



UNIVERSITÀ DEGLI STUDI DI TRENTO

DIPARTIMENTO DI INGEGNERIA E SCIENZA DELL'INFORMAZIONE

38100 Povo — Trento (Italy), Via Sommarive 14
<http://disi.unitn.it/>

BRAIN-COMPUTER EVOLUTIONARY MULTI-OBJECTIVE
OPTIMIZATION (BC-EMO): A GENETIC ALGORITHM
ADAPTING TO THE DECISION MAKER

Roberto Battiti and Andrea Passerini

August 2009

Technical Report # DISI-09-060

Brain-Computer Evolutionary Multi-Objective Optimization (BC-EMO): a genetic algorithm adapting to the decision maker

Roberto Battiti and Andrea Passerini

Dipartimento di Ingegneria e Scienza dell'Informazione,
Università di Trento, Via Sommarive, 14 - 38050 Povo (Trento) - Italy
{battiti,passerini}@disi.unitn.it

August 2009

Abstract

The centrality of the decision maker (DM) is widely recognized in the Multiple Criteria Decision Making community. This translates into emphasis on seamless human-computer interaction, and adaptation of the solution technique to the knowledge which is progressively acquired from the DM.

This paper adopts the methodology of Reactive Optimization (RO) for evolutionary interactive multi-objective optimization. RO follows to the paradigm of “learning while optimizing”, through the use of online machine learning techniques as an integral part of a self-tuning optimization scheme.

User judgments of couples of solutions are used to build robust incremental models of the user utility function, with the objective to reduce the cognitive burden required from the DM to identify a *satisficing* solution. The technique of support vector ranking is used together with a k -fold cross-validation procedure to select the best kernel for the problem at hand, during the utility function training procedure. Experimental results are presented for a series of benchmark problems.

1 Introduction

Many evolutionary algorithms have been developed in the last years, starting at least from the eighties [1], to solve multiobjective optimization problem (MOOPs). A MOOP can be stated as:

$$\text{maximize } \mathbf{f}(\mathbf{x}) = \{f_1(\mathbf{x}), \dots, f_m(\mathbf{x})\} \quad (1)$$

$$\text{subject to } \mathbf{x} \in \Omega \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^n$ is a vector of n decision variables; $\Omega \subset \mathbb{R}^n$ is the *feasible region* and is typically specified as a set of constraints on the decision variables; $\mathbf{f} : \Omega \rightarrow \mathbb{R}^m$ is made of m objective functions which need to be jointly maximized.

Objective vectors are images of decision vectors and can be written as $\mathbf{z} = \mathbf{f}(\mathbf{x}) = \{f_1(\mathbf{x}), \dots, f_m(\mathbf{x})\}$. Problem 1 is ill-posed whenever objective functions are conflicting, a situation which typically occurs in real-world applications. In these cases, an objective vector is considered optimal if none of its components can be improved without worsening at least one of the others. An objective vector \mathbf{z} is said to *dominate* \mathbf{z}' , denoted as $\mathbf{z} \succ \mathbf{z}'$, if $z_k \geq z'_k$ for all k and there exists at least one h such that $z_h > z'_h$. A point $\hat{\mathbf{x}}$ is Pareto optimal if there is no other $\mathbf{x} \in \Omega$ such that $\mathbf{f}(\mathbf{x})$ dominates $\mathbf{f}(\hat{\mathbf{x}})$. The set of Pareto optimal points is called *Pareto set* (PS). The corresponding set of Pareto optimal objective vectors is called *Pareto front* (PF).

Evolutionary algorithms (EAs) work with a population of tentative solutions and they are therefore ideally suited to search for a set of Pareto-optimal solutions to be presented to the decision maker (DM). In this paradigm, evolutionary multiobjective optimization algorithms (EMOAs) aim at *building a representative set of points near the Pareto front*.

Some of the most successful EMOAs [2, 3, 4] rely on Pareto dominance classification as a fitness measure to guide selection of the new population. To adequately cover the PF, two important criteria are the proximity of the points to the PF and the diversity and spread of solutions, aiming at a uniform distribution on the PF. The work in [5] highlights a tradeoff between proximity and diversity that cannot be solved *a priori* without considering the specific demands of the decision maker. Furthermore, the definition of “uniform distribution” hides the fact that uniformity depends on the metric, on the scaling of objective values, in other words on the user preferences which cannot be fixed *a priori* [6].

The work [7] indicates that resorting to Pareto dominance classification to assign fitness becomes ineffective for increasing number of objectives and proposes a refined preference ordering based on the notion of order of efficiency [8]. The reason is that the proportion of Pareto-optimal elements in a set grows very rapidly with the dimension m , so that most solutions become in practice undistinguishable unless other criteria are added.

The curse of dimensionality also strikes when the decision maker is considered. The assumption that a typical DM likes the idea of being presented with hundreds or thousands of representative solutions on a multi-dimensional Pareto front is very far from reality. Although creating this representative set is appealing from a mathematical and research-oriented point of view, the reality of applied decision making must consider the following priorities.

- Bounded rationality and information bottleneck. The typical DM (being a typical human person) cannot deal with more than a very limited number of information items at a time. After Simon’s seminal papers in the fifties, aspects of bounded rationality permeate theories of choice, see for example the review in [9]. Human beings develop *satisficing* decision procedures, which are sensible given their constraints regarding human memory and information processing capabilities.
- Learning on the job. Most DMs cannot or do not like to explicitly formulate their objectives at the beginning. This is already recognized in the MOO formulation, where a combination of the individual objectives into a single preference function is not executed. Knowledge elicitation stops

at the building blocks given by the individual z_k 's. In addition, the traditional MOO community recognizes the need for interactive MOO, decision processes where learning by the DM and computer-supported solution processes work in coordination. According to [10], through interactive MOO the DM is building a conviction of what is possible and confronting this knowledge with her preferences that also evolve.

- Simple questions, qualitative judgments. The number of questions that have to be asked to the DM before a satisfactory solution is identified is a crucial performance indicator of interactive methods. This demands for selecting appropriate questions, for extracting as much information as possible from the answers, for building approximated models which may reduce the need to bother the DM. The complexity of the questions is also an issue. All DMs can be assumed to produce qualitative judgments, like "I prefer solution A to solution B." Asking for quantitative evaluations of the kind "By how much do you prefer A over B?" is in some cases inappropriate, artificial, or even impossible. Asking for *explanations* of choices is even more difficult: most DM are typically more confident in judging and comparing than in explaining.
- Uncertainty and inconsistency. The assumption that a fixed, deterministic and error-free preference structure of the DM is available is often not realistic. Imprecisions, contradictions, changes of judgment over time are the characteristics of most human decision processes.

The above priorities demand a shift of paradigm, from building a set of solutions which is representative of the true PF, to the interactive construction of a sequence of solutions where the DM is a crucial learning component in the optimization loop, a component characterized by limited rationality and scarce question-answering capabilities.

The motivation for our long-term investigation is in the systematic use of machine learning techniques for online learning schemes in optimization processes. In particular, the objective of Reactive Optimization (RO) [11] is to design problem solving schemes with an internal online feedback loop for the self-tuning of critical parameters. Internal parameters and algorithmic choices are adapted to the particular instance being solved.

In the context considered in this work, the objective of the learning process is the approximated construction of a utility function U to be optimized by the DM, who is also the source of learning signals. As in the traditional MOO context, the function to be optimized is not completely unknown like in a black-box context (see for example the response surface methodology and the related kriging and design of experiments approaches [12]), but is to be modeled after starting from the building blocks z_k characterizing positive features of a generic solution. Preference models are built from the DM input by using the support vector ranking method. The functional form of the preference function is not fixed *a priori*, like it is in the weighted sum or Tchebycheff approaches (see for example [13] for a clear explanation of scalarization methods) but is itself learnt during the process in a reactive fashion.

