Modeling Behavior-Realistic Artificial Decision-Makers to Test Preference-Based Multiple Objective Optimization Methods

DOI:
10.4230/DagRep.5.1.96

Document Version
Final published version

Modeling Behavior-Realistic Artificial Decision-Makers to Test Preference-Based Multiple Objective Optimization Methods

Juergen Branke, Salvatore Corrente, Salvatore Greco, Miłosz Kadziński,
Manuel López-Ibáñez, Vincent Mousseau, Mauro Munerato, and Roman Słowiński

Report of Working Group #1
Dagstuhl Seminar 15031
"*Understanding Complexity in Multiobjective Optimization*", 2015

1 Introduction

Traditionally, in Multiple Objective Optimization (MOO), two separate methodological streams have been developed: evolutionary and interactive ones [2]. On the one hand, the role of Evolutionary MOO (EMO) is to approximate the entire Pareto front. On the other hand, Interactive MOO (IMO) deals with identification of the most preferred solution. IMO techniques require participation of a Decision Maker (DM) who is expected to provide her subjective preference information. The recent trend in MOO consists in merging the interactive and evolutionary approaches (for reviews, see [2,8,3]). This is achieved by integrating preference information into the EMO algorithms already during their optimization runs. The appealing effect of such integration consist in focusing the search on the area of the Pareto front which is most suitable to the DM.

Whenever DM preferences are used for guiding the search in MOO methods, the theoretical analysis [4] and experimental assessment of such algorithms is challenging, because it requires setting up a test environment that includes a model of the DM's behavior. Traditionally, artificial DMs have been simulated as a pre-defined value (utility) function for decision making. For example, the two user's functions used in an experimental setting in [3] assumed either linear weighting or a Tchebycheff-like aggregation of the objectives. In some other works, uncertainty of the DM interacting with an algorithm has been modeled by adding noise to an assumed function. In any case, the underlying model of an artificial DM is not known to a tested algorithm, but rather used to derive preference information that is subsequently provided at the method's input.

By contrast, the literature in (multiple criteria) decision making clearly identifies several cognitive biases, psychological phenomena, and inaccuracies occurring at the stage of problem modeling. Obviously, these highly affect preference elicitation and interactive decision making. Thus, the simple models of DMs most commonly used in the literature for testing IMO algorithms neglect the richness of human behavior and aggregate into a random component a variety of factors that should be rather modeled individually. The important factors that we identified are discussed in the following section. Altogether, they contribute to the idea of implementing a *machine DM* that would take into account the "true" criteria and "true" preference modified appropriately so that to approximate the behavior of real-world DMs.

2 Modeling cognitive biases, psychological phenomena, and inaccuracies of a machine Decision Maker

We call a machine DM, a model of DM biases and other factors that influence the interaction of the DM with the algorithm by modifying the true criteria and the true preference considered by the DM. This model does not actually specify the criteria and preferences considered by the DM, although different models (different machine DMs) may require them to satisfy certain characteristics. We decided to extend the machine DM from [7], which is based on previous work by Stewart [10], by modeling additional cognitive biases. Stewart [10] assumes a true preference function inspired by prospect theory, that is, a weighted sum of sigmoidals, and the biases modeled are omission of objectives, mixing of objectives and noise. We discuss these phenomena along with the newly considered ones in the following subsections.

2.1 Omission of criteria

Omission of criteria consists in neglecting by the algorithm some of the criteria that are internally considered by the DM [10,11]. For example, attributes of the problem that are modeled as constraints might be considered criteria by the DM. As noted by Stewart [11], the selection of the q criteria among m true ones (with $q < m$) can be conducted as follows:

- assign to each criterion g_j a uniformly generated weight w_j;
- order the criteria from the least ($rank = 1$) to the most important ($rank = m$);
- select q criteria randomly with probabilities proportional to the ranks of criteria so that less important objectives have a higher probability of being omitted.

