

Pareto navigator for interactive nonlinear multiobjective optimization

Petri Eskelinen · Kaisa Miettinen ·
Kathrin Klamroth · Jussi Hakanen

Published online: 8 October 2008
© Springer-Verlag 2008

Abstract We describe a new interactive learning-oriented method called Pareto navigator for nonlinear multiobjective optimization. In the method, first a polyhedral approximation of the Pareto optimal set is formed in the objective function space using a relatively small set of Pareto optimal solutions representing the Pareto optimal set. Then the decision maker can navigate around the polyhedral approximation and direct the search for promising regions where the most preferred solution could be located. In this way, the decision maker can learn about the interdependencies between the conflicting objectives and possibly adjust one's preferences. Once an interesting region has been identified, the polyhedral approximation can be made more accurate in that region or the decision maker can ask for the closest counterpart in the actual Pareto optimal set. If desired, (s)he can continue with another interactive method from the solution obtained. Pareto navigator can be seen as a nonlinear extension of the linear Pareto race method. After the representative set of Pareto optimal solutions has been generated, Pareto navigator is computationally efficient because the computations are performed in the polyhedral approximation and for that reason function evaluations

P. Eskelinen (✉)
Helsinki School of Economics, P.O. Box 1210, 00101 Helsinki, Finland
e-mail: petri.eskelinen@hse.fi

K. Miettinen · J. Hakanen
Department of Mathematical Information Technology, University of Jyväskylä,
P.O. Box 35 (Agora), 40014 Jyväskylä, Finland
e-mail: kaisa.miettinen@jyu.fi

K. Klamroth
Department of Mathematics, University of Wuppertal, Gaußstr. 20, 42097 Wuppertal, Germany
e-mail: klamroth@math.uni-wuppertal.de

J. Hakanen
e-mail: jussi.hakanen@jyu.fi

of the actual objective functions are not needed. Thus, the method is well suited especially for problems with computationally costly functions. Furthermore, thanks to the visualization technique used, the method is applicable also for problems with three or more objective functions, and in fact it is best suited for such problems. After introducing the method in more detail, we illustrate it and the underlying ideas with an example.

Keywords Multicriteria optimization · MCDM · Interactive methods · Decision support · Pareto optimality

1 Introduction

Solving multiobjective optimization problems can be understood as finding the most preferred trade-off between conflicting objectives. In the field of multiple criteria decision making (MCDM), the idea is to help a decision maker in finding the best solution among mathematically incomparable compromises, so-called Pareto optimal solutions. During the years, many methods have been developed for this purpose (see, e.g., [Chankong and Haimes 1983](#); [Hwang and Masud 1979](#); [Miettinen 1999](#); [Sawaragi et al. 1985](#); [Steuer 1986](#)). However, their real-life applications are still surprisingly few ([Kaliszewski 2004](#)). One possible explanation could be that the tools developed for decision support may not be illustrative and easy-to-use enough for real decision makers.

Multiobjective optimization methods can be classified, for example, according to the role of the decision maker in the solution process ([Miettinen 1999](#)). Among the plethora of multiobjective optimization methods available, interactive methods are regarded promising because they allow an active participation of the decision maker in the solution process. In this way, the decision maker can direct the search and concentrate on solutions that are most interesting to her/him. For example, according to [Vanderpooten \(1989\)](#), two different conceptions regarding interactive methods can be identified. In searching-oriented methods, a converging sequence of solution proposals is presented to the decision maker. On the other hand, in learning-oriented methods, a free exploration of solutions is possible allowing trial and error. As mentioned in [Miettinen \(1999\)](#), the best way would be to combine these approaches. As a matter of fact, in many decision processes, one can identify two phases: a *learning phase* and a *decision phase*. No matter which style of expressing preference information is used (e.g., desirability of trade-offs, reference points, classification of objective functions etc.), it is often valuable for the decision maker first to be able to learn about the possibilities and limitations of the problem in order to adjust one's hopes on a realistic level and then fine-tune the final solution. It is also important to use concepts the decision maker understands well.

An attempt of developing an intuitive and understandable method for linear multiobjective optimization problems was suggested in [Korhonen and Wallenius \(1988\)](#) as a so-called Pareto race. There, the idea is that the decision maker can navigate in the set of Pareto optimal solutions like driving a car, in other words, move around in the Pareto optimal set according to his/her desires in order to identify the most preferred

trade-off. Thanks to parametric linear programming, Pareto race can show changes in objective function values in real-time. This is a very appealing learning-oriented approach but, as said, it works only for linear problems.

In this paper, we concentrate on (mostly convex) nonlinear problems and introduce a method that enables convenient and real-time navigation in the approximated Pareto optimal set of a nonlinear multiobjective optimization problem. We call this learning-oriented method by the name `Pareto navigator`. One can say that we extend and generalize the navigation ideas used in Pareto race to nonlinear problems. At the same time, we enable the decision maker to direct the search for the most preferred solution in more diverse ways than in Pareto race.

The starting point of our method is a relatively small discrete representation of the set of Pareto optimal solutions. Using these solutions, we generate a polyhedral approximation of the Pareto optimal set in the objective space. In this approximation, the decision maker can navigate according to his/her wishes and search for the most preferred trade-off. Because of the structure of the approximation, computation is fast and movements can be shown dynamically in real-time. Thus, our approach is particularly useful for problems where function evaluations are costly and time-consuming.