The rest of the work is organized as follows. In Section 2 we discuss the problem of modeling user preferences for interactive MOO and introduce our learning approach. Section 3 describes the support vector ranking method for

learning user preferences. The overall algorithm is detailed in Section 4 and related works are discussed in Section 5. An extensive experimental evaluation is reported in Section 6. Finally in Section 7 we draw some conclusions and propose possible directions for future research.

2 Brain-Computer Optimization: Learning user preferences in EMO approaches

Solving a MOO problem typically means providing a human decision maker with the solution she believes optimal according to a certain utility criterion allowing her to choose among competing Pareto optimal alternatives. Such utility criterion can of course be partially inconsistent, difficult to formalize and subject to revision according to the solutions provided by the optimization algorithm. Approaches to MOO can be roughly divided into the two broad categories of non-interactive and interactive ones [14]. The former can be further divided into *a priori* approaches, where the decision maker is required to formulate her preferences in advance, and *a posteriori* approaches, where the algorithm recovers a representative subset of the Pareto optimal set from which the user selects the preferred solution. *A priori* methods have the drawback of requiring the user to pre-specify her preferences, for instance as a set of weights on different objectives, which is typically rather hard for a human decision maker. *A posteriori* methods, on the other hand, imply a laborious selection among a large set of candidate solutions. Interactive approaches try to overcome some of these difficulties by keeping the user in the loop of the optimization process and progressively focusing on the most relevant areas of the Pareto front guided by the user feedback. The focus on fusing the capabilities of evolutionary computation (EC) with human evaluations reaches an extreme point with the interactive EC proposed in [15], where the fitness function is *replaced* by a human user. Our investigation follows an intermediate point, where knowledge of the objectives z_k is assumed *a priori*, and the DM is queried in order to build an explicit and robust nonlinear model of her preferences, to be used as an integral component of the problem-solving process.

Formalizing user preferences into a mathematical model is a non-trivial task. A model should be able to capture the qualitative notion of preference and represent it as a quantitative function, while retaining Pareto optimality properties. The simplest possible model is a linear value function, in which (positive) weights can be conceived as encoding the relative importance of different objectives:

$$U(\mathbf{z}) = \sum_{k=1}^m w_k z_k$$

However, its appropriateness is rather questionable [16]. Indeed, whenever objective functions correlate with each other, the most intuitive approach of giving highest weight to the most important criterion can lead to completely unsatisfactory solutions [17, 18]. Furthermore, assuming that satisfaction increases linearly with objective function decrease is not straightforward. Relying on the concept of decreasing limiting performance from mathematical economics, Podinovskii [19] argues that improvements on poorly satisfied objectives should be

considered more relevant than equally sized improvements on better satisfied ones. Most approaches generalize the linear utility function to a weighted- L_p metric of the following form:

$$U(\mathbf{z}) = - \left(\sum_{k=1}^m w_k |z_k^* - z_k|^p \right)^{1/p} \quad (3)$$

where \mathbf{z}^* is a reference ideal objective vector obtained by separately maximizing each objective function subject to the feasible region, i.e. $z_k^* = \max_{\mathbf{x} \in \Omega} f_k(\mathbf{x})$. A popular choice in this setting is that of the Tchebycheff (or L_∞) metric, leading to the following augmented weighted Tchebycheff program (AWTP):

$$\min \quad \alpha + \rho \sum_{k=1}^m (z_k^{**} - z_k) \quad (4)$$

subject to:

$$\begin{aligned} w_k (z_k^{**} - z_k) &\leq \alpha \\ \mathbf{w} \in \mathbb{R}^m, w_k &\geq 0, \sum_{k=1}^m w_k = 1 \\ f_k(\mathbf{x}) &= z_k, \mathbf{x} \in \Omega \\ k &= 1, \dots, m \end{aligned}$$

where \mathbf{z}^{**} is the utopian objective vector obtained adding a small positive scalar to the ideal vector \mathbf{z}^* . Let's comment: if ρ is set to zero, minimizing α amounts to minimizing the maximum weighted distance $w_k(z_k^{**} - z_k)$ between each individual objective and the utopian target. The augmentation by the term $\rho \sum_{k=1}^m (z_k^{**} - z_k)$, with ρ being a small positive scalar, generates *properly* Pareto optimal solutions, a more robust subset of Pareto points (a finite improvement in one objective is possible only at the expense of a *reasonable* worsening in other objectives).

The interactive weighted Tchebycheff procedure (IWTP) [20] consists of solving multiple AWTP, one for each choice of weight vectors \mathbf{w} , in order to present the DM with a small set of well-spaced, representative sample solutions. The DM is asked to select the most preferred among the proposed solutions, a new set of weights is produced which is consistent with the DM choice and a new set of solutions is generated, repeating the procedure until the DM is satisfied with the results. A more recent paper using a Tchebycheff utility function [21] focuses on minimizing the number of questions asked to the DM. Pairwise comparisons of solutions are used to generate constraints on the Tchebycheff weights and to identify multiple disjoint regions in weight space. Representative weights from the different regions are then generated to search for new *challenger* solutions to present to the DM.

Given the critique for linear weighting schemes, related also to the nonlinear preference of "compromise" solutions which is characteristic of many human decision activities, we explicitly consider nonlinear dependencies in this work. For instance, a generic polynomial utility function such as:

$$\begin{aligned}
U(\mathbf{z}) &= \sum_{i=1}^m w_i z_i + \sum_{i=1}^m \sum_{j \geq i} w_{ij} z_i z_j \\
&+ \sum_{i=1}^m \sum_{j \geq i} \sum_{h \geq j} w_{ijh} z_i z_j z_h + \cdots \\
&+ \sum_{i=1}^m \cdots \sum_{k \geq m-1} w_{i \dots km} z_i \cdots z_k z_m
\end{aligned} \tag{5}$$

would account for strong non-linear relations between objectives. However, we do not specify a priori a certain form for the utility function, but rather employ a Reactive Optimization scheme to determine the appropriate model to be used *while* the algorithm runs on a specific instance. The selection of a linear versus nonlinear model as well as the type of nonlinear model is decided in an automated way, depending on the user interaction and on the preference for simple models explaining the data which is the bread and butter of modern machine learning techniques. Our solution aims at:

- 1) being able to learn an *arbitrary* utility function from examples of preference information interactively provided by the DM
- 2) requiring DM intervention only through *holistic judgments* (comparisons of complete solutions instead of specification of detailed parameters like weights of trade-offs), by ranking sets of competing instances or specifying pairwise preferences between candidate solutions
- 3) naturally accounting for incomplete, imprecise and contradictory feedback from the DM
- 4) directly employing the learned utility function in order to guide the search for refined solutions

We rely on an adaptation of the well-known support vector machines [22] classification algorithm to preference learning in order to learn the utility function from user preference information. An EMO algorithm is run alternating a search phase, guided by a fitness measure based on the learned utility function, and a refinement phase where the DM is queried for feedback on candidate solutions and the utility function is updated according to such feedback.