In this scenario, the machine DM derives its preferences from the m-objective space, whereas the algorithm is allowed to refer the $q < m$ objectives only, i.e.:

$$\boldsymbol{g} \in \mathbb{R}^m \quad \text{(DM)} \quad \Rightarrow \quad \hat{\boldsymbol{g}} \in \mathbb{R}^{q<m} \quad \text{(algorithm)}. \tag{1}$$

Inversely, the machine DM may neglect some of the m criteria known to the algorithm by constructing its preferences on the basis of $q < m$ criteria only. In this case:

$$\boldsymbol{g} \in \mathbb{R}^{q<m} \quad \text{(DM)} \quad \Rightarrow \quad \hat{\boldsymbol{g}} \in \mathbb{R}^m \quad \text{(algorithm)}. \tag{2}$$

This bias can be modeled analogously to the previous one.

2.2 Mixing of criteria

Even if the criteria internally considered by the DM are preferentially independent, they may have been inadvertently corrupted when modeling the problem by mixing them in such a way that violates preferential independence [6,10,11]. Stewart suggests to obtain the new criteria in the following way [11]:

$$\hat{g}_k = (1-\gamma)g_j + \gamma g_{j+1}, \tag{3}$$

where $\gamma \in [0, 1]$ is a mixing parameter.

In the same spirit, even if the criteria have been defined so that to satisfy the requirement of preferential independence, one may introduce interaction components to the machine DM's value function. For example, Greco et al. [5] considered two particular types of such components corresponding to "bonus" and "penalty" values for positively or negatively interacting pairs of criteria. A bonus is added to (a penalty is subtracted from) the comprehensive value if a pair of criteria is in a positive (negative) synergy for performances of the considered alternatives on the two criteria. These effects may be considered as mutual strengthening or mutual weakening effects, respectively, which are both easily integratable into the model of a machine DM.

2.3 Mental fatigue

A great share of MOO methods require the DM to provide the preference information incrementally. On the one hand, this allows both avoiding the necessity of dealing with a large set of preference information pieces already at the initial stages of the interaction as well as controlling the impact of each piece of information (s)he supplied on the delivered results. On the other hand, a lengthy preference elicitation process may result in a mental fatigue of the DM. Such fatigue is defined as a temporary inability to maintain maximal cognitive performance from prolonged periods of cognitive activity (in our case, answering questions that would guide the search) [`http://en.wikipedia.org/wiki/Fatigue_(medical)`, last accessed: 10/03/2015]. Obviously, its onset depends upon an individual DM, but in general it is considered to be gradual. Thus, we have decided to model it as a noise factor $\sigma(k)$ that depends on the number of queries (k) to the DM. We found an exponential model $\sigma_0 \cdot e^{\alpha \cdot k}$ as appropriate for this purpose. Note that a closely related cognitive bias may consist in modeling mistakes of the machine DM just by negating or inverting the preferences derived from its model at random intervals.

2.4 Bounded rationality

The limited abilities of the DMs concerning information manipulation and computation have been accounted in the literature within the extensive studies on *bounded rationality* [9]. Indeed, the observed real-world decisions often violate the normative principles according to which all the relevant information should be taken into account. Various phenomena indicating that only a limited part of the available information is accounted in practical decision problems have been framed within so called *decision strategies* or *choice heuristics*. Using reverse engineering, these heuristics can be used for modeling the behavior of a machine DM with bounded rationality. For example, we may refer to:

- the satisficing heuristic [9] which (1) considers the solutions one after another, in a random way, (2) compares the value on each criterion of the current solution to a predefined level, and (3) selects the first alternative which passes this test; this procedure may potentially neglect a large part of the solution set;
- the elimination by aspect heuristic [12] which compares all solutions to a pre-defined aspiration level at each criterion starting from the most important one until a single alternative remains; thus defined, this approach considers a limited number of criteria.

2.5 Anchoring

Anchoring is a cognitive bias that describes the common human tendency to rely too heavily on the first piece of information offered (the "anchor") when making decisions. During decision making, anchoring occurs when individuals use an initial piece of information to make subsequent judgments [`http://en.wikipedia.org/wiki/Anchoring`, last accessed: 10/03/2015].

