Machine Decision Makers as a Laboratory for Interactive EMO

Manuel López-Ibáñez[1] and Joshua Knowles[2]

[1] IRIDIA, Université Libre de Bruxelles (ULB), Brussels, Belgium
manuel.lopez-ibanez@ulb.ac.be
[2] School of Computer Science, University of Manchester, UK
j.knowles@manchester.ac.uk

Abstract. A key challenge, perhaps the central challenge, of multi-objective optimization is how to deal with candidate solutions that are ultimately evaluated by the hidden or unknown preferences of a human decision maker (DM) who understands and cares about the optimization problem. Alternative ways of addressing this challenge exist but perhaps the favoured one currently is the interactive approach (proposed in various forms). Here, an evolutionary multi-objective optimization algorithm (EMOA) is controlled by a series of interactions with the DM so that preferences can be elicited and the direction of search controlled. Multi-criteria decision-making (MCDM) has a key role to play in designing and evaluating these approaches, particularly in testing them with real DMs, but so far quantitative assessment of interactive EMOAs has been limited. In this paper, we propose a conceptual framework for this problem of quantitative assessment, based on the definition of machine decision makers (machine DMs), made somewhat realistic by the incorporation of various non-idealities. The machine DM proposed here draws from earlier models of DM biases and inconsistencies in the MCDM literature. As a practical illustration of our approach, we use the proposed machine DM to study the performance of an interactive EMOA, and discuss how this framework could help in the evaluation and development of better interactive EMOAs.

Keywords: machine decision makers, artificial decision makers, MCDM, interactive EMO, performance assessment

1 Introduction

Good introductions to the current state of research in interactive evolutionary multi-objective optimization algorithms (iEMOAS) [3, 10, 21] indicate that several important issues remain to be tackled when combining methods from multi-criteria decision making (MCDM) and evolutionary multi-objective optimization (EMO). We think that a most fundamental one is that few interactive approaches have been evaluated in a quantitative way, in particular in a way that would reveal how effective and efficient they are at finding a solution the decision-maker (DM) finds satisfactory or preferred. More than this, we think that quantitative assessment of methods involving decision-making remains controversial[3], and this controversy has perhaps led to a position where it is standard to propose new methods, and to justify them on theoretical grounds, but not to test them in a way that reveals practical properties. Testing is difficult (and controversial) because of the need for a DM who interacts with the system (and cares about the result), and the fact that DMs are all different [6, 27]. Nevertheless, there is ongoing research pursuing the goal of quantitative assessment by means of empirical analysis [7, 16]. In one of the first works pursuing a theoretical quantitative analysis, Brockhoff et al. [5] derived bounds on both the runtime and the number of queries to the DM for two different local-search based EMOAs, under the simplifying assumption that the DM would be able to answer a query whenever the EMOA had an incomparable pair of solutions to choose between. A common criticism of such quantitative analysis is that they rely on strongly simplifying assumptions about human DMs.

Simulating human DMs is not an easy task and most work in MCDM has focused on how to model and elicit the DM's preferences. Few works have considered the simulation of human factors and other non-idealities of the decision-making process. Kornbluth [17] incorporates "unsureness" of the DM as a range of values of the utility function for which the DM cannot make a decision. Morgan [19] describes the development of an even more realistic model to simulate expert decision-making in dynamic and time-critical environments (e.g., air combat); yet, the model itself is not fully described and its applicability to other contexts remains unclear. Perhaps the most extensive simulation of DM biases and other non-idealities has been conducted by Stewart [22, 23, 24], who studied how the ranking of efficient solutions is affected by these non-idealities. Nonetheless, the usual approach when evaluating interactive EMOAs in the recent literature is to add random noise to the DM's preferences [7, 16].

The step which we propose (and begin to investigate here) beyond recent studies is to use *more realistic* machine decision makers, and to perform a more detailed and multi-factorial analysis of the relationships between performance, DM satisfaction (i.e., finding and recommending a solution close to the ideal one the DM would choose), DM biases, and the EMO/MCDM interactive approach.

This paper is structured as follows. Section 2 presents our proposed conceptual framework for quantitative assessment of interactive multi-objective optimization. In particular, we propose the concept of machine DMs as a problem- and preference-independent framework for the simulation of realistic DMs. Section 3 describes a possible instantiation of our framework based on previous work from the EMO and MCDM literature. The goal of this instantiation is not only to serve as an example, but also to show how variations of the human non-idealities can have strong and surprising effects on the behavior of an interactive EMOA. Towards this goal, we present in Section 4 experimental results that confirm these effects. Finally, Section 5 discusses in more depth related work

[3] This is apparent from several discussions at the Dagstuhl series of seminars concerned with MCDM and EMO [1].

and the context of our proposal within the ongoing effort to combine MCDM and EMO approaches.

2 A conceptual framework of Machine Decision Making in Interactive Multi-objective Optimization

One of the difficulties when comparing interactive EMOAs is that competing algorithms not only differ in their interaction style and the preference models they can handle, but also on the underlying MOEA that ultimately provides the alternatives that are considered by the DM. Thus, it is difficult to assess whether any observable differences are due to one or another aspect, or their precise combination. The other major difficulty, as discussed above, is how to simulate a realistic DM in a way that allows us to understand the influence that human biases and other non-idealities have on the behavior of iEMOAs.

In order to overcome the above difficulties, we propose a conceptual framework for the quantitative analysis of iEMOAs. The architecture of our proposed framework is shown in Figure 1. It is composed of three main modules: a machine decision maker, an interactive module and an EMOA. Traditionally, research on interactive EMOAs has focused on the combination of the two latter components and considered the machine DM as a preference function, perhaps with some added noise.