3 Learning to rank

Given a sequence of competing instances $(\mathbf{z}_1, \dots, \mathbf{z}_\ell) \in \mathcal{Z}^\ell$ and an order relation \succ such that $\mathbf{z}_i \succ \mathbf{z}_j$ if \mathbf{z}_i should be preferred to \mathbf{z}_j , a *ranking* (y_1, \dots, y_ℓ) of the ℓ instances can be specified as a permutation of the first ℓ natural numbers such that $y_i < y_j$ if $\mathbf{z}_i \succ \mathbf{z}_j$. Learning to rank consists of learning a ranking function f from a datasets $\mathcal{D} = \{(\mathbf{z}_1^{(i)}, \dots, \mathbf{z}_{\ell_i}^{(i)}), (y_1^{(i)}, \dots, y_{\ell_i}^{(i)})\}_{i=1}^s$, of sequences with their desired rankings. Different approaches have been proposed in the literature to deal with ranking tasks (see [23] for a review). A common approach consists of learning a *utility* function $U : \mathcal{Z} \rightarrow \mathbb{R}$ measuring the importance of the instance, with the aim that $U(\mathbf{z}_i) > U(\mathbf{z}_j) \iff \mathbf{z}_i \succ \mathbf{z}_j$. In the following we will focus on an effective solution for learning a utility function based on suitable adaptations of the support vector machine algorithm [24, 25, 26]. The solution will naturally

account for situations in which only partial information on pairwise ordering is available for training. We will start with a brief discussion on support vector machines for classification [22] and extend the formulation to handle ranking tasks.

3.1 Support vector classification

Let $\mathcal{D} = (\mathbf{z}_i, y_i)_{i=1}^s$ be a training set of binary labelled examples, $y_i = 1$ or $y_i = -1$. Standard Support Vector Machines [22] for classification learn a decision function $h(\mathbf{z}) = \text{sign}(\langle \mathbf{w}, \mathbf{z} \rangle + b)$ trading-off fitting of the training data with large margin separation of classes, by solving the following optimization problem:

$$\begin{aligned} \min_{\mathbf{w} \in \mathcal{Z}, b \in \mathbb{R}, \boldsymbol{\xi} \in \mathbb{R}^s} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^s \xi_i & (6) \\ \text{subject to:} \quad & y_i(\langle \mathbf{w}, \mathbf{z}_i \rangle + b) \geq 1 - \xi_i \\ & \xi_i \geq 0 \\ & i = 1, \dots, s \end{aligned}$$




Figure 1: Classification problem solved by support vector machines. The solid line represents the separating hyperplane, while dotted lines are hyperplanes with confidence margin equal to one. Black points are unbound SVs, grey points are bound SVs and extra borders indicate bound SVs which are also training errors. All other points do not contribute to function to be minimized. Dotted lines indicate the margin error ξ_i for bound SVs.

The first term in the minimization accounts for a large margin separation, with margin given by $2/\|\mathbf{w}\|$, while the second accounts for penalties ξ_i for examples i not separated with the required margin (see Figure 1 for an example). The tradeoff factor C has to be adapted to the data, for example a noisy measurement process for \mathbf{z} implies that margin errors ξ_i be given a smaller weight (a smaller C value).

Whenever non-linear separation surfaces are needed in order to obtain a satisfactory separation, examples are projected into a suitable higher dimensional

feature space via a mapping function Φ (see Figure 2) and the same optimization problem is solved, simply replacing \mathbf{z}_i with $\Phi(\mathbf{z}_i)$. In the dual formulation of the optimization problem, projected examples always occur within dot products $\langle \Phi(\mathbf{z}_i), \Phi(\mathbf{z}_j) \rangle$. These can be replaced by an equivalent *kernel function* $k(\mathbf{z}_i, \mathbf{z}_j) = \langle \Phi(\mathbf{z}_i), \Phi(\mathbf{z}_j) \rangle$ on input examples which efficiently computes the dot product without explicitly doing the projection, which can even be infinite dimensional. The resulting decision function can be represented as a linear combination of kernel functions on support vectors: $h(\mathbf{z}) = \text{sign} \sum_{\mathbf{z}_i \in SV} \alpha_i y_i K(\mathbf{z}_i, \mathbf{z})$. Kernel functions can be seen as similarity measures generalizing dot products to arbitrary domains and allowing to decouple learning algorithms from example representation. The above decision function can be rationalized as follows: the final decision is obtained from a linearly weighted combination of the outputs y_i for selected training examples, the weights being influenced by the similarity between the current example \mathbf{z} and the training example \mathbf{z}_i . *Universal* kernels [27] are powerful classes of kernel functions which can uniformly approximate any arbitrary continuous target function. A popular class of universal kernels is the Gaussian kernel, which is written as:

$$K(\mathbf{z}, \mathbf{z}') = \exp(-\gamma \|\mathbf{z} - \mathbf{z}'\|^2) \quad (7)$$

where γ is a positive constant. Kernel functions for structured data such as string, trees and graphs have been developed in the literature. In some cases the “kernel trick” transforms the computational complexity in a radical way, from practically unaffordable scalar products to efficient computations. Details on kernel machines can be found in several textbooks, e.g., see [28].




Figure 2: Higher dimensional feature space projection via mapping function Φ for non-linearly separable problem.

3.2 Support vector ranking

The SVM formulation can be easily adapted to learning the utility function for ranking: $U(\mathbf{z}) = \langle \mathbf{w}, \Phi(\mathbf{z}) \rangle$. The constant term b can be omitted because only differences of U values matter for the ranking. Given a ranking dataset $\mathcal{D} = \{(\mathbf{z}_1^{(i)}, \dots, \mathbf{z}_{\ell_i}^{(i)}), (y_1^{(i)}, \dots, y_{\ell_i}^{(i)})\}_{i=1}^s$, it suffices to impose constraints on the

correct pairwise ordering of instances within a sequence: $U(\mathbf{z}_h^{(i)}) > U(\mathbf{z}_k^{(i)}) \iff y_h^{(i)} < y_k^{(i)}$. The resulting minimization problem can be written as:

$$\min_{\mathbf{w} \in \mathcal{Z}, \boldsymbol{\xi} \in \mathbb{R}^*} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^s \sum_{h_i, k_i} \xi_{i, h_i, k_i} \quad (8)$$

subject to:

$$\begin{aligned} \langle \mathbf{w}, \Phi(\mathbf{z}_{h_i}^{(i)}) \rangle - \langle \mathbf{w}, \Phi(\mathbf{z}_{k_i}^{(i)}) \rangle &\geq 1 - \xi_{i, h_i, k_i} \\ \xi_{i, h_i, k_i} &\geq 0 \\ h_i, k_i : y_{h_i}^{(i)} &< y_{k_i}^{(i)} \\ i &= 1, \dots, s \end{aligned}$$

Note that the formulation allows to naturally account for ties and partial rankings, as constraints are only included whenever two examples should be ranked differently.

The problem is equivalent to the standard SVM formulation in (6) if we consider examples as pairs $\delta\Phi(\mathbf{z}_{h_i}^{(i)}, \mathbf{z}_{k_i}^{(i)}) = \Phi(\mathbf{z}_{h_i}^{(i)}) - \Phi(\mathbf{z}_{k_i}^{(i)})$, all pairs labels as positive and no bias term as it cancels out in the difference. In the dual formulation, the utility function thus becomes a linear combination of kernel functions on example pairs:

$$U(\mathbf{z}) = \sum_{(\mathbf{z}_{h_i}^{(i)}, \mathbf{z}_{k_i}^{(i)}) \in SV} \alpha_{i, h_i, k_i} \left(K(\mathbf{z}_{h_i}^{(i)}, \mathbf{z}) - K(\mathbf{z}_{k_i}^{(i)}, \mathbf{z}) \right)$$

By noting that each example $\mathbf{z}_j^{(i)}$ can appear multiple times in the summation (depending on the number of constraints it is involved in), we can rewrite the utility function more compactly as a linear combination of kernel functions on examples:

$$U(\mathbf{z}) = \sum_{\mathbf{z}_j^{(i)} \in SV^*} \alpha_{i, j} K(\mathbf{z}_j^{(i)}, \mathbf{z})$$

where:

$$\alpha_{i, j} = \sum_{\substack{(\mathbf{z}_{h_i}^{(i)}, \mathbf{z}_{k_i}^{(i)}) \in SV \\ : \mathbf{z}_j^{(i)} = \mathbf{z}_{h_i}^{(i)}}} \alpha_{i, h_i, k_i} - \sum_{\substack{(\mathbf{z}_{h_i}^{(i)}, \mathbf{z}_{k_i}^{(i)}) \in SV \\ : \mathbf{z}_j^{(i)} = \mathbf{z}_{k_i}^{(i)}}} \alpha_{i, h_i, k_i}$$

therefore reducing the computational cost of calculating the utility function.