There are two levels of anchoring: a psychological or judgmental level, where there is no notion of gains or losses, and a reference-based level, where the DM defines her reference point according to earlier interactions and resists changing it. As a particular case of anchoring, we considered shifting the reference levels at each interaction according to the median value of each criterion for solutions shown to the DM (or the best solution found). However, we concluded that such a shift may have different interpretations depending on whether dynamic changes in true preference are desirable or not. Thus, we considered two models, where $U()$ is the true preference of the DM and $\hat{U}()$ is the perceived preference that determines the interaction with the algorithm:

- **Static (stable) model**, where interaction does not change the true reference point. In this model, anchoring means that interaction shifts the perceived reference point in $\hat{U}()$ from the true reference point in $U()$. The goal of the algorithm is to minimize the error with respect to the true (static) preference.
- **Dynamic model**, where interaction allows the DM to adjust her reference point (learn), that is, reference point changes in the true preference $U(\boldsymbol{z})$. In this model, anchoring means a resistance to change in $\hat{U}()$, when $U()$ changes. The goal of the algorithm is to minimize error with respect to the true preference at the last iteration. Such dynamic model may be treated as an example of evolving DM preferences, when the internal model of the DM changes as a result of the interaction with an algorithm.

We also tentatively discussed an additional *dynamic model with delayed adjustment of reference point* (Fig. 1), where the reference point is updated as:

$$\tau_i^{t^*} = \tau_i^{t_0} + (z_i^{t_0} - \tau_i^{t_0}) \cdot \frac{t^* - t_0}{\delta}$$

where $\delta > 0$ is a delay in adjusting preferences (anchoring). In this model, the goal is to minimize the error with respect to the true preference model at the last iteration + δ.

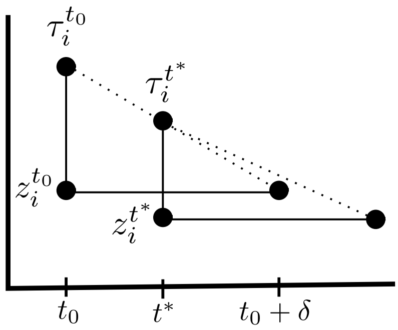

Fig. 1. Dynamic model with delayed adjustment of reference point

2.6 Loss aversion

The best solution identified so far in the course of an interaction with the MOO method may be treated by the DM as a reference point. When further exploring the objective space, the DMs tend to collate the newly constructed solutions with her actual reference. Such comparisons may be affected by a loss aversion bias, which implies that the impact of a difference on a criterion is greater when that difference is evaluated as a loss than when the same difference is evaluated as a gain [13]. Such asymmetry in perception of gains and losses with respect to the reference point $R = [r_1, \ldots, r_j, \ldots]$ may be easily modeled by transforming the DM's true function u_j in the following way:

$$R_j(x_J) = \begin{cases} u_j(x_j), & \text{if } x_j \geqslant r_j \\ \lambda_j u_j(x_j) - (1 - \lambda_j)u_j(r_j), & \text{if } x_j < r_j, \end{cases} \tag{4}$$

where λ_j is a coefficient of loss aversion for criterion g_j. In Figure 2, we provide exemplary indifference curves illustrating the application of thus defined transformation to a two-objective additive linear value function. These isoquants demonstrate that, when observing improvements on both objectives, the perception of the DM is unchanged, whereas the loss in performance at one objective negatively affects the overall quality of the solution from the point of the DM.

3 Conclusions and future work

The assumption that a true, not directly observable, preference exists is controversial on itself. One consequence of such a model is that, when attempting to avoid biases that distort this function, we are basically telling the DM that her behavior is somehow wrong. We recognize that this is a contentious issue, however, for simulation purposes, the existence of such a true preference is a useful working hypothesis which enables the analysis of how different biases affect interactive algorithms.

When modeling the machine DM, we can draw inspiration from previous literature on theoretical models and empirical studies with human DMs in (multiple criteria) decision making, behavioral economics, judgmental psychology, and cognitive science. In

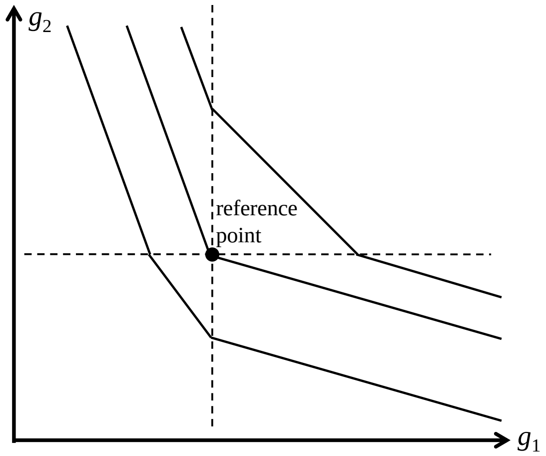

Fig. 2. Exemplary indifference curves illustrating loss aversion with respect to the reference point

this perspective, we feel that a thorough analysis of how DMs actually behave may gain yet another stream of applied research. That is, apart from having a relevant practice of preference elicitation and designing efficient preference elicitation procedures, we may design the procedures for deriving preference information to be provided at the input of tested algorithms.