Because we wish to maintain the generality of our approach, we do not connect it to any particular method to generate the representation. Possible methods that can be used are given, for example, in [Deb \(2001\)](#), [Klamroth et al. \(2002\)](#), [Lotov et al. \(2004\)](#), [Ruzika and Wiecek \(2005\)](#). However, we emphasize that the representative set can be generated off-line before involving the DM. On the other hand, instead of response surface or kriging type of approaches to approximate individual objective functions (see, e.g., [Köksalan and Plante 2003](#)), we directly approximate the Pareto optimal set in the objective space in order to decrease computational complexity. Once the decision maker has identified an interesting region or solution, (s)he can ask for a more accurate approximation of the Pareto optimal set (i.e., more points in the approximation or concentrate the approximation in the desired region as e.g., in [Klamroth and Miettinen 2008](#)) or see (in some sense) the closest Pareto optimal solution in the original problem. Then one can continue navigating or continue with another interactive method after having learned about the feasible trade-offs.

The way the polyhedral approximation of the Pareto optimal set is generated in the Pareto navigator method is similar to what is presented in [Monz et al. \(2008\)](#) and utilized in the corresponding implementation ([Craft 2007](#)). However, in Pareto navigator, the navigation on the approximation is carried out in a completely different manner. In [Monz et al. \(2008\)](#), by using an approximation, the DM can study how improvement in the value of some objective affects the values of the other objectives. Differing from this objective-wise what-if analysis, in our approach, the navigation is carried out in a dynamic way, that is, instead of just setting an aspiration levels for some objective, the DM is allowed to determine a search direction and movement speed in this direction. The approximated Pareto optimal solutions available in the given direction are displayed to the DM in real-time with a continuously changing visualization. Thus, our approach concentrates on reflecting the dynamics of trade-offs between objectives.

The strengths of our approach include the applicability to computationally costly problems as well as a very intuitive user interface. Actually, in our method, the DM can

see approximative information about trade-offs in general, instead of typical point-wise trade-off considerations. In other words, the DM can see trade-off information between different conflicting objectives in real-time and conveniently control in which direction to move, that is, which objectives to improve or which objectives to allow getting worse.

Our approach enables the decision maker to learn about the interdependencies among the objectives in the problem and set one's expectations on a realistic level. Even though we are dealing with an approximation, the decision maker can get to see (in some sense) the closest Pareto optimal solution of the original problem corresponding to any approximate solution (by projecting the solution identified). We demonstrate our approach and a graphical user-interface of a software implementation with an example.

The rest of this paper is organized as follows. First, we introduce some notations and concepts used in Sect. 2. In Sect. 3, the algorithm of Pareto navigator is presented. Section 4 is devoted to implementation issues including a user interface and an example that demonstrates how Pareto navigator can be used. In Sect. 5, we shortly discuss some issues related to the development of the method and future research. We finish with some concluding remarks in Sect. 6.

2 Some notations and concepts

In this paper, we deal with multiobjective optimization problems of the form

$$\begin{aligned} & \text{minimize} && \{f_1(\mathbf{x}), \dots, f_k(\mathbf{x})\} \\ & \text{subject to} && \mathbf{x} \in S, \end{aligned} \quad (1)$$

where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ ($i = 1, \dots, k$) and $\mathbf{x} \in \mathbb{R}^n$ are called *objective functions* and a *decision (variable) vector*, respectively. A *set of feasible decision vectors* is denoted by S . The image of a feasible decision vector $\mathbf{x} \in S$ under mapping $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^k$ from the *decision variable space* \mathbb{R}^n to the *objective space* \mathbb{R}^k is called a *feasible objective vector* and denoted by $\mathbf{z} = \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_k(\mathbf{x}))^T$. The components of objective vectors are called *objective values*.

In this paper, we define optimality using the concept of Pareto optimality.

Definition 1 A decision vector $\mathbf{x}^* \in S$ and an objective vector $\mathbf{f}(\mathbf{x}^*)$ are said to be *Pareto optimal* if there does not exist another decision vector $\mathbf{x} \in S$ such that $f_i(\mathbf{x}) \leq f_i(\mathbf{x}^*)$ for all $i = 1, \dots, k$ and $f_j(\mathbf{x}) < f_j(\mathbf{x}^*)$ for at least one index j .

Note that the Pareto optimal solutions are in a mathematical sense incomparable without additional information. Furthermore, problem (1) typically has infinitely many Pareto optimal solutions. The set of all the Pareto optimal objective vectors is called the *Pareto optimal set*. Later on, this will also be referred to as an *actual Pareto optimal set*.

Usually in the MCDM field, the aim of solving a multiobjective optimization problem is to find a single feasible decision vector which is considered as the final solution for problem (1). However, we need some external information to decide which of the Pareto optimal solutions is the *most preferred* one. A *decision maker (DM)* is a person

who has knowledge about the problem in question and can express preference information related to Pareto optimal solutions. Naturally, what is to be regarded as most preferred depends on the DM involved.

It is often useful for the DM to know lower and upper bounds for the values of objective functions appearing in the Pareto optimal set. The *ideal objective vector* $\mathbf{z}^* \in \mathbb{R}^k$ consists of the optimal values z_i^* for each objective function f_i with respect to S . In other words, these are the best values that the individual objective functions can attain in the Pareto optimal set. If the objective functions are conflicting, which usually is the case, the ideal objective vector is infeasible. The *nadir objective vector* \mathbf{z}^{nad} gives upper bounds for the values of the individual objective functions in the Pareto optimal set. However, because the set of Pareto optimal decision vectors is unknown beforehand, we usually need to use an approximated nadir objective vector (see, e.g., [Miettinen 1999](#)) (in the case of more than two objectives). Because in this paper we assume that we have an approximation of the Pareto optimal set available in the form of a representation, we can use the extreme values of individual objectives present in this set to approximate the ideal and nadir vectors. Thus, there is no need to use additional function evaluations for calculating ideal or nadir vectors.