4 The BC-EMO algorithm

Support vector ranking has a number of desirable properties making it a suitable candidate for learning user preferences. First, it accepts supervision in terms of pairwise preferences, a much more affordable request for a human decision maker than a quantitative quality score. Second, it is well grounded on learning theory: its trading off data fitting and complexity of the learned hypothesis allows to effectively deal with noisy observations, a situation which is quite

likely to occur when receiving feedback from a human user with only partial knowledge on the domain at hand. Third, the ability to implicitly project data into a higher dimensional feature space via the kernel trick provides the needed flexibility in order to best approximate the underlying preference model of the specific user. The polynomial value function in eq. (5), for instance, can be easily implemented by an m -th degree inhomogeneous polynomial kernel: $K(\mathbf{z}, \mathbf{z}') = (1 + \langle \mathbf{z}, \mathbf{z}' \rangle)^m$.

The problem boils down to learning the utility function from examples of user preferences, and employing the learned function to guide the search for the globally optimal solution. The most natural approach would be to turn the MOO problem into a single objective optimization problem by directly optimizing the learned utility function. However, this requires the learned function to retain Pareto optimality, which is guaranteed if the function has non-negative weights. This non-negativity constraint cannot be easily incorporated into the general SVM formulation while retaining the advantages of the kernel trick. A simple and effective alternative can be conceived when the preference function is employed in conjunction with evolutionary optimization algorithms. EMOAs [2, 3] typically rely on Pareto dominance classification as a fitness measure to guide selection of the new population. Previous works [7] already highlighted that such a choice can be rather suboptimal for increasing number of objectives and proposed a refined preference ordering based on the notion of order of efficiency [8]. Here we exploit the same idea replacing order of efficiency with value of the learned utility function. An unordered population combining parents and offsprings of the current generation is sorted in steps by:

- 1) collecting the subset of non-dominated individuals in the population
- 2) sorting them according to the learned utility function
- 3) appending to the sorted set the result of repeating the procedure on the remaining dominated individuals.

Algorithm 1 describes this ranking procedure which clearly retains Pareto optimality. The procedure is terminated as soon as the desired set of s individuals is obtained. Any EMOA can be equipped with such a preference ranking procedure in order to guide the selection of the next population. Note that the combined ranking according to dominance and utility function preference is employed whenever comparisons between candidate individuals have to be made in creating the next generation. Algorithm 2 describes the procedure of a generic training iteration, which combines training or refinement of the utility function and preference ranking according to the learned function. Training consists of:

- 1) selecting a set of *exa* training individuals as the best ones according to the current criterion, where the utility function is replaced by random selection (within non-dominated individuals) at the first training iteration;
- 2) collecting pairwise preferences for these individuals from the decision maker and adding them to the set of training instances (empty at the first training iteration);
- 3) performing a kernel selection phase by a k-fold cross validation procedure in order to choose the best kernel for the problem at hand;
- 4) training the utility function on the overall set of examples with the chosen kernel.

The learned utility function is then used to sort the remaining set of candidate individuals.

Algorithm 1 Preference ordering based on utility function

```
1: procedure PREFORDER( $P_i, U, s$ )

   Input:
    $P_i$  unordered population
    $U$  utility function
    $s$  size of output population
   Output:
    $P_o$  ordered population

2:   identify Pareto non-dominated individuals  $P_i^*$  in  $P_i$ 
3:   build  $P_o$  ordering individuals in  $P_i^*$  according to  $U$ 
4:   if  $\text{len}(P_o) \geq s$  then
5:     return first  $s$  elements of  $P_o$ 
6:   else
7:     collect Pareto dominated individuals  $P_i^d \leftarrow P_i \setminus P_i^*$ 
8:     if  $P_i^d \neq \emptyset$  then
9:        $P_o^d \leftarrow \text{PREFORDER}(P_i^d, U, s - \text{len}(P_o))$ 
10:      concatenate the two lists  $P_o \leftarrow P_o + P_o^d$ 
11:     end if
12:     return  $P_o$ 
13:   end if
14: end procedure
```

The overall procedure, which we name BC-EMO as an acronym for *Brain-Computer Evolutionary Multi-Objective Optimization*, is described in Algorithm 3 for a generic EMO algorithm. The algorithm parameters are: the number of allowed training iterations (*maxit*), the number of training individuals for iteration (*exa*), the number of generations before the first training iteration (*gen₁*) and between two successive training iterations (*gen_i*), a minimal performance requirement for prematurely stopping the training phase (*thres*), the size of population for the EMO algorithm (*s*). Additional parameters can be required depending on the specific EMO algorithm employed (see the experimental section). The algorithm alternates a training phase where the current utility function is updated according to the DM feedback, and a search phase guided by the preference ordering procedure. When either the maximum number of training iterations or the desired accuracy level are reached, an additional search phase is conducted producing the final ordered population.

Our BC-EMO algorithm is a generic formulation which can be implemented on top of any EMO algorithm. In this work we employed the NSGA-II [40] EMOA¹. NSGA-II runs in its original formulation, including the crowded-comparison operator for guaranteeing a sufficiently diversified population, for *gen₁* generations. Our preference model is then trained according to the DM feedback and the modified preference ordering procedure in Alg. 1 replaces both the ordering of the new population and the selection criterion of the binary tournament selection operator. The crowding distance mechanism is switched off at this point as we are interested in directing the generation of new individuals in the most

¹available at <http://www.iitk.ac.in/kangal/codes.shtml>

Algorithm 2 Training procedure at a generic EMO iteration

1: **procedure** TRAIN(P_i, U_i, s, exa)

Input:

P_i unordered population
 U_i current version of the utility function
 s size of output population
 exa number of training individuals for iteration

Output:

P_o ordered population
 U_o refined version of the utility function
 res_o estimated performance of refined utility function

2: $P_{tr} \leftarrow \text{PREFORDER}(P_i, U_i, exa)$
3: obtain pairwise preferences for P_{tr} from the DM
4: sort P_{tr} according to user preferences
5: add P_{tr} to the current list of training instances
6: Choose best kernel K by k-fold cross validation
7: $U_o \leftarrow$ function trained on full training set with K
8: $res_o \leftarrow$ k-fold cv estimate of function performance
9: $P_o^* \leftarrow \text{PREFORDER}(P_i \setminus P_{tr}, U_o, s - \text{len}(P_{tr}))$
10: $P_o \leftarrow P_{tr} + P_o^*$
11: **return** P_o, U_o, res_o
12: **end procedure**

interesting regions according to the DM preferences. The preference learning module is based on the support vector machine package SVMlight².

5 Related Work

Interactive multiobjective optimization has been an active research area in the last decades. In the field of EMO [4], a survey of preference-handling methods is presented in [29], and a recent survey of interactive EMO algorithms is given in [30]. The objective of this section is not to give an exhaustive review of the different papers but to compare our approach with some significant and related approaches dealing with modeling user preferences and interacting with the DM.