Let us assume a multi-objective optimization problem with m objectives, and let $\boldsymbol{z} = (z_1, \ldots, z_m) \in \mathbb{R}^m$ represent an objective vector. Moreover, let us assume that there is an ideal preference function $U(\boldsymbol{z}) \in \mathbb{R}$ that must be maximized in order to satisfy the DM. Then, in our proposed framework, the machine DM simulates a *true preference function* $U(\cdot)$, but in addition it simulates several biases that distort the expression of the true preference function. As a result, the interaction module does not have access to the true preference function and instead it interacts (either directly or indirectly) with the resulting *imperfect preference function* $\hat{U}(\cdot)$. Another characteristic of the proposed framework is that the machine DM may also distort the true set of objectives (the *true criteria* $\boldsymbol{z} \in \mathbb{R}^m$) such that they are different from the criteria optimized by the EMOA (the *modeled criteria* $\hat{\boldsymbol{z}} \in \mathbb{R}^{m'}$). The rationale for this is that the way the DM sees the problem is not necessarily the way that the EMOA is able to optimize it. If both non-idealities are present, they are combined such that the imperfect preference function is evaluated on the modeled criteria.

A particular instantiation of each module can be defined by setting their parameters (the parameter layer in Fig. 1). An instantiation of the framework can then be applied to a given preference function and optimization problem, and the effect of different parameter settings can hopefully help us to identify general effects on the performance of various interactive EMOAs. Thus, the main goal of the proposed framework is to enable a factorial analysis, where the effect of the different modules can be analyzed separately and independently from a particular preference function and optimization problem being tackled.

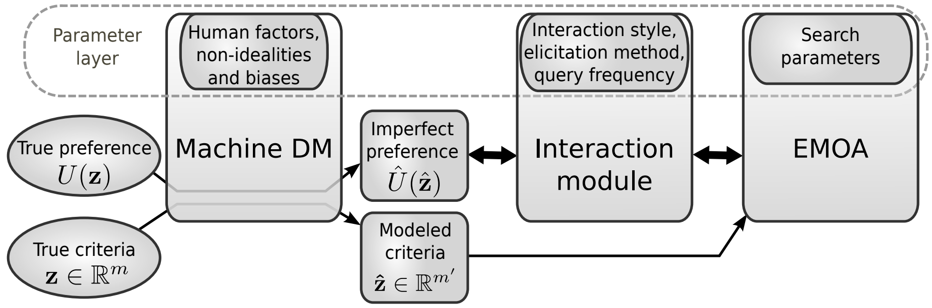

Fig. 1. A conceptual framework for the quantitative assessment of iEMOAs. The framework comprises three modules: the machine decision maker, the EMOA, and the interaction approach, which mediates between the other two. The parameter layer (top) encapsulates the settings of each module to be controlled, which enables experiments in a range of scenarios to be conducted. One goal of this framework is to generalize over the actual DM's preferences (*true preference*) and the actual optimization problem at hand (*true criteria*). Moreover, although the EMOA interacts with the DM via the interaction module, the problem being optimized by the EMOA (*modeled criteria*) is not necessarily identical to the problem seen by the machine DM (*true criteria*).

3 An example instantiation of our framework

Rather than proposing our own interactive EMOA (iEMOA) here, we chose to recast previous work as an instantiation of our proposed framework. This allows us to re-analyze results from the literature in a new light, while also showing how we can combine advanced iEMOAs with previous studies in MCDM.

The EMO Algorithm. The EMO algorithm used here is BC-EMOA [2], which is a variant of NSGA-II that learns the DM's preference function using support vector ranking (SVR).

The Interaction Module. BC-EMOA offers several alternatives for interaction. Here we consider the one recommended by the authors, which is based on periodically presenting a set of solutions to the DM, who must rank them according to her preferences. This means that, when the interaction module is active, the EMOA algorithm does not have direct access to the preference values.

The Machine DM. The machine DM used in this work follows the simulation of non-idealities proposed by Stewart [22]:

- **Omitted objectives.** From the m objectives considered by the DM, $q < m$ objectives are not known, that is, they are not modeled by the iEMOA. This may be due to a failure to identify relevant objectives when modeling the problem. The $m - q$ objectives ($\hat{z}_k$, $k = 1, \ldots, m - q$) that are known are selected randomly with a probability proportional to their true weights.

– **Mixed objectives.** The objectives modeled might actually correspond to the aggregation of two or more of the objectives internally considered by the DM. The machine DM simulates this "mixing" by making the $m - q$ known objectives a combination of two true objectives such that $\hat{z}_k = (1-\gamma)z_{c_k} + \gamma z_{c_{k+1}}$, where $\gamma \in [0, 1)$ is called the *mixing* parameter, and c_k is the position of objective k in the random selection described above (the position for not selected objectives does not matter).
– **Imperfect preference function.** Instead of the true preference function $U(\boldsymbol{z})$, the machine DM uses a transformed function $\hat{U}(\hat{\boldsymbol{z}})$, which is similar to the true preference function except that
  - the q objectives that are not modeled are always set to the same value,
  - and the addition of Gaussian noise $N(0, \sigma^2)$, where σ is a parameter of the machine DM.

4 Experimental Analysis

4.1 Experimental Setup

In order to assess the impact of the non-idealities and compare our results with those reported in the literature, we use similar parameters for BC-EMOA as in the original paper [2], that is, the algorithm runs for 500 generations, with a population of 100 individuals, and the standard crossover and mutation operators of NSGA-II. After the first 200 generations, the algorithm presents 10 solutions to the machine DM, who must rank them according to her preferences (either the true preference function or the imperfect one). This interaction occurs at most three times, with 20 generations between each interaction.

