 5.

2 Basic Concepts

In Multi–objective optimization (MOO), we deal with problems whose statement is in the following way:

$$Minimize\ F(x) = \langle f_1(x), \ldots, f_m(x) \rangle \quad (1)$$

$$subject\ to:\ x \in X \subset R^n \quad (2)$$

X is the set of all feasible solutions to the problem. In MOO it is not clear "a priori" how to identify the the best solution. We assume that there is a trade off between the objective functions that gives the highest benefit to the end user, but this trade off is unknown. In order to identify a unique optimal solution, we need the help of a decision maker (DM), the person who has the final word when choosing a solution.

In this work we are interested in an "a posteriori" approach, where the DM preferences are not clear. In this approach we present the DM with a set of candidate solutions, with the hope that when the DM analyze some attainable solutions, he/she will clarify his/her preferences. Note that the criteria of the DM will depend on the set of solutions we present to him/her. So, it is important find not any random set of candidate solutions, but the most promising and efficient solutions we can find. But first, we must use a criterion to decide what is a promising solution.

One of the most used criteria to compare vectors is the Pareto Optimality Criteria (POC). POC is defined through a binary relation between vectors known as dominance. For two vectors $a, b \in R^m$, a dominates b, denoted $a \succ b$, if $\forall i \in \{1, 2, \ldots, m\},\ a_i \leq b_i\ \wedge\ \exists j \in \{1, 2, \ldots, n\}\ \mid a_j < b_j$, a_i stands for the i–th element of a. In the case where a does not dominate b and b does not dominate a, we say that a and b are not comparable. In the context of a multi–objective problem, we say that candidate solution x dominates candidates solution y if $F(x) \succ F(y)$. Using the POC, we can discard dominated solutions, because they are not optimal. A Pareto Optimal solution is a feasible solution that is not dominated by any other feasible solution. The set of all Pareto Optimal solutions is known as Pareto Set. The image of the Pareto Set under the vector of objective functions, $F(x)$, is known as Pareto Front (PF). The image of X under $F(x)$ is known as Z. It is common to call "the objective functions space" to R^m, the space where we locate the image of the candidate solutions under $F(x)$.

The quality of a candidate solution depends on its values of the objective functions, so in the rest of this work we base our discussion in the image of the candidate solutions in objective function space.

We want to present the DM with as much information as possible of the PF, because this set contains the most promising solutions. The Pareto Front may have many topologies, it can be discrete, continuous, finite, etc. When the PF is finite and small, it is possible to find all its elements. In this work we consider that the PF is too big to find all its elements. We can collect, at most, a sample of it.

There are many methodologies to collect samples of the Pareto Front. In this work we are interested in a special set of algorithms known as Multi–Objective Evolutionary Algorithms (MOEAs). MOEAs are inspired in the theory of evolution, with concepts like "population", "natural selection", "mutations", etc. For more detail about MOEAs, we recommend to read [2]. For the goals of this work, we only need to know that MOEAs are optimization algorithms that work with multi–objective problems and their output is an approximation of the Pareto Front. Also, MOEAs are stochastic algorithms, since for different runs they may return different approximations. In the rest of this work, we call "approximation set" or "approximation" to a set that tries to approximate the Pareto Front. As the Pareto Front consists of a set of mutually non comparable vectors, we consider that the elements of an approximation are mutually non comparable.

There are many MOEAs published, and a common question is how to evaluate whether a MOEA performs better than another one? There are many factors related with the answer of these questions, but one of the most important is the quality of the approximation they generate. This lead us to the question of what makes "better" an approximation? When an approximation is better than another one?

We consider that a decision maker may do two things with an approximation set: a) he/she will choose an element of the approximation as a final solution, or b) he/she will use the approximation to clarify its preferences in order to make a more specific search.

The most desirable characteristics of an approximation set are two: convergence and diversity. Convergence refers to how close is the approximation to the Pareto Front. Ideally, we want the approximation to be a subset of the Pareto Front. Diversity refers to how well distributed are the elements of the approximation among the Pareto Front. If the DM chooses a final solution from the approximation, a good distribution increases the chances to find a solution with a desired trade–off between objective functions. If the DM needs to clarify his/her preference, a good diversity gives a better idea of the shape of the Pareto Front and the different options of trade–off between objective functions.

In the rest of this paper we focus on diversity. The diversity of an approximation has been very difficult to measure. There are several studies that apport a lot in the subject but still there is a lot of room for improvement. In the next section we present some ideas about how to clarify the concept of diversity in order to create a system to measure the diversity of an approximation set.