As motivated in the introduction, we consider here interactive approaches and, in particular, learning-oriented methods. The purpose of interactive methods is to aid the DM in finding a Pareto optimal solution which corresponds to the preferences of the particular DM in the best possible way. In interactive methods, a solution pattern is formed and iteratively repeated allowing the DM to adjust one's preferences and concentrate on solutions (s)he finds interesting. This means that the DM is directing the search according to her/his desires.

In interactive multiobjective optimization, there are several possibilities for the DM to express preference information ([Miettinen 1999](#)). For example, the DM can indicate desired changes in the objective function values of the current Pareto optimal solution by specifying a *classification*. Widely used classes are the following: objective functions whose values should be improved, are satisfactory or are allowed to impair. Note that if some objective function value is improved, then some other one must be allowed to impair in order to get another Pareto optimal solution. Another possibility to express preferences is to use reference points. A *reference point* $\bar{\mathbf{z}} \in \mathbb{R}^k$ consists of aspiration levels \bar{z}_i , $i = 1, \dots, k$, that represent desired values for the objective functions ([Wierzbicki 1980](#)). Note that the reference point does not need to be feasible. A reference point can also be extracted from the classification information ([Miettinen and Mäkelä 2002, 2006](#)). In that case, the components of the reference point $\bar{\mathbf{z}}$ are the ideal value z_i^* , the current value and the nadir value z_i^{nad} for the classes described above, respectively.

3 Pareto navigator

We propose a new interactive learning-oriented method, Pareto navigator, for multiobjective optimization. We use a polyhedral approximation of the Pareto optimal set in order to enable the DM to explore the Pareto optimal set. With Pareto navigator, the DM can conveniently study trade-offs between the conflicting objectives dynamically

in real-time and locate regions that are interesting to her/him. Therefore, we have a good learning tool to study the behavior of the problem. Because the exploration takes place in the approximated instead of the actual Pareto optimal set, it can be done with a low cost even for computationally expensive problems. Thus, our approach is particularly useful in complex real-world problems where function evaluations come, for instance, from some simulation tool and may necessitate solving systems like partial differential equations.

As already introduced, our method starts with a relatively small set of Pareto optimal solutions that is used to represent the actual Pareto optimal set. The polyhedral approximation generated by these solutions in the objective space will be referred to as an *approximated Pareto optimal set*. Solutions in this set will be called *approximated Pareto optimal solutions*. Using the approximated Pareto optimal set enables studying changes in objective function values in real-time when moving from one solution to another. This is realized using parametric programming. Once the DM has located an interesting solution or area from the approximation, (s)he is shown the *corresponding solution* in the actual Pareto optimal set. By a corresponding solution we here mean an actual Pareto optimal solution that is in some sense closest to the approximated solution. If desired, one can continue, for instance, with some other multiobjective optimization method which allows more detailed comparison of trade-offs between the Pareto optimal solutions (see, e.g., [Miettinen 1999](#)). If the approximated Pareto optimal solution and the corresponding actual Pareto optimal solution are far from each other, the DM can also ask for a more accurate approximation of the Pareto optimal set to be generated (e.g., by increasing the number of solutions that the polyhedral approximation is based on). It is also possible to improve the approximation locally, as suggested in [Klamroth and Miettinen \(2008\)](#), and study the polyhedral approximation related to a subset of the original Pareto optimal set. Our simplified setting allows us to decrease the cognitive burden set on to the DM while (s)he is making an overall evaluation about what kind of solutions may be available and which of them are interesting.

Usually, in interactive multiobjective optimization methods, the interdependencies between objectives are analyzed in a rather local sense, that is, mostly pointwise (e.g., trade-off rates at some solution). Instead, in Pareto navigator, the idea is to concentrate on capturing a global understanding of the possibilities and limitations in the problem considered. In other words, we provide means to study the overall behavior of Pareto optimal solutions. This is an essential part of the learning phase, discussed in the introduction. If a separate decision phase is needed after using Pareto navigator, one can switch to some other interactive method which typically are designed to support the decision phase.

The Pareto navigator method consists of two phases: *initialization* and *navigation*. The initialization phase is purely technical where we first produce a (relatively small) representative set of Pareto optimal solutions of problem (1). Using these solutions we generate a polyhedral approximation for the Pareto optimal set in the objective space. After the initialization phase is completed we are ready to start the navigation phase where the DM is in charge. In Fig. 1, the algorithm of the Pareto navigator method is presented from the DM's perspective using steps 1–6. Next, we describe more detailed actions related to these steps. In this context, we present an implementable framework

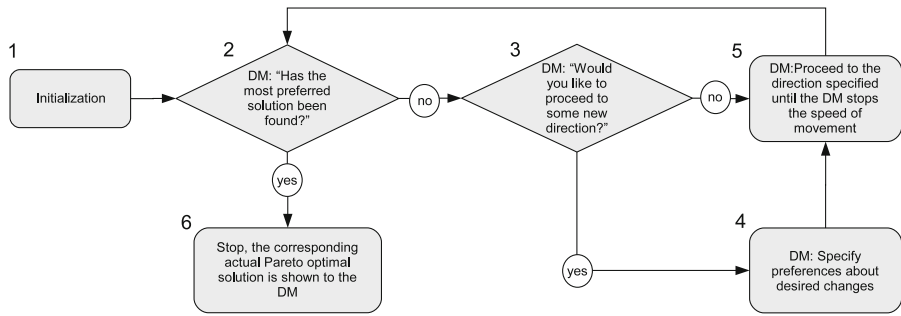


Fig. 1 Flowchart of Pareto navigator from the DM’s perspective

for Pareto navigator based on a convex polyhedral approximation of Pareto optimal set. However, as discussed in Sect. 5, the ideas related to the navigation can be applied also in the case of nonconvex polyhedral approximation.