In the experimental analysis, we compare the best solution returned by three variants of the algorithm:

– **G**, or the *gold standard* variant [2], uses the true utility function $U(\boldsymbol{z})$ directly without any interaction. The MOEA optimizes the true criteria $\boldsymbol{z}$.
– **M**, or the *modeled* variant, uses the imperfect utility function $\hat{U}(\hat{\boldsymbol{z}})$ directly without any interaction. In addition, the MOEA uses the modeled criteria $\hat{\boldsymbol{z}}$ instead of the true criteria $\boldsymbol{z}$. When $q = 0$, $\gamma = 0$ and $\sigma = 0$, then $\hat{U}(\hat{\boldsymbol{z}})$ and $U(\boldsymbol{z})$ are equal, thus **G** and **M** are equivalent.
– **I**, or the *interactive* variant, interacts with the machine DM as described above. The machine DM uses the imperfect utility function to rank solutions and the MOEA optimizes the modeled criteria.

As for the machine DM parameters, we consider all combinations of the following settings: $m - q \in \{0, 1, 2\}$, $\gamma \in \{0.0, 0.05, 0.1, 0.2\}$ and $\sigma \in \{0.0, 0.05, 0.1, 0.2\}$, which are in the range of those considered by Stewart [22].

As a simple benchmark problem suite, we consider DTLZ1, DTLZ2, DTLZ6 and DTLZ7 [11] with $m = 5, 7$ objectives and $n = 2m$ variables. As for the true preference function, although in principle any arbitrary additive value function could be used [22], for simplicity, we consider here a machine DM with a linear scalarizing function and three different sets of randomly generated weights. The application of the machine DM described above to such a preference function is straightforward.

We repeat each run 10 times with different random seeds, but in order to reduce variance, we use the same set of 10 seeds for all the variants compared. Since some of the variants are equivalent, they will produce the same results.

We assess the results according to normalized true utility value $U(\cdot)$ of the most preferred solution found computed as follows. For a given problem and preference function, we compute the maximum and minimum values of the true preference function ever found over all runs of the algorithms, then we normalize the true preference values corresponding to the most preferred solution returned by each run to the interval $[0, 1]$, as $U'(\boldsymbol{z}) = \frac{(U(\boldsymbol{z}) - U_{\min})}{(U_{\max} - U_{\min})}$, such that 0 corresponds to the worst value and 1 to the best one.

4.2 Experimental Results

We present the results in terms of plots (e.g., Figure 2) of the mean U' for each of the three variants **G**, **M** and **I** described above. Each point corresponds to the mean value of the 10 runs for each of the three preference settings. The error bars denote a 95% confidence interval around the mean.

Figure 2 shows the results of omitting objectives (parameter q) without noise or mixing objectives for all problem instances with $m = 5$ (left) and $m = 7$ (right) objectives. When no objective is omitted ($q = 0$), the variants **G** and **M** are equivalent. As soon as we omit one or two objectives, there is a significant drop in the normalized preference value of the solutions found by **M** and **I**, which for $q = 2$ becomes more than two times worse for some problems. Interestingly, the difference between **I** and **M** often decreases for larger q. This is explained by the fact that by dropping objectives, learning the imperfect preference function becomes easier, despite the fact that the imperfect function is further away from the true one.

Figure 3 shows the results of mixing objectives, where the degree of mixing is controlled by parameter γ. The similar values obtained by **G** and **M** suggest that $\gamma = 0.1$ is not high enough to produce a noticeable effect. On the other hand, further increasing γ up to 0.2 produces an effect that is problem-dependent: In the case of DTLZ2, it makes more difficult for the interactive approach **I** to produce as good preferred solution as either **G** or **M**, whereas in the case of DTLZ7 and $m = 5$, the value $\gamma = 0.2$ produces the opposite effect.

Figure 4 shows the effect of adding noise to the utility function via parameter σ. With a sufficient high noise, we can observe that the variant using the imperfect preference (**M**) deteriorates noticeably with respect to the gold variant **G**. However, the effect on the actual interactive variant **I** is limited. In fact, for DTLZ1 and $m = 5$, a high σ does actually help **I** to approximate better the most preferred solution. Except for a few cases (DTLZ2 and DTZL7 with $m = 7$), the noise does not seem strong enough to produce any differences between **G** and **M**, and neither it has a remarkable effect on **I**.