Artificial neural networks have been used in a number of works [31, 32, 33] to learn user preferences in an interactive fashion. In the interactive FFANN procedure [31], a feed-forward artificial neural network is trained to approximate the user utility function using a sample of user evaluated non-dominated solutions, and employed as the objective function of a non-linear optimization problem. However, there is no guarantee that the improved solutions found by the non-linear optimization problem are non-dominated. FFANN were later [32] used within an IWTP to replace the diversity selection criterion with a user preference criterion as predicted by the FFANN. While the resulting procedure is guaranteed to always produce non-dominated solutions, the role of FFANN is

²available at <http://svmlight.joachims.org/>

Algorithm 3 The BC-EMO algorithm

1: **procedure** BC-EMO($maxit, exa, gen_1, gen_i, thres, s$)

Input:

$maxit$ maximum number of allowed training iterations

exa number of training individuals for iteration

gen_1 generations before first training iteration

gen_i generations between two training iterations

$thres$ performance requirement to stop training

s size of population

Output:

P final ordered population

2: $res \leftarrow 0, it \leftarrow 0, U \leftarrow \text{RAND}$
3: run the plain EMO for gen_1 generations
4: collect last population P
5: **while** $it \leq maxit \wedge res < thres$ **do**
6: $P, U, res \leftarrow \text{TRAIN}(P, U, s, exa)$
7: run the EMO for gen_i generations
8: guided by PREFORDER with U
9: collect last population P
10: **end while**
11: run the EMO for remaining number of generations
12: guided by PREFORDER with U
13: **return** the final population P
14: **end procedure**

limited to halving the set of solutions to be presented to the user at each iteration, and they cannot directly guide the search procedure. Huang et al [33] developed an alternative approach where a FFANN is directly employed to produce refined weight vectors for the next AWTP. The network is trained using weight vectors as inputs and as outputs the utility of the solution obtained by the AWTP using the corresponding weight vector. An optimization problem is then solved in order to determine the set of weights maximizing the output of the trained network, and the obtained weights are used in the next iteration of the IWTP.

While retaining the guarantee to produce non-dominated solutions, our approach has two main advantages. First, the IWTP based methods assume to be able to constraint the search to the correct subspace relying on a linear set of weights, one for each objective; on the other hand, our approach directly employs the learned utility function to guide the search and can be applied to highly non-linear combinations of objectives such as those in eq. (5). By automatically tuning the kernel to the problem at hand it is possible to effectively approximate utility functions of varying complexity as will be shown in the experimental evaluation. Second, our method can be trained with pairwise preference supervision as opposed to quantitative scores. As previous approaches pointed out [31, 32], pairwise preferences can be converted into scores following the method of the Analytical Hierarchical Process (AHP) [34]: a reciprocal com-

parison matrix is constructed from the pairwise comparisons and the normalized principal eigenvector of the matrix is computed. Each component of the normalized eigenvector, which is also named priority vector in this context, can be viewed as the score of the corresponding solution and used as the desired target. However, while the procedure perfectly fits the need to order a set of instances according to their pairwise comparisons, it is not straightforward that a function trained to fit such scores on a certain training set will correctly generalize to unseen instances. The function is required to match scores which depend on the training set generating the matrix, while the ranking constraints in problem (8) simply require the learned function to correctly sort the examples.

A number of approaches exist combining interactive EMO with pairwise preference information from the DM. Quan et al [35] developed an approach based on the notion of imprecise value function [36]. Pairwise preferences are translated into a set of constraints on the weights of a linear combination of normalized objectives, which determines a constrained subspace W . Two arbitrary solutions can be compared by solving two linear programming problems, minimizing $\sum_k w_k [v_k(z_k) - v_k(z'_k)]$ and its inverse, where $v_k(\cdot)$ is a normalization function. If either of the two problems has a positive minimum, preference between the two solutions can be uniquely determined, and in turn used to assign fitness to the individuals. Figueira et al. [37] also developed methods relying on the set of all *additive* utility functions consistent with user preferences:

$$U(\mathbf{z}) = \sum_{k=1}^m u_k(z_k)$$

where single-objective functions $u_k(\cdot)$ are general non-decreasing functions. Such utility functions are not used to guide the search of the evolutionary algorithm, but rather to recover desired solutions from the whole Pareto front returned by the search. Necessary (resp. possible) pairwise rankings are produced for pairs of solutions \mathbf{z} and \mathbf{z}' such that $U(\mathbf{z}) \geq U(\mathbf{z}')$ for all (resp. at least one of) the utility functions consistent with user preferences. These rankings are employed to generate a new sample of reference solutions together to their user preferences by interacting with the DM, and the process is repeated until the DM is satisfied with the obtained solutions. The method also accounts for intensity of preference, both global and at the level of single objectives. Phelps and Köksalan [38] employ the weighted L_p utility function in eq. (3) trained from pairwise preferences according to the middlemost weights technique [39]:

$$\begin{aligned} \max \quad & \epsilon & (9) \\ \text{subject to:} \quad & & \\ & \sum_{k=1}^m w_k = 1, w_k \geq \epsilon, \forall k = 1, \dots, m & \\ & \sum_{k=1}^m w_k (z_k^* - z_{i,k})^t \leq \sum_{k=1}^m w_k (z_k^* - z_{j,k})^t - \epsilon & \\ & \forall \mathbf{z}_i \succ \mathbf{z}_j & \end{aligned}$$

where ϵ is a preference margin and $z_{i,k}$ indicates the k^{th} component of objective vector \mathbf{z}_i . Constraints are iteratively removed in chronological order in case

of infeasibility. The method shares a number of similarities with our approach: the use of pairwise similarities as constraints in learning the utility function, the maximization of a preference margin and the use of the learned utility function as a fitness measure for the evolutionary algorithm. All these methods have the advantage of directly learning a function which retains Pareto optimality, with the drawback of constraining in different ways the set of allowed utility functions. On the other hand, our method aims at exploiting the ability of kernel machines to learn arbitrary functions, together with their noise robustness implied by the trade-off between data fitting and complexity of the learned hypothesis, which naturally accounts for imprecise and contradictory user preferences. These features closely correspond to real-world situations characterized by strongly non-linear and DM-dependent utility functions, and by the complex evolution of a decision process characterized by possible imprecisions and even partial contradictions.

6 Experimental evaluation

The experimental evaluation is focused on demonstrating the effectiveness of the BC-EMO algorithm in approximating user preferences and correctly guiding the search towards the most preferred solution. Given this focus, we did not attempt to fine-tune non-critical parameters which were fixed to reasonable values for all experiments. We chose a population size of 100, 500 generations, probability of crossover equal to 1 and probability of mutation equal to the inverse of the number of decision variables. The number of initial generations (gen_1) was set to 200, while the number of generations between two training iterations (gen_i) was set to 20. A perfect fit ($thres = 1$) was required in order to prematurely stop training. The number of training iterations and examples per iteration were varied in the experiments as detailed later on in the section. Concerning the learning algorithm, we fixed the C regularization parameter in eq. (8) to 100. A preliminary analysis showed it to be a reasonable value considering the cost of the resulting optimization problem and an assumption of no noise in the DM evaluation. In general, this parameter can be tuned to the problem at hand as it depends on the kernel employed, the amount of noise and the number of available training instances.

We selected a set of MOO test problem classes from the popular DTLZ [41] suite. Each problem class challenges different aspects of MOO and can be scaled to any number of decision variables and objectives. Table 1 describes the test problem classes we employed. We retained the original formulation of the problems as minimization tasks. Note that \mathbf{x}_m indicates the components of \mathbf{x} from m up to n , according to the notation used in [41]. The choice of this subset of problems is motivated by the form of their Pareto fronts which allowed us to design non-linear tasks, as will be detailed later on in the section. For test problem DTLZ1, we slightly restricted the feasible set for decision variables with respect to the original formulation (from $[0,1]$ to $[0.25,0.75]$) in order to rule out sparse objective vectors (i.e. those with one or few non-zero entries) from the Pareto front and make the preference learning task more challenging when increasing the number of objectives.