3.1 Initialization phase

Before the DM can start to use Pareto navigator, an initialization procedure is carried out. In Fig. 1, this phase is denoted more generally as a step 1 but technically it contains the following two separate substeps.

1a. A discrete representation of the Pareto optimal set and a corresponding polyhedral approximation are produced in the objective space.

This step needs no interaction with the DM. Like pointed out in Sect. 1, there exist many methods that can be used to generate the discrete representation. However, it must be emphasized that with evolutionary approaches the resulting solutions may only be nondominated with respect to the final population but not actually Pareto optimal (the actual Pareto optimal set may dominate the polyhedral approximation considered). The production of the representation is computationally the most demanding task. However, this can be done well before the DM is involved in the method.

Let us assume that we have a discrete set of Pareto optimal solutions related to problem (1). As an example, we describe one possible way to build a polyhedral approximation for the actual Pareto optimal set using these solutions. Sometimes it is enough to consider a convex polyhedral approximation for the Pareto optimal set of problem (1) and in such a case we can utilize a convex hull of the Pareto optimal solutions available. A convex hull can be expressed in the form $Az \leq b$ and this inequality characterizes a polyhedral set. In other words, the inequality holds for objective vectors $z \in \mathbb{R}^k$ which belong to the convex hull. Later on, when step 5 is described, we explain how a convex hull of the form $Az \leq b$ can be used to obtain approximated Pareto optimal solutions.

In order to generate visualizations, we also need information about the ranges of the objective values in the Pareto optimal set. We can approximate ideal and nadir objective values as discussed in Sect. 2, or alternatively, it is possible to ask the DM for the best and the worst objective values to be considered.




Fig. 2 Value path visualization of the initial set of Pareto optimal solutions

1b. In what follows, the DM is involved. First, the DM is asked to select a starting point for the navigation phase.

For example, (s)he can specify a reference point which is then projected in the approximated Pareto optimal set. The projection can be made, for example, by solving a parametric programming problem (3) introduced in step 5. Alternatively, (s)he can select the most preferred actual Pareto optimal solution from the set used in step 1a. In this, a visualization like the one given in Fig. 2 can be helpful for the DM.

3.2 Navigation phase

The overall idea of the navigation phase is to allow the DM to move around in the approximated Pareto optimal set in those directions where (s)he feels the most promising solutions could be located. After the desired direction has been found, the method starts, in real-time, to produce approximated Pareto optimal solutions by moving into the direction determined. The approximated solutions are produced with a parametric linear programming procedure using a relatively small step size such that when the solutions are shown to the DM, (s)he experiences the motion as continuous. The speed of movement (i.e., how rapid the change in objective function values is) is determined by the DM. At any time the DM wishes, (s)he can adjust the speed and change the direction. In this way, the DM rapidly evaluates each approximated Pareto optimal solution produced along the determined direction. In practice, the objective function values can be visualized, for example, using a bar chart. In other words, the DM sees the lengths of the bars changing continuously as the corresponding objective function values change. Once the DM has found the most promising solution in the current direction, (s)he can stop the movement. This solution is called a *current solution*. In what follows, we describe steps 2–6 forming the navigation phase of Pareto navigator (see Fig. 1) in more detail.

2. The objective values of the current solution are visualized to the DM using a bar chart and the DM is asked the question “Has the most preferred solution been found?”. If the DM is satisfied with this solution, we proceed to step 6. Otherwise, we continue with step 3.

Note that in principle the DM may feel that the solution selected in the initialization phase is good enough as the final solution and we can stop without the navigation phase. However, we assume here that the DM is interested in learning about the problem and studying further solution possibilities.

3. If the DM is not satisfied with the current solution (i.e., (s)he answered “no” to the question in step 2), (s)he is further asked “Would you like to proceed to some new direction?”. A negative answer means that the DM wants to continue in the current

direction. Then we proceed to step 5. In the opposite case the DM can change one’s preferences and then a new search direction is determined.

Note that the DM must always specify preferences related to desired changes in objective values after the initialization phase. After this, the initial search direction can be determined correspondingly.

4. Now, the DM wants to change the search direction and (s)he is asked to specify preferences on how the current solution should be improved. The preference information is requested from the DM, for instance, in the form of a classification or as a reference point.

With the preference information, the DM indicates what kind of changes would make the current solution more preferred. This information is used to determine a new search direction from the current solution. Basically, any form of preference information can be given as long as it is possible to extract a unique search direction from it. For example, if a reference point $\bar{\mathbf{z}}$ is used, we can set the search direction as $\mathbf{d} = \bar{\mathbf{z}} - \mathbf{z}^c$, where \mathbf{z}^c is the current solution. Remember that a reference point can also be extracted from classification information, as described in Sect. 2. Determining the search direction plays a significant role when reflecting the preferences of the DM. Thus, special attention should be paid in selecting a convenient and intuitive form for the DM to specify preferences.

5. After the search direction has been determined, the method starts to produce approximated Pareto optimal solutions from the current solution \mathbf{z}^c . The DM is able to determine a desired speed of movement or (s)he can stop the movement at any solution desired. Then we continue with step 2.