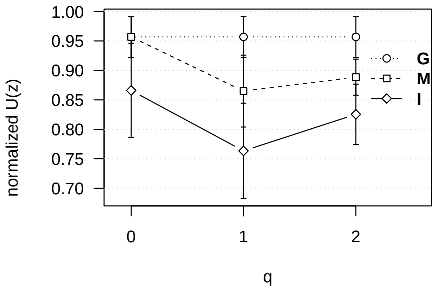

(a) DTLZ1, $m = 5$

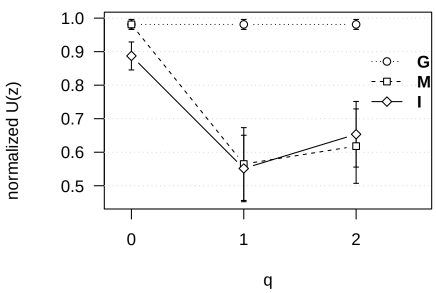

(b) DTLZ2, $m = 5$

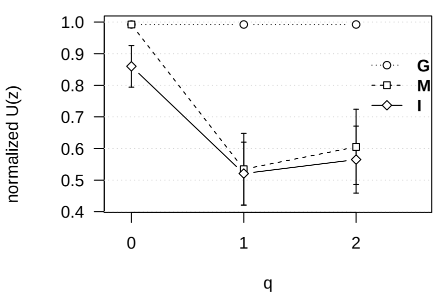

(c) DTLZ6, $m = 5$

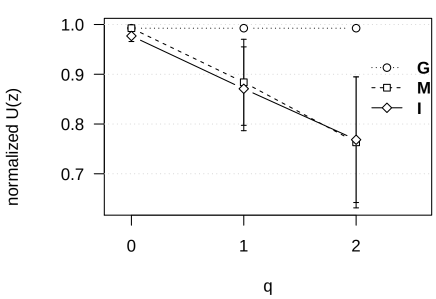

(d) DTLZ7, $m = 5$

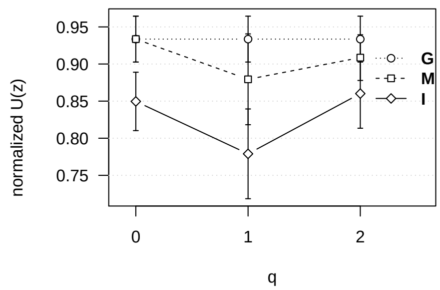

(e) DTLZ1, $m = 7$

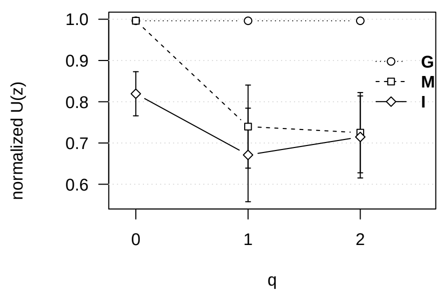

(f) DTLZ2, $m = 7$

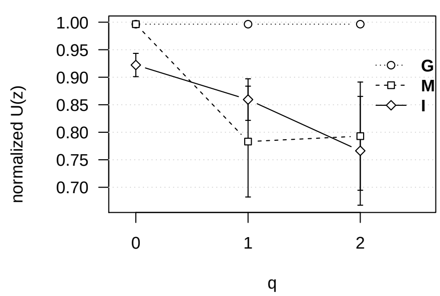

(g) DTLZ6, $m = 7$

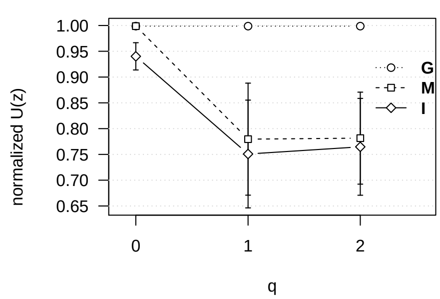

(h) DTLZ7, $m = 7$

Fig. 2. Results when omitting q out of m objectives without noise ($\sigma = 0.0$) nor mixing of objectives ($\gamma = 0.0$).