The first set of experiments is aimed at showing the effectiveness of the method in early focusing on the correct search area with very few queries to

Table 1: Test problem classes employed in the study (from [41])

Name	Description
DTLZ1	$\min_{\mathbf{x} \in \Omega} \mathbf{f}(\mathbf{x})$ $\Omega = \{\mathbf{x} \mid 0.25 \leq x_i \leq 0.75 \forall i = 1, \dots, n\}$ $f_1(\mathbf{x}) = 1/2 \cdot (1 + g(\mathbf{x}_m)) \cdot x_1 x_2 \cdots x_{m-1}$ $f_2(\mathbf{x}) = 1/2 \cdot (1 + g(\mathbf{x}_m)) \cdot x_1 x_2 \cdots (1 - x_{m-1})$ \vdots $f_m(\mathbf{x}) = 1/2 \cdot (1 + g(\mathbf{x}_m))(1 - x_1)$ $g(\mathbf{x}_m) = 100 \cdot \left[\mathbf{x}_m + \sum_{x_i \in \mathbf{x}_m} (x_i - 0.5)^2 - \cos(20\pi(x_i - 0.5)) \right]$
DTLZ6	$\min_{\mathbf{x} \in \Omega} \mathbf{f}(\mathbf{x})$ $\Omega = \{\mathbf{x} \mid 0 \leq x_i \leq 1 \forall i = 1, \dots, n\}$ $f_1(\mathbf{x}) = x_1$ \vdots $f_{m-1}(\mathbf{x}) = x_{m-1}$ $f_m(\mathbf{x}) = (1 + g(\mathbf{x}_m))h(f_1, f_2, \dots, f_{m-1}, g)$ $g(\mathbf{x}_m) = 1 + \frac{9}{ \mathbf{x}_m } \sum_{x_i \in \mathbf{x}_m} x_i$ $h = m - \sum_{i=1}^{m-1} \left[\frac{f_i}{1 + g} (1 + \sin(3\pi f_i)) \right]$
DTLZ7	$\min_{\mathbf{x} \in \Omega} \mathbf{f}(\mathbf{x})$ $\Omega = \{\mathbf{x} \mid 0 \leq x_i \leq 1 \forall i = 1, \dots, n\}$ $f_j(\mathbf{x}) = \frac{1}{\lfloor \frac{n}{m} \rfloor} \sum_{i=\lfloor (j-1) \frac{n}{m} \rfloor + 1}^{\lfloor j \frac{n}{m} \rfloor} x_i \forall j = 1, \dots, m$ $g_j(\mathbf{x}) = f_m(\mathbf{x}) + 4f_j(\mathbf{x}) - 1 \geq 0 \forall j = 1, \dots, m - 1$ $g_m(\mathbf{x}) = 2f_m(\mathbf{x}) + \min_{i,j=1, i \neq j}^{m-1} [f_i(\mathbf{x}) - f_j(\mathbf{x})] - 1 \geq 0$

the DM, for different test problems and number of objectives in the setting of linear utility functions. For each problem class, we generated a number of test instances by varying the number of objectives m from 2 to 10 and by setting the size of the input decision vector to $n = 2m$ for DTLZ1 and DTLZ6 and $n = 10m$, as suggested in [41], for DTLZ7. A total of 27 test instances was thus considered. Furthermore, we generated linear utility functions by randomly choosing weights in the range (0,1].

Figure 3 reports a representative set of results for test problems with four, eight and ten objectives for the three problem classes. The evaluation measure is the approximation error in percentage with respect to a *gold standard* final

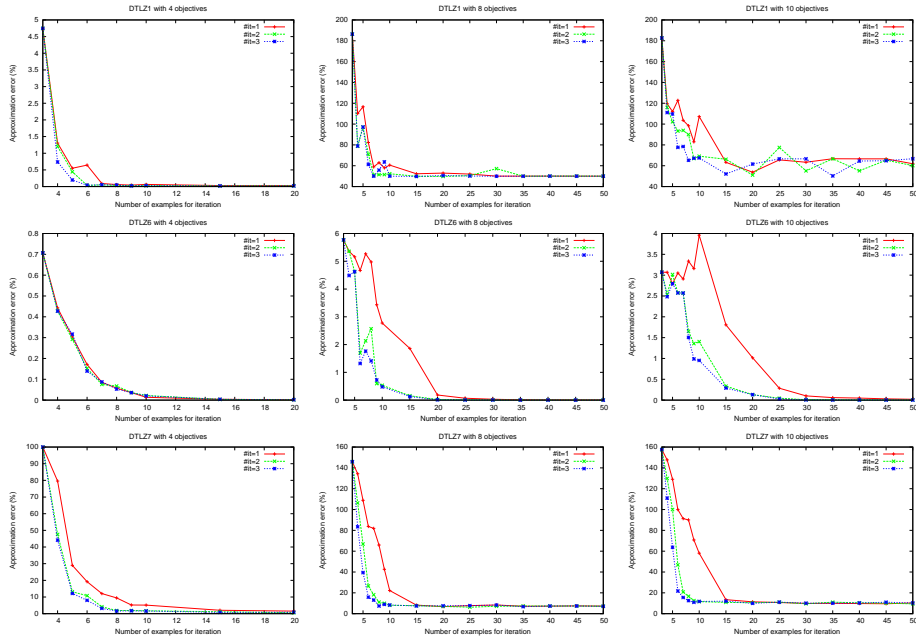


Figure 3: Learning curves for an increasing number of training examples per iteration with one (red/solid), two (green/dashed) and three (blue/dotted) training iterations. Each row reports results for a different problem class: DTLZ1, DTLZ6 and DTLZ7 respectively. Each column reports results for a different number of objectives, respectively four, eight and ten. Results are medians over 100 runs. Note that y-axes have different ranges.

solution, the one which is obtained by guiding the algorithm with the true utility function. Note that the ranking function does not need to correctly estimate the value of the utility function (and it does not in general), but only to rank good solutions higher than bad ones in order to correctly guide the search. Each graph reports three learning curves for an increasing number of training examples per iteration (*exa*), with one, two and three iterations (*maxit*) respectively. Results are the medians over 100 runs with different random seeds for the search of the evolutionary algorithm.

Problem classes present quite different characteristics: problems in class DTLZ1 (first row) were easily solved with just five training examples for up to four objectives, but were the most difficult to solve for a high number of objectives. Problems in class DTLZ6 (second row) were the easiest to solve: an approximation error of less than 1% could be achieved with at most 15 examples and two training iterations even in the case of ten objectives. Problems in class DTLZ7 (third row) required more examples than those in DTLZ1 in general but lead to better approximations with a high number of objectives. If one considers the number of training iterations, a second iteration was beneficial in many cases to possibly correct an early suboptimal model, especially with few training examples. Further training iterations did not provide substantial advantages.

Note that the DM is only required to compare examples within each training

iteration, as we assume that the last generation will represent the most relevant solutions to her. The learning algorithm (see Section 3.2) combines all these rankings into a single optimization problem: each training iteration i provides l_i ordered examples and a total of $s = \text{maxit}$ iterations are available. Given the algorithm formulation in eq. 8, there is no need for a full ranking of solutions, but the DM can provide partial information as sets of pairwise comparisons. In the presented experiments, for simplicity, we assumed a complete ranking.