As far as producing approximated Pareto optimal solutions is concerned, we solve a parametric linear programming problem to generate them in real-time. To be more specific, we consider the formulation

$$\begin{aligned} &\text{minimize} && \max_{i=1,\dots,k} w_i(z_i - \bar{z}_i(\alpha)) \\ &\text{subject to} && \mathbf{A}\mathbf{z} \leq \mathbf{b}. \end{aligned} \tag{2}$$

Problem (2) uses an achievement (scalarizing) function employed in the reference point method (Wierzbicki 1980, 1982, 1986) where $w_i, i = 1, \dots, k$, are the scaling coefficients and the reference point $\bar{\mathbf{z}}(\alpha) = (\mathbf{z}^c + \alpha\mathbf{d}) \in \mathbb{R}^k$ can be moved parametrically to the specified search direction $\mathbf{d} \in \mathbb{R}^k$ by altering the value of the parameter $\alpha \in \mathbb{R}$ (negative values of α allow us to move backwards). Here we use scaling coefficients $w_i = 1/(z_i^{\text{nad}} - z_i^*)$, for $i = 1, \dots, k$ (assuming $z_i^{\text{nad}} \neq z_i^*$). The linear constraints $\mathbf{A}\mathbf{z} \leq \mathbf{b}$ of problem (2) are forming a convex hull for a set of Pareto optimal solutions, as described in step 1a, and actually we are now projecting a reference point $\bar{\mathbf{z}}(\alpha)$ to the nondominated facets of the convex hull. This gives an approximated Pareto optimal solution. The parametric problem (2) above can be considered in an equivalent linear form by adding a new variable $\zeta \in \mathbb{R}$ (see, e.g., Benayoun et al. 1971). This leads to the following formulation

$$\begin{aligned} &\text{minimize} && \mathbf{c}^T \mathbf{z}' \\ &\text{subject to} && \mathbf{A}'\mathbf{z}' \leq \mathbf{b}' \\ &&& \mathbf{z}' \in \mathbb{R}^{k+1} \end{aligned} \tag{3}$$

where

$$\mathbf{c} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \mathbf{z}' = \begin{pmatrix} \zeta \\ z_1 \\ \vdots \\ z_k \end{pmatrix}, \quad \mathbf{A}' = \begin{pmatrix} -\frac{1}{w_1} & 1 & 0 & \dots & 0 \\ -\frac{1}{w_2} & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -\frac{1}{w_k} & 0 & 0 & \dots & 1 \\ 0 & a_{11} & a_{12} & \dots & a_{1k} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & a_{q1} & a_{q2} & \dots & a_{qk} \end{pmatrix}, \quad \mathbf{b}' = \begin{pmatrix} \bar{z}_1(\alpha) \\ \vdots \\ \bar{z}_k(\alpha) \\ b_1 \\ \vdots \\ b_q \end{pmatrix},$$

where q is the number of linear constraints in problem (2). The solution of problem (3) is an approximated Pareto optimal solution in the polyhedral approximation (i.e., due the projection it belongs to one of the nondominated facets of the convex hull). By utilizing a suitable optimizer we can generate new solutions very fast because we are solving this linear parametric programming problem.

6. Once the DM has found a satisfactory approximated Pareto optimal solution, we stop the search. The approximated solution is projected to the actual Pareto optimal set and the resulting solution is shown to the DM (with computationally demanding problem this may take some time).

The projection can be realized, for example, by setting the selected solution as a reference point and using some achievement scalarizing function to project it (see, e.g., [Miettinen 1999](#); [Wierzbicki 1986](#)). Note that the DM can ask for the actual Pareto optimal solution corresponding to any current solution. If the problem is convex, the approximated solution is always feasible. This means that the objective values in the corresponding actual Pareto solution are always at least as good as in the solution selected by the DM. If the DM is willing to continue navigation after this, we continue with step 2. If desired, it is possible to add the computed actual Pareto optimal solution to the approximation to make it more accurate. This, naturally, necessitates regenerating the polyhedral approximation, that is, we must go back to step 1.

After the learning phase has been completed with the Pareto navigator method (after step 6), the DM can stop (if the most preferred solution has been found) or proceed to the decision phase and continue, for example, with some interactive method (see, e.g., [Chankong and Haimes 1983](#); [Hwang and Masud 1979](#); [Miettinen 1999](#); [Steuer 1986](#)), as mentioned earlier. If the DM has specified preferences in the form of reference points or classifications, it may be natural and intuitive to continue with interactive reference point ([Jaszkiewicz and Slowiński 1999](#); [Wierzbicki 1982](#)) or classification based methods ([Miettinen and Mäkelä 1995, 2000, 2006](#); [Nakayama and Sawaragi 1984](#)), respectively. Alternatively, the DM can continue the learning phase and, for example, ask for a more accurate approximation of the Pareto optimal set to be generated in the neighborhood of the selected solution (see, e.g., [Klamroth and Miettinen 2008](#)). Then we go to step 1 to generate a new set of Pareto optimal solutions in the specified region and initiate Pareto navigator again.

4 Using Pareto navigator

In this section, we demonstrate how the Pareto navigator method can be used from the DM’s point of view. To be more specific, by using an example problem, we describe step-by-step what kind of interaction can take place during the navigation phase outlined in Fig. 1. However, it must be emphasized that visualizations used and the way how the DM indicates his/her preferences are only suggesting one possible approach, and in this respect the user interface can be customized also in other ways to meet the needs of the DM. We consider the following simple problem:

$$\begin{aligned}
 &\text{minimize} && \left\{ \begin{array}{c} -x_1 - x_2 + 5 \\ \frac{1}{5}(x_1^2 - 10x_1 + x_2^2 - 4x_2 + 11) \\ (5 - x_1)(x_2 - 11) \end{array} \right\} \\
 &\text{subject to} && 3x_1 + x_2 - 12 \leq 0 \\
 &&& 2x_1 + x_2 - 9 \leq 0 \\
 &&& x_1 + 2x_2 - 12 \leq 0 \\
 &&& 0 \leq x_1 \leq 4, \quad 0 \leq x_2 \leq 6.
 \end{aligned}$$

Actually, this example problem is not convex due the third objective function but the example demonstrates that this does not affect the applicability of Pareto navigator.