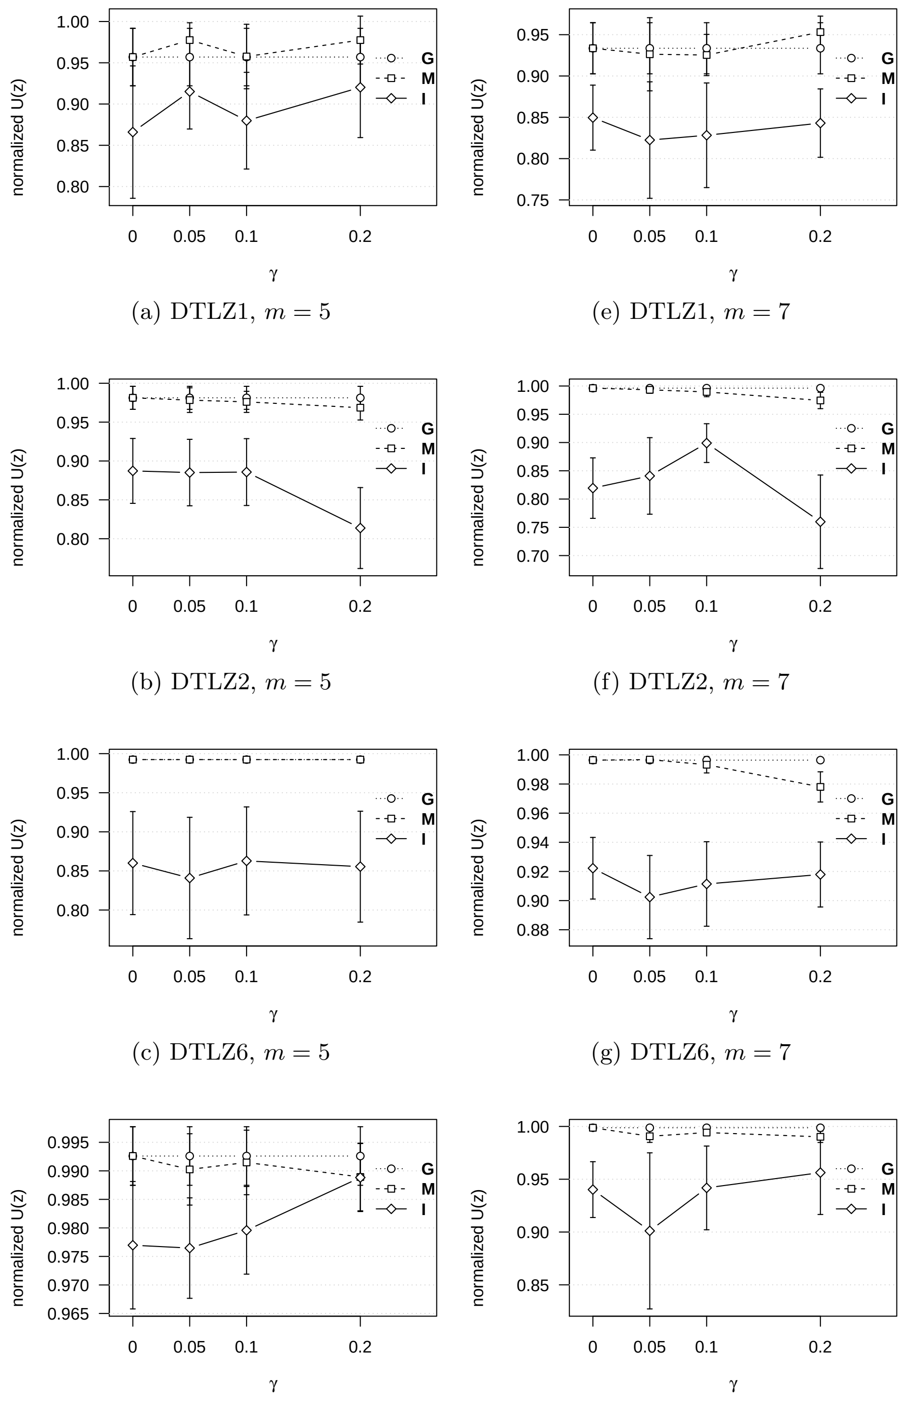

(a) DTLZ1, $m = 5$
(e) DTLZ1, $m = 7$
(b) DTLZ2, $m = 5$
(f) DTLZ2, $m = 7$
(c) DTLZ6, $m = 5$
(g) DTLZ6, $m = 7$
(d) DTLZ7, $m = 5$
(h) DTLZ7, $m = 7$

Fig. 3. Results when mixing objectives (parameter γ) without omitting objectives ($q = 0$) and without noise ($\sigma = 0.0$).

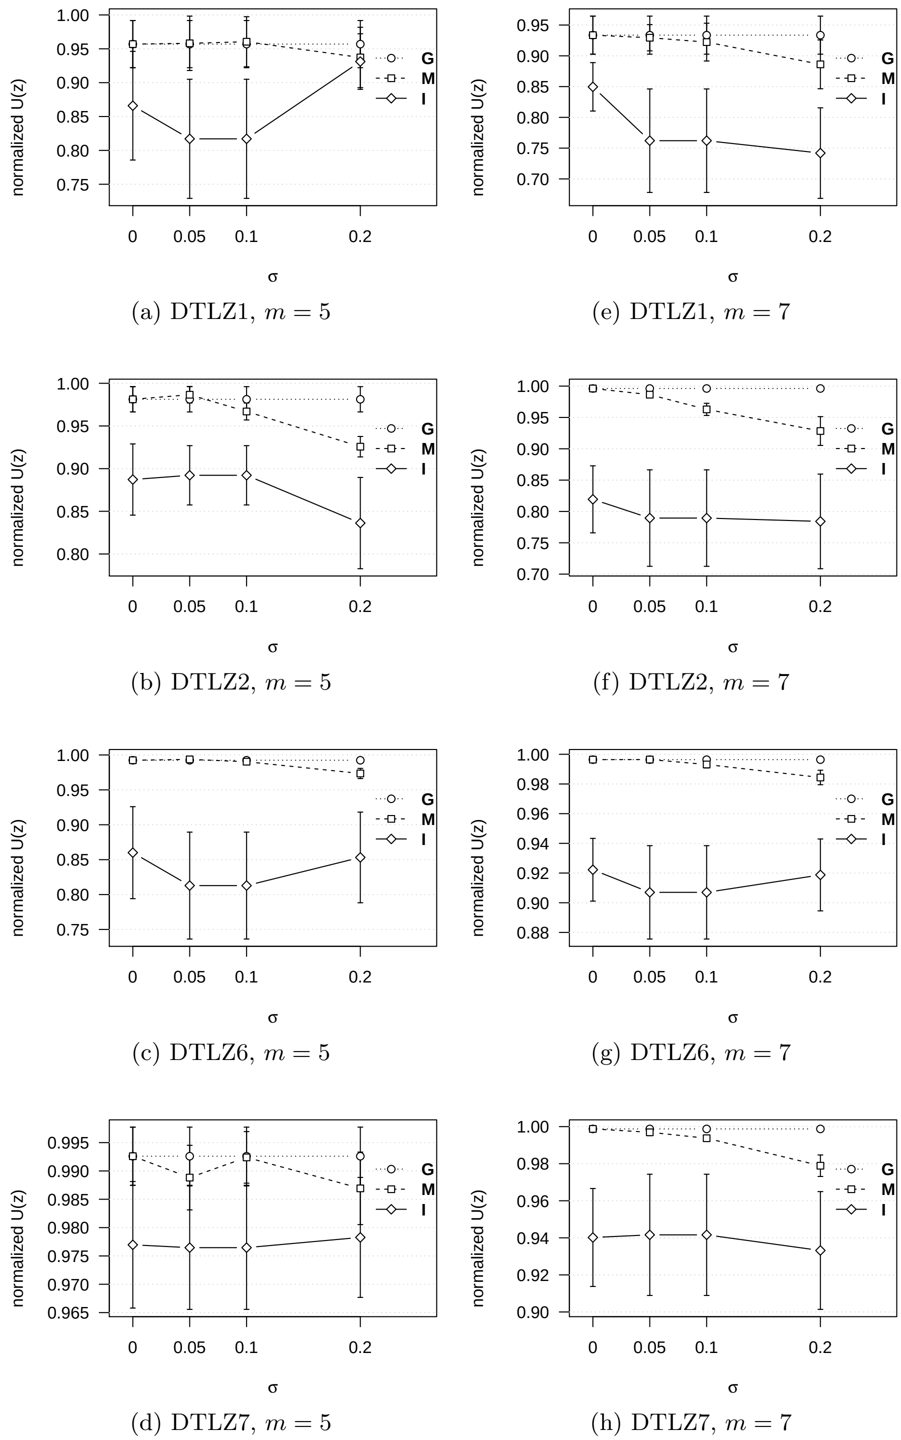

(a) DTLZ1, $m = 5$
(e) DTLZ1, $m = 7$
(b) DTLZ2, $m = 5$
(f) DTLZ2, $m = 7$
(c) DTLZ6, $m = 5$
(g) DTLZ6, $m = 7$
(d) DTLZ7, $m = 5$
(h) DTLZ7, $m = 7$

Fig. 4. Results when adding noise to the utility function (parameter σ) without omitting objectives ($q = 0$) and mixing them ($\gamma = 0.0$).