3 Diversity on the Pareto Front

Convergence is considered the most important property of an approximation, and diversity comes in second place. Another way to state it, is that we first want to reach the Pareto Front, but once we are there, we want to collect the most diverse information. In the rest of this discussion, we assume that the approximations we are evaluating converged to the Pareto Front, or in other words, they are subsets of the Pareto Front. Diversity has been difficult to measure, to the point that some studies about comparison methods for MOEAs have ignored this property. We are convinced that diversity must be considered and here are some of the reasons:

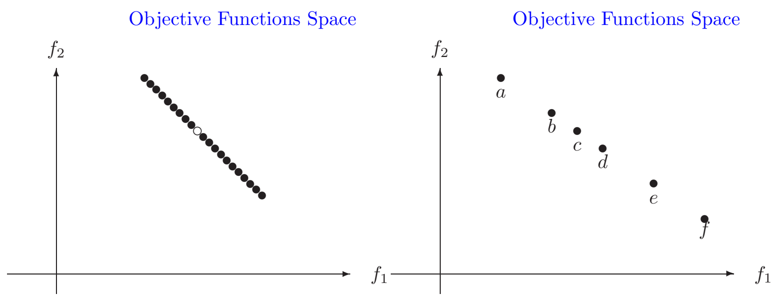

Fig. 1. Left panel, the approximation A ($\bullet$) has better diversity than approximation B ($\circ$). Right panel, an example of a Pareto Front.

- The final goal of evaluating the quality of an approximation set is to create a methodology to compare two (or more) MOEAs. Most MOEAs in literature use implicit or explicit mechanisms to generate approximation sets with good diversity. Ignoring diversity in a comparison method may result in unfair comparisons.
- The goal of presenting a set of solutions to the DM is to give him/her a good idea of the different possible trade off between objective functions. So, the DM can have a better idea of the attainable solutions and can make a decision based on more and better information. Consider the two NSs in Figure 1, left panel. A has lots of elements and B has only one, assume that both sets are in the Pareto Front and they are disjoint. Appealing to the intuitive notion of convergence, it can be argued that both sets have the same convergence, but they have a huge difference in diversity. What approximation would we choose between A and B to present it to the DM?

Diversity has been interpreted in several ways, some examples are:

1. We want the vectors in the approximation set to be as evenly distributed as possible.
2. We want to have as many solutions as possible in the Pareto Front.
3. It is desirable to know the extreme values of the Pareto Front.

Several interpretations in previous list, can contradict our intuitive concept of diversity. For example, suppose that the vectors in Figure 1, right panel, are the Pareto Front of a multi–objective problem. According to item 1 in the list above, the approximation $A = \{a, b, d, e, f\}$ has a better diversity than the whole Pareto Front, but this is incorrect because the whole Pareto Front has more diversity than any of its subsets. Another example is item 2 in previous list. If the Pareto Front consist only in one vector, this interpretation of diversity is not adequate. A similar conclusion can be derived from item 3. These contradiction are a consequence of the way we interpret diversity.

In order to elaborate a new interpretation of diversity, we introduce some concepts. The preferences of the DM are influenced by the approximation that we show to him/her. The more and better information we present to the DM, the more "optimal" his/her decisions. Imagine the hypothetical situation where we know the whole PF and the DM is able to analyze all its elements. Under these conditions is valid to assume that the DM can identify perfectly its preferences[1] and select the really optimal solution. We call to the preferences derived under these conditions "the True DM preferences" or "True preferences" and the optimal solution derived under these conditions "the True optimal solution".

Different Decision Makers may have different True preferences. Even the same DM may have different True preferences over the time or under different conditions. We represent a arbitrary True preference with $\mathfrak{p}$. Each True preference has associate its corresponding True optimal solution, denoted $z_{\mathfrak{p}}$[2]. We assume that the True optimal solution of a True DM preference is an element of the Pareto Front. We denote with $\mathfrak{P}$ the set of all possible True preferences. The Pareto Front is the set of all possible True optimal solutions. When using the Pareto Optimality Criteria, there is an assumption that is very important to remember:

General Assumption 1 *The Decision Maker prefers non–dominated solutions over dominated ones. In other words, for $x, y \in X$, if x dominates y, then the DM considers x a better solution than y.*

There is another assumption related with diversity, that is widely accepted in the a posteriori approach:

General Assumption 2 *There is no preference, a priori, for any zone of the Pareto Front. All elements of the Pareto Front are equally important.*

All Pareto Optimal solutions have the same potential to be the True optimal solution. There are some situations where the General Assumption 2 is not accepted. For example, some algorithms gives preference to specific zones of the Pareto Front known as "knees" [1]. In these cases, more information about the DM preferences is introduced "a priori". We consider General Assumption 2 as valid, because most MOEAs are designed under this assumption.