1a. To begin with, we apply an approximation method in the initialization phase and as a result we get a list of actual Pareto optimal solutions (see Table 1) to be used to produce a polyhedral approximation. We use the adaptive approximation method Klamroth et al. (2002). Here we have selected to use seven solutions. After using the adaptive approximation method, we also get approximations for the ideal and the nadir objective vectors. In what follows, we describe actions of a DM and explain how Pareto navigator can be used to search for a promising solution in the objective space.

1b. The DM specifies the starting point for Pareto navigator by selecting the most interesting actual Pareto optimal solution from the initial set listed in Table 1. The same set of solutions is also shown with a value path visualization in Fig. 2 (remember that the starting point could alternatively be specified with a reference point.) Let us now assume that the solution 4 is the most appealing to the DM.

2. In what follows, we refer to the solution selected as $A = (1.38, 0.62, -35.33)^T$ (see Fig. 3) and assume that the DM wants to examine its surroundings.

Table 1 Initial set of Pareto optimal solutions

	f_1	f_2	f_3
1	-2.00	0.00	-18.00
2	-1.00	4.60	-25.00
3	0.00	-3.10	-14.25
4	1.38	0.62	-35.33
5	1.73	1.72	-38.64
6	2.48	1.45	-42.41
7	5.00	2.20	-55.00

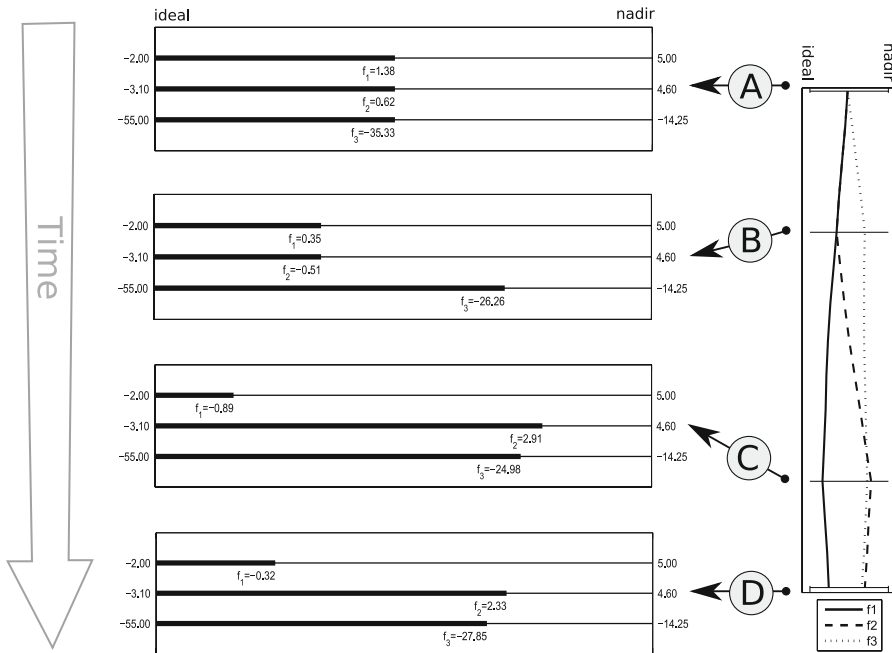


Fig. 3 The progress of the method

3. The initial solution *A* has now been selected but we do not yet have a search direction. After the initialization step, it is obligatory for the DM to specify preferences in order to set the search direction.

4. When studying *A*, the DM is interested in finding solutions where the values of objectives f_1 and f_2 are improving while the value of objective f_3 can be relaxed. By using this classification information, an initial search direction from the selected solution *A* can be generated. To be more precise, the classification made produces a reference point $\bar{z}^1 = (z_1^*, z_2^*, z_3^{nad})^T$ which is used together with the current solution *A* to produce the search direction $d^1 = (\bar{z}^1 - A)$.

5. Pareto navigator starts to generate solutions in the direction determined and updates in real-time the bar chart visualization depicting approximated solutions that are generated. Based on the continuously changing lengths of bars in the bar chart visualization, the DM can rapidly see what kind of solutions are available in the current direction (which corresponds to the given classification). When an interesting area has been reached, the DM stops the movement. Note that the DM is also allowed to move backwards in the current direction if the DM feels that (s)he already passed an area which seemed more interesting than the current one.

2. Let us now assume that the DM has arrived from the starting solution *A* to the solution $B = (0.35, -0.51, -26.26)^T$, where (s)he has stopped the movement (see Fig. 3). However, solution *B* is not satisfactory to the DM.

3. While moving from *A* to *B* the DM has learned how the solutions are behaving with respect to the classification given at point *A*. The DM thinks that continuing in

this direction is no more interesting and wants to change the direction by giving a new classification indicating how objective values should be changed from those obtained at B .

4. The DM classifies the objectives at B in such a way that the value of objective f_1 should be improved even more and therefore allows objective f_2 to degrade. On the other hand, the DM feels that the value of objective f_3 should maintain its current level. The classification gives us the reference point $\bar{z}^2 = (z_1^*, z_2^{\text{nad}}, B_3)^T$ and the search direction $\mathbf{d}^2 = (\bar{z}^2 - B)$ is generated.

5. The method starts to produce approximated solutions to the new search direction. The DM analyzes in real-time the bar chart visualization until an interesting solution is achieved.

2. The DM stops at solution $C = (-0.64, 1.82, -25.95)^T$ (see Fig. 3) and decides to explore the surrounding solutions further.