Finally, Figure 5 shows the combined effect of both noise ($\sigma = 0.2$) and mixing of objectives ($\gamma = 0.2$) when omitting $q = 0, 1, 2$ objectives. The results are mostly as expected, that is, the combined effect is a sharp decrease of the normalized utility with respect to the gold variant **G**. The other major difference with the corresponding plots in Fig. 2 is that, in this case, the interactive approach that learns the preferences (**I**) is sometimes better than the approach that directly uses the imperfect preference of the machine DM (**M**), e.g., in DTLZ2. In addition, the results for DTLZ7 when $q = 0$ suggest that this problem is quite easy for BC-EMOA. Yet, when omitting one or two objectives, there is an enormous degradation of the utility value, despite the fact that there is almost no difference between the values obtained by the interactive algorithm **I** and the variant using directly the imperfect preference **M**. Thus, this indicates that, although learning the preference function is relatively easy in this problem, even in the presence of noise or mixing of objectives, the omission of two objectives will lead the algorithm to a completely wrong answer with respect to the true preference.

5 Discussion and Related Work

Although nearly 15 years old, the review by Coello Coello [9] is worthy of attention for the number of attempts at merging methods from MCDM with EMO approaches already proposed by that time. Of course, in the intervening period there has been much further progress in EMO and in preference-based optimization and a proliferation of algorithms and interaction schemes. However, we believe the advancement in methods for the *assessment* of interactive EMO methods (in terms of their ability to satisfy a decision maker, or decision makers of different types) is less clear and there seems to be much less work in this direction.

As reviewed by Coello Coello [9], MCDM splits in broad terms into the French and American schools, and there is in general much difference of opinion about how preferences should be elicited, the benefits and vagaries of different schemes, and any number of difficulties associated with this task. This is all further compounded when one considers how to marry an MCDM approach eliciting and modelling preferences interactively (or otherwise) to an EMO algorithm, which generates a highly non-deterministic trajectory through the candidate solution space. But sidestepping this undoubted complexity, we can ask a simpler question. If a DM had a certain set of preferences and if they were expressible in some fixed way (as in the American school), how could we assess whether a particular modeling of this DM — the outcome of the method of preference elicitation — was successful? We could answer this simple question by agreeing to use not a human, but a machine DM for which we have full access to the underlying preference model. What makes it a *machine DM* simulating a human DM, and not merely a simple preference structure, is that we can also add into it more complex actions like biases, inconsistencies, learning and so forth (but all controllable by us). Using this approach, we can then evaluate how well the DM's