In the next set of experiments we aimed at testing the ability of the method to automatically adapt to non-linear user preferences. Our aim is simulating many real world problems where the most preferred solution is a compromise among the different objectives. We thus generated test cases making sure that the utility function projected the most preferred solutions in a central area of the Pareto front, while retaining the Pareto dominance property. Designing such utility functions is not trivial. We heuristically generated them for each test problem by the following process: 1) running the plain NSGA-II algorithm to reach an approximation of the Pareto front 2) computing a reference point as the average of the final population 3) computing its distance to all other points 4) randomly generating weights of a second order polynomial 5) computing the utility value for all points according to the generated polynomial 6) keeping the polynomial if the rank correlation between the utility values and the distances to the reference point was higher than 0.4. This should guarantee that solutions distant from the center are less preferred.

BC-EMO was allowed to choose by internal 3-fold cross validation among linear and second degree polynomial kernel as well as Gaussian kernel with γ selected in the set $\{e^{-3}, e^{-2}, e^{-1}, 1, e^1, e^2, e^3\}$.

Figure 4 (left) shows an example of Pareto front reached by the plain NSGA-II algorithm for the DTLZ1 problem with two objectives. Figure 4 (right) shows the value of Pareto front according to the following polynomial utility function:




Figure 4: Problem DTLZ1 with two objectives: (left) Pareto front for a sample run of plain NSGA-II without user preference; (right) preference values of the Pareto front according to the non-linear utility function in eq. (10).

$$0.28 \cdot z_1^2 + 0.29 \cdot z_2 z_1 + 0.38 \cdot z_2^2 + 0.05 \cdot z_1 \quad (10)$$

Note that the solution which is most preferred by the DM ($z_1 = 0.249, z_2 = 0.251$) lies in the middle of the Pareto front. We ran our BC-EMO algorithm on this test problem and compared to two alternative versions: using a fixed linear kernel versus automatically tuning the kernel by internal k-fold cross validation.

Figure 5 shows the approximation results obtained for an increasing number of training examples. Results are medians over 10 runs. The linear kernel is completely unable to improve the performance regardless of the amount of training examples, stopping at an error of approximately 8%. On the contrary, tuning the kernel allows to reproduce the user preferences accurately enough to drive the search towards the desired solution. Let's note that in the initial point, corresponding to three examples, a linear kernel is chosen by default (insufficient data for a model selection by cross validation).

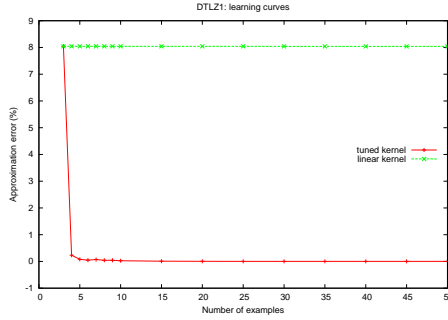


Figure 5: Approximation error for the non-linear DTLZ1 problem as a function of the number of training instances: tuned kernel vs linear kernel. Results are medians over 10 runs.

Problem DTLZ6 presents a highly disconnected Pareto front, as shown in Figure 6 (left) where plain NSGA-II was again employed to generate the sample. Figure 6 (right) shows the value of Pareto front according to the following polynomial utility function:




Figure 6: Problem DTLZ6 with two objectives: (left) Pareto front for a sample run of plain NSGA-II without user preference; (right) preference values of the Pareto front according to the non-linear utility function in eq. (11).

$$0.05 \cdot z_2 z_1 + 0.6 \cdot z_1^2 + 0.38 \cdot z_2 + 0.23 \cdot z_1 \quad (11)$$

This utility function generates two separate non-linear regions, with the global minimum located in the left region ($z_1 = 0.192, z_2 = 3.622$). Learning curves for the resulting problem are shown in Figure 7. The linear utility function converges to a suboptimal solution, which is actually located at the

left edge of the first region ($z_1 = 0, z_2 = 4$), while tuning the kernel allows to converge to the global optimum.




Figure 7: Approximation error for the non-linear DTLZ6 problem as a function of the number of training instances: tuned kernel vs linear kernel. Results are medians over 10 runs.

Finally, problem DTLZ7 presents a Pareto front similar to the one of DTLZ1, as shown in Figure 8 (left). Figure 8 (right) shows the non-linear surface obtained by using the following polynomial utility function:




Figure 8: Problem DTLZ7 with two objectives: (left) Pareto front for a sample run of plain NSGA-II without user preference; (right) preference values of the Pareto front according to the non-linear utility function in eq. (12).

$$0.44 \cdot z_1^2 + 0.33 \cdot z_1 + 0.09 \cdot z_2 z_1 + 0.14 \cdot z_2^2 \quad (12)$$

Figure 9 reports learning curves for the problem, again showing that a linear kernel is totally unable to locate the correct region, stopping at approximately 33% error, while the tuned kernel version converges to the solution ($z_1 = 0.148, z_2 = 0.407$) with less than ten training examples.

Increasing the number of iterations did not change results significantly. In particular, using a linear kernel leads to a wrong direction when generating new individuals and further training iterations are useless to correct this initial behaviour.

It is interesting to note that all non-linear cases show a similar behaviour in terms of form of the approximating model. In detail, seriously suboptimal solutions are found in the few cases in which a linear kernel is selected. This




Figure 9: Approximation error for the non-linear DTLZ7 problem as a function of the number of training instances: tuned kernel vs linear kernel. Results are medians over 10 runs.

happens by default with only three training examples and in rare cases with a very small numbers of examples: see for instance the spike for six examples in Figure 9. On the other hand, excellent fits are obtained when a second degree polynomial kernel is correctly chosen in the tuning phase. However, very good approximations are also achieved most of the times with Gaussian kernels, especially when enough training examples are available.

7 Conclusions

We presented a preference-based EMO algorithm (BC-EMO) characterized by its ability to learn an arbitrary utility function from a decision maker (DM) who expresses preferences between couples of selected solutions. The method to build a flexible preference model, possibly highly nonlinear, is based on Support Vector Machines and derived tools from the machine learning community. The optimization methodology of Reactive Optimization based on the paradigm of *learning while optimizing* is adopted in two directions: the progressive tuning of a preference model following a DM interactive evaluation and personal learning path, and the automated adaptation of the model form to one which is most appropriate, in a cross-validated manner, to the data collected during the interaction.

The method is robust as it can potentially withstand incomplete, imprecise and even contradictory feedback by the DM. The alternation of evolutionary optimization and DM ranking can be organized according to a flexible schedule, depending on the willingness by the DM to interact more times during the solution process and by intrinsic characteristics of the problem to be solved and the complexity of the user preferences.

The presented experimental results demonstrate the feasibility and effectiveness of the BC-EMO algorithm on a variety of benchmark tasks, with both linear and non-linear user preferences.

This research work leaves many avenues for future explorations. Future extensions of this investigation can consider a detailed analysis of the robustness of the proposed technique in different noisy scenarios. In addition, when considering promising techniques developed in the machine learning community, *active or query learning* methods are a viable candidate to further reduce the

number of comparisons required by the DM. The reduction can occur through the strategic choice of future examples to be presented for evaluations, depending on the past feedback received and therefore on the current preference model. Finally, *transfer learning* techniques could provide added benefits when solving a series of related problems: we plan to evaluate the effectiveness of transferring the previous experience accumulated in two different contexts: new and related problems solved by the same DM (for example problems sharing a subset of the objectives), or the same problem solved by DM's with different preferences. Last but not least, an extension of this work to discrete multi-objective optimization tasks is of interest, by using the effective kernel functions recently introduced for structured discrete structures like graphs and trees.