Since the ultimate goal of modeling machine DM consists in using them for analysis and comparison of different methods, their models should be extended to various types of preference information, interaction and true preference models, in order to achieve as much generality as possible. During a group discussion, we decided to focus on how to model DM biases in the context of ranking and pairwise comparisons of solutions, nonetheless, we agreed that it is a worthwhile goal to understand how to simulate DM biases in the context of various types of interaction and preference information, including aspiration levels (goal programming), reference points, trade-offs, select 1 of n, sorting into categories, scoring, intensities of preferences, order of objectives, and desirability functions.

Our plan is to apply the proposed machine DM to NEMO-Choquet [1], which is an interactive evolutionary multiple objective algorithm based on the Choquet integral. Our intuition is that NEMO-Choquet should be able to cope with various biases, such as the mixing of objectives. In the medium term, we wish to do experiments that examine the trade-off between number of questions and quality of information, which decreases because of the fatigue, with respect to different types of questions (pairwise vs. ranking vs. aspiration levels vs. . . .). Future machine DMs should also simulate more biases such as bounded rationality heuristics in order to simulate more realistic behaviors.

References

1. Branke, J., Corrente, S., Greco, S., Słowiński, R., Zielniewicz, P.: Using Choquet integral as preference model in interactive evolutionary multiobjective optimization. Tech. rep., WBS, University of Warwick

(2014)
2. Branke, J., Deb, K., Miettinen, K., Słowiński, R. (eds.): Multi-objective Optimization: Interactive and Evolutionary Approaches, Lecture Notes in Computer Science, vol. 5252. Springer, Heidelberg, Germany (2008)
3. Branke, J., Greco, S., Słowiński, R., Zielniewicz, P.: Learning value functions in interactive evolutionary multiobjective optimization. IEEE Transactions on Evolutionary Computation 19(1), 88–102 (2015)
4. Brockhoff, D., López-Ibáñez, M., Naujoks, B., Rudolph, G.: Runtime analysis of simple interactive evolutionary biobjective optimization algorithms. In: Coello Coello, C.A., et al. (eds.) Parallel Problem Solving from Nature, PPSN XII, Lecture Notes in Computer Science, vol. 7491, pp. 123–132. Springer, Heidelberg, Germany (2012)
5. Greco, S., Mousseau, V., Słowiński, R.: Robust ordinal regression for value functions handling interacting criteria. European Journal of Operational Research 239(3), 711–730 (2014)
6. Keeney, R.L.: Analysis of preference dependencies among objectives. Operations Research 29, 1105–1120 (1981)
7. López-Ibáñez, M., Knowles, J.D.: Machine decision makers as a laboratory for interactive emo. In: Evolutionary Multi-criterion Optimization, EMO 2015. Lecture Notes in Computer Science, Springer, Heidelberg, Germany (2015), (To appear)
8. Rachmawati, L., Srinivasan, D.: Preference incorporation in multiobjective evolutionary algorithms: A survey. In: Proceedings of the 2006 Congress on Evolutionary Computation (CEC 2006), pp. 3385–3391. IEEE Press, Piscataway, NJ (Jul 2006)
9. Simon, H.A.: A behavioral model of rational choice. The Quarterly Journal of Economics 69(1), 99–118 (1955)
10. Stewart, T.J.: Robustness of additive value function methods in MCDM. Journal of Multi-Criteria Decision Analysis 5(4), 301–309 (1996)
11. Stewart, T.J.: Evaluation and refinement of aspiration-based methods in MCDM. European Journal of Operational Research 113(3), 643–652 (1999)
12. Tversky, A.: Choice by elimination. Journal of Mathematical Psychology 9(4), 341–367 (1972)
13. Tversky, A., Kahneman, D.: Loss aversion in riskless choice: a reference-dependent model. The Quarterly Journal of Economics 106(4), 1039–1061 (1991)