Inspired by the framework of Hansen and Jaszkiewicz [5], imagine that for any multi–objective problem, there exist a density function $p(\mathfrak{p})$ over the elements of $\mathfrak{P}$ that gives the probability of $\mathfrak{p}$ to be the True preference chosen by the DM. Imagine a multi–objective problem whose Pareto Front is the one in Figure 1, right panel. Also, imagine that we select randomly DM preferences from $\mathfrak{P}$ based on the density function $p(\mathfrak{p})$ associated to the multi–objective problem. For each selected DM preference , we locate its corresponding $z_{\mathfrak{p}}$ in the Pareto Front. And we count how many times each element of the Pareto Front was the optimal solution for . With this information, we construct a density function $q(z)$ defined

[1] We call Decision Maker preferences to the information from the DM that can allow us to identify a unique optimal solution to multi–objective problem.

[2] This framework is similar to that of Hanzen and Jaszkiewicz [5], where they represented the DM preferences using utility functions.

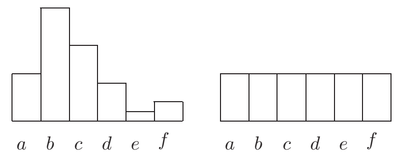

Fig. 2. Two probability ditributions.

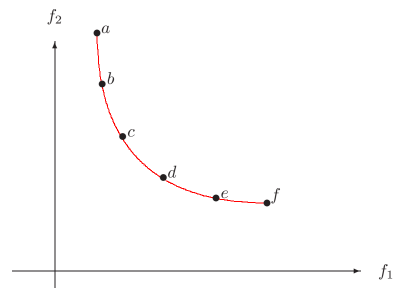

Fig. 3. A continuous Pareto Front.

over the elements of the Pareto Front that gives the probability of an element of the Pareto Front to be the True Pareto optimal solution.

Suppose that the density function $q(z)$, for the multi–objective problem described before, is as shown in Figure 2, left panel. We can conclude from this figure, that vector b is more likely to be the True optimal solution, followed by vectors c and a, respectively. If $q(z)$ is as shown in Figure 2, right panel, then we see that all the elements of the Pareto Front have the same probability to be the True optimal solution.

Now, consider the case where the Pareto Front is a continuous line, like in Figure 3 (points a, b, c, d, e and f are landmarks). Again, we can imagine a density function $q(z)$, on the Pareto Front, derived from a density function $p(\mathfrak{p})$, on the user preferences. Suppose that, in order to draw the density function, we align the elements of the Pareto Front horizontally. If the resulting density function $q(z)$ is as in Figure 4, left panel, we conclude that the most promising zones of the Pareto Front are those between b and c and in the neighborhood of e. If $q(z)$ is like in Figure 4, right panel, we see that all the elements of the Pareto Front have the same probability to be a True optimal solution of an arbitrary DM preference.

Different probability functions represent different information about the DM preferences. From the point of view of entropy, a constant probability function $q(z)$ represents the minimum quantity of information, where no predictions can be done. This is the condition that we assume when the DM preferences are unknown a priori. Based on this discussion we give an alternative definition to General Assumption 2.

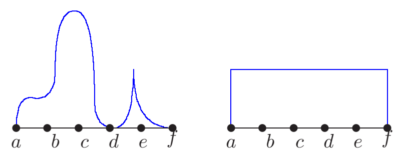

Fig. 4. Two density functions for a continuous Pareto Front.

Alternative Definition of General Assumption 2 1 *The density function $q(z)$, that gives the probability of a vector $z \in PF$ to be the True optimal solution for the DM, is a constant function.*

Now, imagine a Pareto Set with an infinite number of elements, like the one in Figure 3. Suppose that we obtain an approximation from a MOEA. A MOEA always generate a finite set of solutions. If the size of an approximation is finite, and the size of the Pareto Front is infinite, and the Alternative Definition of General Assumption 2 is true, what is the probability for True optimal solution to be an element of the approximation? The answer is, or course, zero. Except for the cases where the Pareto Front has a relatively small number of elements, the True optimal solution will not be an element of any approximation. So, when choosing a final solution from an approximation, the Decision Maker will not obtain the best solution $z_{\mathtt{p}}$. At most he/she will get an alternative solution, hopefully similar to $z_{\mathtt{p}}$. Based on this analysis, we propose the following definition of diversity:

Definition 1 *The diversity of a subset of the Pareto Front is its capacity to proportionate a similar alternative solution to any element of the Pareto Front.*

Remember, the Pareto Front is the set of all possible True optimal solutions. Unfortunately, Definition 1 is ambiguous and it needs to be interpreted in order to create a useful indicator of diversity. What makes an alternative solution similar to another one? How we measure the capacity to proportionate an alternative solution? We consider that two solution are similar if they provide a similar benefit to the Decision Maker. Based on intuition, we make the following assumption:

Assumption 1 *The similarity between two solutions is inversely proportional to their distance in objective function space.*

The nearer the solutions, the more similar they are. There is not a guaranty that Assumption 1 is always valid. But equivalent assumptions are frequently accepted in evolutive computation. For example, Assumption 1 is implicity accepted when the concept of niches is applied to multi–objective optimization. Assumption 1 is also accepted in the diversity mechanism of many MOEAs. For the rest of this work, we accept Assumption 1.

Now, we need to find a way to evaluate the capacity of an approximation to provide similar solutions to the elements of the Pareto Front.

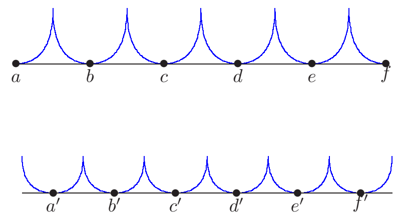

Fig. 5. Penalty functions for two different configurations of vectors.

4 Some Ideas For Indicators of Diversity

4.1 A Penalty Function

In this subsection, we construct a penalty function whose value increases as the capacity of an approximation to provide alternative solutions decreases.

Consider the Pareto Front in Figure 3. In order to simplify the following explanation, imagine that this Pareto Front can be represented with a continuous line in R^2. Also imagine that we "strain" and align horizontally this Pareto Front, preserving the relative distance between its elements trough the continuous line. Now, let A be a subset of the Pareto Front. For each element z of the Pareto Front, we add a penalty based on the distance of the nearest element of A. For example, we can use $d(z,a)^2$ as penalty, where $d(x,y)$ is the Euclidean distance between x and y. Our penalty (or error) diversity measure is defined as:

$$C(A) = \int_{z \in PF} d(z, a_{near})^2 dz \tag{3}$$

where a_{near} is the element of A nearest to z. The graphic of a penalty function can be seen in Figure 5, up, where a, b, c, d, e and f are an approximation set.

It is a common assumption that the extreme elements of the Pareto Front must be found. According to our penalty diversity measure this is not the best option. For example, in Figure 5, top, we have six vectors (set A) evenly distributed. The elements of A are also extended to reach the extreme elements of the Pareto Front. In the same figure, bottom, we have other six vectors (set B), whose elements are evenly distributed. But, the extreme vectors of B are not coincident with the extreme vectors of the Pareto Set. They are separated by half the distance between two adyacent elements of B. The value of Formula 3 is higher in A compared to its value in B. Therefore set B is better according to Formula 3. A reason for this, is that the maximum distance between an element of the Pareto Front and the nearest element of A is higher than the corresponding distance in B. So, A has higher penalty values.

Two desirable properties in a quality indicator for approximation sets, are monotony and relativity. A quality indicator with monotony improves the evaluation of an approximation as we add new non–dominated vectors. A quality indicator with relativity gives to the Pareto Front a better evaluation than any other approximation set.

Formula 3 has monotony. To see this imagine that we add a new element a to an approximation A. Let C be the elements of the Pareto Front nearer to a than to any other element of A. All vectors in C will have a reduction in their penalty value after the addition of a. As a consequence, the value of Formula 3 is reduced. Formula 3 also has relativity, because its value for the Pareto Front is zero, the minimum possible.

The penalty quality indicator described above, has many interesting theoretical properties. But, unfortunately, there are many practical restrictions for its implementation. First, we need to know the Pareto Front in order to use it, and for more than two dimensions, it is not possible to "strain" the Pareto Front. A partial solution to this problem is to project the approximations to a plane P and to define a domain $P' \subset P$. Next, Formula 3 is calculated using P' and the projected elements of the approximation. P' is a substitute for the Pareto Front, and the projection to the plane P is a substitution to "straining" the Pareto Front. For a discussion of how to project a non–dominated set to a plane, see [4]. The main problem of the approach just described, is how to define the domain for the integral in Formula 3. Different domains may result in different evaluations. Approximation A may be better than approximation B for one domain and the other way for another domain. Also, in order to detect the nearest elements of the domain to each element of an approximation, it is necessary to make a Voronoi tessellation, which is very expensive in higher dimensions.