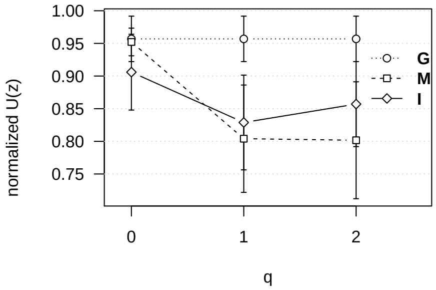

(a) DTLZ1, $m = 5$

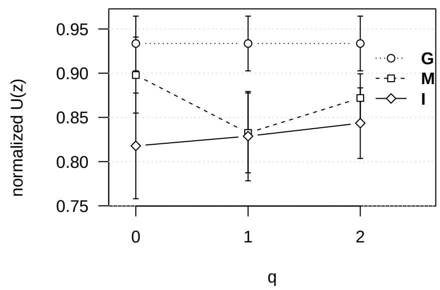

(e) DTLZ1, $m = 7$

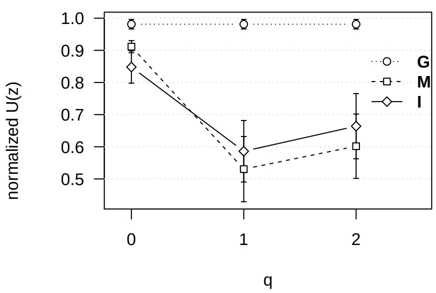

(b) DTLZ2, $m = 5$

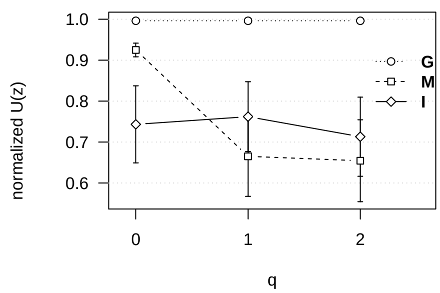

(f) DTLZ2, $m = 7$

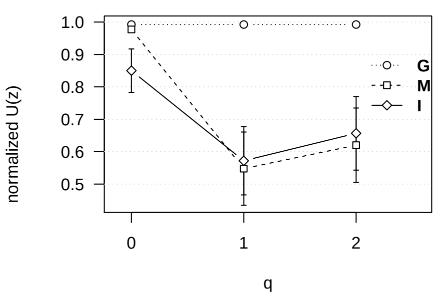

(c) DTLZ6, $m = 5$

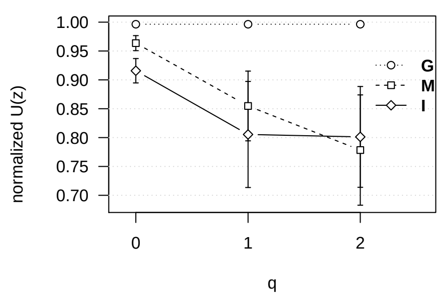

(g) DTLZ6, $m = 7$

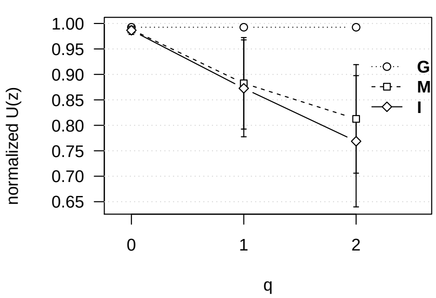

(d) DTLZ7, $m = 5$

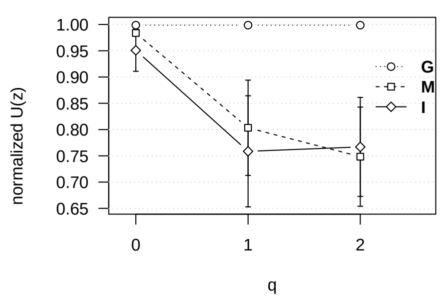

(h) DTLZ7, $m = 7$

Fig. 5. Results when omitting q out of m objectives with both noise ($\sigma = 0.2$) and mixing of objectives ($\gamma = 0.2$).

true preference is captured by the preference elicitation scheme, or in an EMO system, we can measure how well the search finds a solution that satisfies the DM's true preferences.

We are by no means the first to approach quantitative assessment in this way. Indeed, we were inspired in this work most strongly by previous work by Stewart [22,23,24], and based several aspects of our approach on it. In particular, modeling several non-idealities of real human decision making, Stewart [22] measured the robustness in the ranking of alternatives obtained when attempting to model the DM as a stable additively-independent (non-linear) value function. Working with simulations involving 7 criteria and 100 nondominated alternatives, a series of sensitivity analyses showed that the 'elicited' value functions worked well in preserving the 'true' ranking of solutions provided that (i) the value function was modeled as piecewise linear with a sufficient number of pieces (4 seemed sufficient for value functions derived from Prospect Theory [15]), (ii) the criteria are close to additively independent, and (iii) not too many criteria are missed out (missing one or two was not very detrimental).

While seminal as a simulation to quantify the robustness and limits of additive value functions, the paper (ibid.) does not consider either the elicitation process per se, or search (optimization). Thus the problem of evaluating iEMOAs is a good deal more involved. We hope that we have remained true to Stewart's aim of using (machine) DMs that mimic real human behaviours in important ways, while showing how we can begin to assess complete interactive EMO algorithms in ways that matter for the DM who would be using them.

The choice of iEMOA used in the study, BC-EMO, is somewhat arbitrary; although, the fact that it can handle up to five different types of preference information means that future analysis can compare the results presented here with those obtained with more complex preference models. Nonetheless, our framework is intended to be able to abstract away from any particular EMO algorithm, DM's preference model and preference elicitation technique. Thus, it would be very interesting to extend our analysis to other leading methods in interactive EMO such as [4, 8, 12, 13, 14, 20, 28] within the same framework, and this is our longer term goal. In particular, it would be important to study potential mismatch between methods that use reference points and a DM with preferences that do not rely on such, or vice versa. Evaluating very sophisticated approaches, like robust ordinal regression [13] or those based on machine learning [7], would shed light on whether these approaches really work in practice, and under what violations of DM assumptions do they also begin to exhibit poorer performance. And as EMO algorithms become more adept at handling many objectives, it will be increasingly important to put proposals for preference-based many-objective methods such as [28] to more stringent testing.

From a broader perspective still, the problem of eliciting human preferences while searching [25, 26] is not limited to mathematical optimization per se. A common scenario where human DMs are exercised on a daily basis are web or database searches for things such as books, holidays, and property. In these cases, human users both seem to learn about or construct their preferences through the interaction, and also learn the querying system. Work in this area (ibid.) seems advanced compared to work on iEMOAs in terms of the interaction systems proposed, and also concerning evaluation of them. We see a bright future for iEMOAs if and only if we can embrace a similar focus on the objective evaluation of our MCDM/EMO hybrids, including realistic and generalized DM models (machine DMs), and a more stringent testing regime that presents a variety of challenges to the working of these methods.

6 Conclusion

We have given an illustration how EMO/MCDM interactive approaches can be evaluated quantitatively using a conceptual framework based on machine DMs incorporating human-like non-idealities. Whilst iEMOAs are becoming more sophisticated, incorporating more advanced methods both from EMO and MCDM, it has remained unclear up to now how we can quantitatively assess the improvement. We proposed and demonstrated a few parameters of machine DMs for assessing iEMOA performance quantitatively. Importantly, the robustness of iEMOAs can be evaluated using this approach and we show with an example that existing iEMOAs are not very robust in the face of these lesser tested non-idealities. In particular, BC-EMOA is very sensitive to the omission of objectives and much less sensitive to their mixing. In fact, we observed that there are non-obvious interactions between various parameters of the machine DM. For instance, under some circumstances, the mixing of objectives may actually help the iEMOA to not get confused by an imperfect preference function. From these foundations, we hope to build a more comprehensive testing facility for the interactive EMOA community.