Acknowledgment

We acknowledge K. Deb, A. Pratap, S. Agarwal and T. Meyarivan for making available the software NSGA-II used as a starting point to develop our experiments, and T. Joachims for sharing the code of SVMlight. Part of this research has been supported by the Project Triton funded by Provincia Autonoma di Trento.

References

- [1] J. D. Schaffer, "Multiple objective optimization with vector evaluated genetic algorithms," in *Proceedings of the 1st International Conference on Genetic Algorithms*. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 1985, pp. 93–100.
- [2] N. Srinivas and K. Deb, "Multiobjective optimization using nondominated sorting in genetic algorithms," *Evolutionary Computation*, vol. 2, pp. 221–248, 1994.
- [3] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach," *Evolutionary Computation, IEEE Transactions on*, vol. 3, no. 4, pp. 257–271, Nov 1999.
- [4] K. Deb, *Multi-objective optimization using evolutionary algorithms*. Wiley, 2001.
- [5] P. 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.
- [6] J. Branke, "Consideration of Partial User Preferences in Evolutionary Multiobjective Optimization," in *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Springer-Verlag Berlin, Heidelberg, 2008, pp. 157–178.
- [7] F. di Pierro, K. Soon-Thiam, and D. Savic, "An investigation on preference order ranking scheme for multiobjective evolutionary optimization," *Evolutionary Computation, IEEE Transactions on*, vol. 11, no. 1, pp. 17–45, Feb. 2007.

- [8] I. Das, “A preference ordering among various pareto optimal alternatives,” *Structural and Multidisciplinary Optimization*, vol. 18, no. 1, pp. 30–35, 1999.
- [9] J. March, “Bounded rationality, ambiguity, and the engineering of choice,” *The Bell Journal of Economics*, pp. 587–608, 1978.
- [10] K. Miettinen, F. Ruiz, and A. Wierzbicki, “Introduction to Multiobjective Optimization: Interactive Approaches,” in *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Springer-Verlag Berlin, Heidelberg, 2008, pp. 27–57.
- [11] R. Battiti, M. Brunato, and F. Mascia, *Reactive Search and Intelligent Optimization*, ser. Operations research/Computer Science Interfaces. Springer Verlag, 2008, vol. 45.
- [12] D. Jones, “A Taxonomy of Global Optimization Methods Based on Response Surfaces,” *Journal of Global Optimization*, vol. 21, no. 4, pp. 345–383, 2001.
- [13] 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.
- [14] J. Branke, K. Deb, K. Miettinen, and R. Słowiński, Eds., *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Berlin, Heidelberg: Springer-Verlag, 2008.
- [15] H. Takagi *et al.*, “Interactive evolutionary computation: Fusion of the capabilities of EC optimization and human evaluation,” *Proceedings of the IEEE*, vol. 89, no. 9, pp. 1275–1296, 2001.
- [16] B. Roy and V. Mousseau, “a theoretical framework for analysing the notion of relative importance,” *Journal of Multi Criteria Decision Analysis*, 1996.
- [17] R. Steuer, *Multiple Criteria Optimization: Theory, Computation, and Application*. Wiley, New York, 1986.
- [18] L. Tanner, “Selecting a text-processing system as a qualitative multiple criteria problem,” *European Journal of Operational Research*, vol. 50, no. 2, pp. 179–187, January 1991. [Online]. Available: <http://ideas.repec.org/a/eee/ejores/v50y1991i2p179-187.html>
- [19] V. V. Podinovskii, “The quantitative importance of criteria with discrete first-order metric scale,” *Autom. Remote Control*, vol. 65, no. 8, pp. 1348–1354, 2004.
- [20] R. Steuer and E. Choo, “An interactive weighted Tchebycheff procedure for multiple objective programming,” *Mathematical Programming*, vol. 26, no. 3, pp. 326–344, 1983.
- [21] R. Dell and M. Karwan, “An interactive MCDM weight space reduction method utilizing a Tchebycheff utility function,” *Naval Research Logistics*, vol. 37, no. 2, 1990.

- [22] C. Cortes and V. N. Vapnik, "Support vector networks," *Machine Learning*, vol. 20, no. 3, pp. 1–25, Sep 1995.
- [23] S. Menchetti, "Learning preference and structured data: Theory and applications," Ph.D. dissertation, Universit degli Studi di Firenze, 2005.
- [24] W. W. Cohen, R. E. Schapire, and Y. Singer, "Learning to order things," *Journal of Artificial Intelligence Research*, vol. 10, pp. 243–270, 1999.
- [25] M. Collins and N. Duffy, "Convolution kernels for natural language," in *Advances in Neural Information Processing Systems 14*. MIT Press, 2001, pp. 625–632.
- [26] M. Collins, N. Duffy, and F. Park, "New ranking algorithms for parsing and tagging: Kernels over discrete structures, and the voted perceptron," in *In Proceedings of ACL 2002*, 2002, pp. 263–270.
- [27] C. A. Micchelli, Y. Xu, and H. Zhang, "Universal kernels," *J. Mach. Learn. Res.*, vol. 7, pp. 2651–2667, 2006.
- [28] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*. New York, NY, USA: Cambridge University Press, 2004.
- [29] C. Coello, "Handling preferences in evolutionary multiobjective optimization: a survey," in *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, vol. 1, 2000.
- [30] A. Jaszkiewicz and J. Branke, "Interactive Multiobjective Evolutionary Algorithms," in *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Springer-Verlag Berlin, Heidelberg, 2008, pp. 179–193.
- [31] M. Sun, A. Stam, and R. Steuer, "Solving multiple objective programming problems using feed-forward artificial neural networks: the interactive fpan procedure," *Manage. Sci.*, vol. 42, no. 6, pp. 835–849, 1996.
- [32] —, "Interactive multiple objective programming using tchebycheff programs and artificial neural networks," *Comput. Oper. Res.*, vol. 27, no. 7-8, pp. 601–620, 2000.
- [33] H. Z. Huang, Z. G. Tian, and M. J. Zuo, "Intelligent interactive multiobjective optimization method and its application to reliability optimization," *IIE Transactions*, vol. 37, no. 11, pp. 983–993, 2005.
- [34] T. Saaty, "Axiomatic foundation of the analytic hierarchy process," *Manage. Sci.*, vol. 32, no. 7, pp. 841–855, 1986.
- [35] G. Quan, G. Greenwood, D. Liu, and S. Hu, "Searching for multiobjective preventive maintenance schedules: Combining preferences with evolutionary algorithms," *European Journal of Operational Research*, vol. 177, no. 3, pp. 1969–1984, 2007.
- [36] C. White, A. Sage, and A. Dozono, "A model of multiattribute decisionmaking and trade-off weight determination under uncertainty," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 14, no. 2, pp. 223–229, 1984.

- [37] J. Figueira, S. Greco, V. Mousseau, and R. Słowiński, “Interactive Multi-objective Optimization using a Set of Additive Value Functions,” *Lecture Notes In Computer Science*, pp. 97–119, 2008.
- [38] S. Phelps and M. Koksalan, “An interactive evolutionary metaheuristic for multiobjective combinatorial optimization,” *Management Science*, pp. 1726–1738, 2003.
- [39] M. Koksalan, M. Karwan, and S. Zionts, “An improved method for solving multiple criteria problems involving discrete alternatives,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 14, no. 1, pp. 24–34, 1984.
- [40] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, “A fast elitist multi-objective genetic algorithm: Nsga-ii,” *IEEE Transactions on Evolutionary Computation*, vol. 6, pp. 182–197, 2000.
- [41] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, “Scalable multi-objective optimization test problems,” in *in Congress on Evolutionary Computation (CEC2002)*, 2002, pp. 825–830.