Due to all those problems, it is desirable to find another function, cheaper to evaluate, that models diversity and that does not need knowledge about the Pareto Front.

4.2 A Benefit Function

Instead of penalizing based on how far are the elements of an approximation to the elements of the Pareto Set, we can add based on "how much" of the Pareto Set is near to an approximation.

Again, imagine a Pareto Front like the one in Figure 3 and suppose that we "strain" this set and align it horizontally. Suppose there exist a distance U, such that if the distance between a candidate solution a and the True optimal solution is less than U, then the DM will not mind choosing a instead of the True optimal solution. We call "the zone of influence" of a (I_a), to the elements of the Pareto Front whose distance from a is smaller than U, see Figure 6, top.

For an approximation set A, the union of all the zones of influence of its elements (I_A) can be use as measure of its diversity. I_A is the set of all the elements of the Pareto Front for which A can apport an alternative solution (assuming that we can find the value of U). We can measure the size of I_A, $\mu(I_A)$, measuring the length of the union of all I_a's. For the approximation in Figure 6, center, $\mu(I_A) = \alpha + \beta + \gamma$.

For a fixed value of U, there is a point where adding more elements to the approximation does not improve the value of $\mu(I_A)$. In order to correct this problem, we can associate a weight $w(z)$ to the elements z of the Pareto Front,

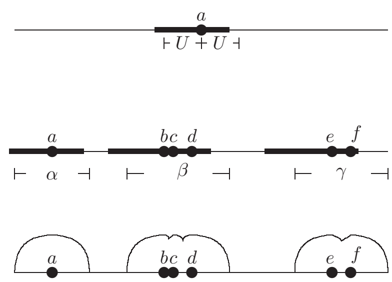

Fig. 6. The radius U define a zone of influence I.

depending on their closeness to an element of A. Many weight functions are possible, but in this work we use the following:

$$w(z) = \begin{cases} \sqrt{U^2 - d(z, a_{near})^2} & \text{if } d(z, a_{near}) \leq U \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

And the benefit quality indicator can be defined as:

$$B(A) = \int_{z \in PF} w(z) dz \tag{5}$$

In Figure 6, bottom, we see a graphical representation of $w(z)$. Formula 5 can be interpreted as the area of the union of semicircles centered in the elements of A. $B(A)$ has the property of monotony, because adding a new vector to A increases the area of the union of semicircles. $B(A)$ also has relativity, because the Pareto Front has the maximum area possible for the union of semicircles.

The restrictions to the practical implementation of $B(A)$ are similar to those of $C(A)$. The main restriction is the need to define a domain for the integral in $B(A)$. In order to solve this problem, we introduce some modification to $B(A)$. First, instead of semicircles, we can associate full circles to the elements of an approximation (imagine that we multiply $w(z)$ by two). This way, in order to evaluate the integral in Formula 5 we only need to calculate the area of the union of circles associated to each element of A. Next, instead of "straining" the Pareto Front, we can associate the circles (balls, from now) to the elements of the approximation in their original positions, as is shown in Figure 7. Denote with $b(U, a)$ the ball of radius U and center a, and with $\mu(Y)$ the measure of the set Y (in 2d it means area, in 3d it means volume, etc.). We define the diversity quality indicator $B2$ as:

$$B2(A) = \mu\Big(\bigcup_{a \in A} b(U, a) \Big) \tag{6}$$

$B2$ depends only in the position of the vectors in A and the balls associated to them. There is no need to know the Pareto Front. By intuition, we can see that the more diverse the elements of an approximation are, the smaller the intersection between balls, and the bigger the value of $B2$. As the diversity

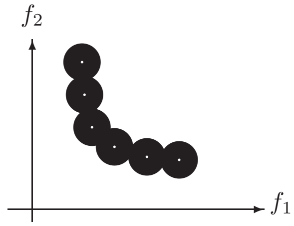

Fig. 7. B2.

of an approximation degrades, the amount of intersections between balls tend to increase and the value of $B2$ is reduced. As the value of $B2$ depends only in the position of the element in the approximation set, $B2$ can be evaluated independently of the topology of the Pareto Front.