3. (S)he feels that the current direction no longer produces solutions that are interesting enough.

4. The DM is satisfied with the value of f_1 but is ready to sacrifice in its value to slightly improve the values of objectives f_2 and f_3 . This classification gives us the reference point $\bar{z}^3 = (z_1^{\text{nad}}, z_2^*, z_3^*)^T$ and the search direction $\mathbf{d}^3 = (\bar{z}^3 - C)$.

5. The DM continues to analyze approximated solutions in the direction determined until (s)he sees a promising solution.

2. The solution $D = (-0.32, 2.33, -27.85)^T$ (see Fig. 3) seems very satisfactory to the DM.

6. The navigation phase is stopped and the final approximated Pareto optimal solution D is now projected to the actual Pareto optimal set, as described in Sect. 3. The corresponding actual Pareto optimal solution $(-0.33, 2.32, -27.91)^T$ is computed and shown to the DM. The DM is satisfied with this solution, which is quite similar to the approximated solution D . The whole solution process can be stopped.

As we can see, the DM was able to learn about the solutions available and conveniently direct the search for the most preferred solution. The whole solution process described above is summarized and depicted in Fig. 3 where the vertical value path visualization on the right side is reflecting what kind of objective vectors the DM has analyzed during the solution process. On the left side of Fig. 3, the solutions A , B , C , and D are snapshots of the vertical value path. These are the solutions where the DM has temporarily stopped the movement and the continuously changing visualization in the bar chart, to determine a new search direction (i.e., classification). It is not possible to properly illustrate real-time movements here. Thus, it must be emphasized that during the solution process the DM analyzes approximated Pareto optimal solutions using only one bar chart visualization where the lengths of the bars are changing in real-time.

Because our simple example contains only three objectives, it is possible to illustrate what happened in the objective space during the solution process described. In Fig. 4, the triangulated area represents the polyhedral approximation (i.e., nondominated facets of the convex hull) which was constructed by using the initial set of Pareto optimal solutions (vertex points in Fig. 4). The approximation contains only nondominated facets so for the DM the movement on this polyhedral approximation feels like moving in the Pareto optimal set. Solutions A , B , C , and D are the same approximated Pareto optimal solutions that were obtained during the solution process.




Fig. 4 The progress in the objective space

Fig. 5 The actual Pareto optimal set




Figure 5 illustrates the actual Pareto optimal set. The black points in Fig. 5 corresponds to the vertex points in Fig. 4. It must be emphasized that in a general case it may be computationally very expensive to produce a visualization of the actual Pareto optimal set (as presented in Fig. 5). Naturally, for problems with more than three objective functions, we cannot generate visualizations like Figs. 4 and 5.

5 Discussion

In what follows, we wish to discuss some topics related to the implementation and future research. First of all, we must emphasize that in this kind of an interactive

method, the user interface plays a very important role. When designing the user interface, it is especially important to consider how the DM is able to indicate preferences (for setting a search direction) to utilize the full strength of the Pareto navigator method. The interaction should be as easy, intuitive and as fast as possible.

It is clear that the more accurate the representative set of Pareto optimal solutions is to begin with, the more reliable the results of Pareto navigator are. However, if the problem is computationally costly, it is better to generate the representative set before involving the DM. Instead of approximation algorithms of the MCDM field (Klamroth et al. 2002; Lotov et al. 2004; Ruzika and Wiecek 2005) we can also use evolutionary approaches (see, e.g., Deb 2001). The advantage of using an approximation algorithm as in Klamroth et al. (2002) is that it also produces an upper bound for the approximation error.

As far as determining the search direction is concerned, we can say that Pareto navigator is a more versatile method than Pareto race. This is because in Pareto race the DM can only say which objective function should be improved, whereas Pareto navigator allows the DM to express desires for all the objective functions about how their values should change. Future research could include a comparison of different methods that can aid the DM to determine directions while navigating in the approximation.

When describing the idea of classification in general, we mentioned that some objective function should be allowed to get worse if some other is desired to be improved. However, in Pareto navigator the DM can violate this, if so desired. In other words, the DM is allowed to ask for improvement in all the objectives. Naturally, it is not possible to improve all objective values simultaneously but the search direction can be set so. In this case, when producing new solutions, the method selects some objective functions to get worse in order to allow the others to improve. However, if the DM wants to feel being more in control, it is recommended that the classification expresses the DM's opinion about which objective functions can impair in value so that the others can improve. This kind of behavior must be also taken into account when designing a user interface.

Our method may have difficulties when dealing with, for example, design problems where the visualization of solutions needs variable values. In such cases, in addition to the objective values, the DM analyzes the quality of solutions via visualization. However, while moving in the approximated Pareto optimal set, the connection to the variable space is temporarily lost. This may be regarded as a shortcoming of Pareto navigator, but this is the price to be paid for enabling a computationally inexpensive navigation phase. Let us point out that the actual Pareto optimal solution together with the corresponding variable values can always be obtained by using the projection (in step 6 of the algorithm).

As mentioned before, once the DM has asked for the actual Pareto optimal solution corresponding to the current approximated one, one does not have to stop the whole solution process. Instead, one can continue the navigation from that point. It is also possible to update the approximation used by including the actual Pareto optimal solutions generated so far in the set that is the basis of the polyhedral approximation and repeating the initialization phase. In this way, the approximation can be made more accurate while the DM is using the method. Naturally, this option is convenient

only if we assume that the computation of the actual Pareto optimal solution does not take too much time.