Adding a new element to an approximation A does not degrade the value of $B2(A)$ (this is known as weak monotony). The value of $B2$ in the Pareto Front is not inferior to any approximation contained in the Pareto Set (this is known as weak relativity).

The measure of the union of M balls can be calculated in $O(Mlog(M))$ time in 2d, and in $O(M^2)$ time in 3d [3]. For more than 3 dimension we do not know an exact algorithm. But the measure of a union of M balls can be probabilistically approximated using a Monte Carlo integration. There is a Monte Carlo integration specially designed to this end whose complexity is linear with respect to M and with respect to the number of dimensions [7].

It is important to find a proper value for U. We want to determine the value of U based on the configuration of the approximation sets we are to compare. Suppose we have M approximations. For each element a of each approximation A_i, we find the nearest element to a (a_{near}) in the same approximation. Next, we calculate the distance between a and a_{near}. We call to this distances d_a. We set the value of U, as the mean value of the d_a's.

4.3 A Small Test

In [6] is proposed a benchmark for quality indicators. We selected three sets from the benchmark, $S1$, $S2$ and $S6$. We choose these sets because they can be ordered from the best to the worst, based on their diversity. $S1$, Figure 8, is the best set. Its elements are evenly distributed and its extension is the highest. $S2$, Figure 8, is the second best, their elements are evenly distributed but has a smaller extension than $S1$. $S6$, Figure 8, has a poor distribution and the same extension as $S2$[3]. After applying Formula 6 the results are: $B2(S1) = 11.32$, $B2(S2) = 6.73$, $B2(S6) = 5.52$, so $B2$ evaluated the sets correctly, giving better scores to the sets with better diversity.

[3] In [6], $S6$ has an slightly better extension than $S2$. In order to guarantee that $S6$ is the worst set, we adjusted it extension to be equal to $S2$.

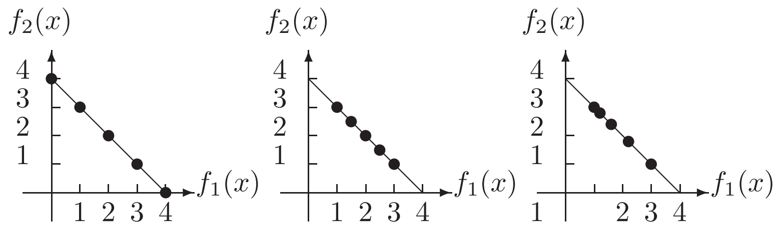

Fig. 8. S1 (left), S2 (center) and S6 (right).

5 Conclusions

In this work we presented a study on diversity. From this study we developed a new definition of diversity. From this definition we created two quality indicators for performance measures with desirable theoretical properties. These quality indicators are hard to implement in practice, but a simplified version of the second quality indicator was defined ($C2$). $C2$ was applied to some test cases published in literature, evaluating correctly the diversity of the sets.

References

1. J. Branke, K. Deb, H. Dierolf, and M. Osswald. Finding knees in multi-objective optimization. In *In the Eighth Conference on Parallel Problem Solving from Nature (PPSN VIII). Lecture Notes in Computer Science*, pages 722–731. Springer-Verlag, 2004.
2. D. V. Carlos A. Coello and G. B. Lamont. *Evolutionary Algorithms for Solving Multi-Objective Problems.* Kluwer Academic/Plenum Publishers, New York, USA, 2002.
3. B. K. B. David Avis and H. Imai. Computing the volume of the union of spheres. *The Visual Computer*, 3(6):323–328, 1988.
4. A. Farhang-Mehr and S. Azarm. Diversity assessment of pareto optimal solution sets: an entropy approach. *E-Commerce Technology, IEEE International Conference on*, 1:723–728, 2002.
5. M. P. Hansen and A. Jaszkiewicz. Evaluating the quality of approximations to the non-dominated set. Technical Report IMM-REP-1998-7, 1998.
6. T. Okabe, Y. Jin, and B. Sendhoff. A Critical Survey of Performance Indices for Multi-Objective Optimisation. In *Proceedings of IEEE Congress on Evolutionary Computation (CEC-2003)*, pages 878–885. IEEE Press, 2003.
7. P. G. Spirakis. The volume of the union of many spheres and point inclusion problems. In *Proceedings on STACS 85 2nd annual symposium on theoretical aspects of computer science*, pages 328–338, New York, NY, USA, 1985. Springer-Verlag New York, Inc.