Even though we have outlined in detail an implementable framework where Pareto navigator is utilizing a convex polyhedral approximation, the general navigation idea can be extended to work with a nonconvex polyhedral approximation as well. For example, as shown in Sect. 4, the applicability of Pareto navigator based on a convex approximation is not restricted to convex problems only. Sometimes studying a convex approximation of a nonconvex Pareto optimal set can offer accurate enough information for the DM. However, the implementation of the navigation idea in the case of highly nonconvex problems (that are too complex to be approximated by a convex polyhedron) is subject of future research.

In this paper, we have used the concept of Pareto optimality to determine what kind of solutions of problem (1) are interesting to the DM. However, despite the name, the Pareto navigator method can also be used in the case of more general dominance structures than Pareto dominance (see, e.g., [Sawaragi et al. 1985](#)).

6 Conclusions

We have described a new interactive learning-oriented method for nonlinear multiobjective optimization, called Pareto navigator. The method enables convenient navigation in the approximated Pareto optimal set. With this method, the DM can see changes in the trade-offs between conflicting objectives dynamically in real time. This intuitive approach gives general understanding of the possibilities and limitations of the problem considered and, thus, supports learning. Even though we are operating in an approximated Pareto optimal set, the actual Pareto optimal solution corresponding to any approximated one can be calculated whenever the DM so desires (but this may take some time depending on the computational complexity of the problem in question). Because most of the computations are carried out in the approximated Pareto optimal set, the method is particularly suitable for computationally challenging real-world problems.

We have demonstrated Pareto navigator with an example with three objective functions. However, our approach is not limited to such problems because bar charts can easily be used as the main means of visualization also when solving problems with more objectives. Particularly in cases when a visualization of the whole Pareto optimal set is not available, Pareto navigator is a useful tool to explore the Pareto optimal set in an interactive and efficient way.

Acknowledgments The authors want to thank Dr. David Craft who kindly offered technical information related to the Pareto Surface Navigator Matlab implementation ([Craft 2007](#)). This research was partially supported by Tekes, the Finnish Funding Agency for Technology and Innovation [Technology Program for Modelling and Simulation (MASI)].

References

Benayoun R, de Mongolfier J, Tergny J, Larichev O (1971) Linear programming with multiple objective functions: STEP method (STEM). *Math Program* 1:366–375

- Chankong V, Haimes YY (1983) Multiobjective decision making: theory and methodology. Elsevier, New York
- Craft D (2007) Pareto Surface Navigator. <http://www.mathworks.com/matlabcentral>, File Id: 13875
- Deb K (2001) Multi-objective optimization using evolutionary algorithms. Wiley, Chichester
- Hwang C-L, Masud ASM (1979) Multiple objective decision making—methods and applications: a state-of-the-art survey. Lecture notes in economics and mathematical systems, vol 164. Springer, Berlin
- Jaszkiewicz A, Slowiński R (1999) The light beam search approach—an overview of methodology and applications. *Eur J Oper Res* 113:300–314
- Kaliszewski I (2004) Out of the mist—towards decision-maker-friendly multiple criteria decision making support. *Eur J Oper Res* 158:293–307
- Klamroth K, Miettinen K (2008) Integrating approximation and interactive decision making in multicriteria optimization. *Oper Res* 56(1):222–234
- Klamroth K, Tind J, Wiecek MM (2002) Unbiased approximation in multicriteria optimization. *Math Meth Oper Res* 56:413–437
- Köksalan M, Plante RD (2003) Interactive multicriteria optimization for multiple-response product and process design. *Manuf Serv Oper Manage* 5:334–347
- Korhonen P, Wallenius J (1988) A Pareto race. *Naval Res Logist* 35:615–623
- Lotov AV, Bushenkov VA, Kamenev GK (2004) Interactive decision maps. Approximation and visualization of Pareto frontier. Kluwer, Dordrecht
- Miettinen K (1999) Nonlinear multiobjective optimization. Kluwer, Boston
- Miettinen K, Mäkelä MM (1995) Interactive bundle-based method for nondifferentiable multiobjective optimization: NIMBUS. *Optimization* 34:231–246
- Miettinen K, Mäkelä MM (2000) Interactive multiobjective optimization system WWW-NIMBUS on the internet. *Comput Oper Res* 27:709–723
- Miettinen K, Mäkelä MM (2002) On scalarizing functions in multiobjective optimization. *OR Spectr* 24:193–213
- Miettinen K, Mäkelä MM (2006) Synchronous approach in interactive multiobjective optimization. *Eur J Oper Res* 170:909–922
- Monz M, Küfer KH, Bortfeld TR, Thieke C (2008) Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Phys Med Biol* 53:985–998
- Nakayama H, Sawaragi Y (1984) Satisficing trade-off method for multiobjective programming. In: Grauer M, Wierzbicki AP (eds) Interactive decision analysis. Springer, Heidelberg, pp 113–122
- Ruzika S, Wiecek MM (2005) Approximation methods in multiobjective programming. *J Optim Theory Appl* 126:473–501
- Sawaragi Y, Nakayama H, Tanino T (1985) Theory of multiobjective optimization. Academic Press, New York
- Steuer RE (1986) Multiple criteria optimization: theory, computation and applications. Wiley, New York
- Vanderpooten D (1989) The interactive approach in MCDA: a technical framework and some basic conceptions. *Math Comput Modell* 12(10–11):1213–1220
- Wierzbicki AP (1980) The use of reference objectives in multiobjective optimization. In: Fandel G, Gal T (eds) Multiple criteria decision making theory and applications. Springer, Berlin, pp 468–486
- Wierzbicki AP (1982) A mathematical basis for satisficing decision making. *Math Modell* 3(25):391–405
- Wierzbicki AP (1986) On the completeness and constructiveness of parametric characterizations to vector optimization problems. *OR Spectr* 8(2):73–87