The experiments reported here are evidently preliminary. We plan in the future to extend the experiments to more complex preference functions. Given the variability of these preliminary results, carrying out a full factorial ANOVA would help to identify the most important factors. Future work will extend the framework of machine DMs to other types of interactions, such as goal programming [24] and aspiration-based techniques [23]. It is currently an open question how to extend the framework to incorporate a wider range of human behaviors and other non-idealities, in particular, the role of learning or evolving preferences by the DM. Finally, the ultimate goal of this framework should be to provide incentives and a way to benchmark interactive EMOAs able to cope with the complex behaviors of human DMs, possibly enabling at some point in the future the automatic design of interactive EMOAs [18].

Acknowledgments. We would like to thank the Research Center Schloss Dagstuhl[4] for hosting the "Learning in Multiobjective Optimization" seminar (id 12041). We also thank the authors of BC-EMOA for providing its source code. Manuel López-Ibáñez acknowledges support from the Belgian F.R.S.-FNRS, of which he is a postdoctoral researcher. Joshua Knowles thanks Theo Stewart for useful discussions.

[4] `http://www.dagstuhl.de/en`

References

1. Auger, A., Brockhoff, D., López-Ibáñez, M., Miettinen, K., Naujoks, B., Rudolph, G.: Which questions should be asked to find the most appropriate method for decision making and problem solving? (Working Group "Algorithm Design Methods"). In: Greco et al. [12], pp. 50–99
2. Battiti, R., Passerini, A.: Brain-computer evolutionary multiobjective optimization: A genetic algorithm adapting to the decision maker. IEEE Transactions on Evolutionary Computation 14(5), 671–687 (2010)
3. 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)
4. Branke, J., Greco, S., Słowiński, R., Zielniewicz, P.: Interactive evolutionary multiobjective optimization driven by robust ordinal regression. Bulletin of the Polish Academy of Sciences: Technical Sciences 58(3), 347–358 (2010)
5. 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.) PPSN 2012, Part I, Lecture Notes in Computer Science, vol. 7491, pp. 123–132. Springer, Heidelberg, Germany (2012)
6. Buchanan, J.T.: An experimental evaluation of interactive MCDM methods and the decision making process. Journal of the Operational Research Society 45(9), 1050–1059 (1994)
7. Campigotto, P., Passerini, A.: Adapting to a realistic decision maker: experiments towards a reactive multi-objective optimizer. In: Blum, C., Battiti, R. (eds.) Learning and Intelligent Optimization, 4th International Conference, LION 4, Lecture Notes in Computer Science, vol. 6073, pp. 338–341. Springer, Heidelberg, Germany (2010)
8. Chaudhuri, S., Deb, K.: An interactive evolutionary multi-objective optimization and decision making procedure. Applied Soft Computing 10(2), 496–511 (2010)
9. Coello Coello, C.A.: Handling preferences in evolutionary multiobjective optimization: A survey. In: Proceedings of the 2000 Congress on Evolutionary Computation (CEC'00), pp. 30–37. IEEE Press, Piscataway, NJ (Jul 2000)
10. Deb, K., Köksalan, M.: Guest editorial: Special issue on preference-based multiobjective evolutionary algorithms. IEEE Transactions on Evolutionary Computation 14(5), 669–670 (Oct 2010)
11. Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable test problems for evolutionary multiobjective optimization. In: Abraham, A., Jain, L., Goldberg, R. (eds.) Evolutionary Multiobjective Optimization, pp. 105–145. Advanced Information and Knowledge Processing, Springer, London, UK (Jan 2005)
12. Greco, S., Knowles, J.D., Miettinen, K., Zitzler, E. (eds.): Learning in Multiobjective Optimization (Dagstuhl Seminar 12041), Dagstuhl Reports, vol. 2, number 1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Germany (2012)
13. Greco, S., Matarazzo, B., Słowiński, R.: Interactive evolutionary multiobjective optimization using dominance-based rough set approach. In: Ishibuchi, H., et al. (eds.) Proceedings of the 2010 Congress on Evolutionary Computation (CEC 2010), pp. 1–8. IEEE Press, Piscataway, NJ (2010)
14. Greenwood, G.W., Hu, X., D'Ambrosio, J.G.: Fitness functions for multiple objective optimization problems: Combining preferences with Pareto rankings. In: Foundations of Genetic Algorithms (FOGA), pp. 437–455. Morgan Kaufmann Publishers, San Francisco, CA (1996)

15. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica 47(2), 263–291 (1979)
16. Köksalan, M., Karahan, I.: An interactive territory defining evolutionary algorithm: iTDEA. IEEE Transactions on Evolutionary Computation 14(5), 702–722 (Oct 2010)
17. Kornbluth, J.: Sequential multi-criterion decision making. Omega 13(6), 569–574 (1985)
18. López-Ibáñez, M., Stützle, T.: The automatic design of multi-objective ant colony optimization algorithms. IEEE Transactions on Evolutionary Computation 16(6), 861–875 (2012)
19. Morgan, P.D.: Simulation of an adaptive behavior mechanism in an expert decision-maker. IEEE Transactions on Systems, Man, and Cybernetics 23(1), 65–76 (1993)
20. Phelps, S., Köksalan, M.: An interactive evolutionary metaheuristic for multiobjective combinatorial optimization. Management Science 49(12), 1726–1738 (2003)
21. Purshouse, R.C., Deb, K., Mansor, M.M., Mostaghim, S., Wang, R.: A review of hybrid evolutionary multiple criteria decision making methods. COIN Report 2014005, Computational Optimization and Innovation (COIN) Laboratory, University of Michigan, USA (Jan 2014)
22. Stewart, T.J.: Robustness of additive value function methods in MCDM. Journal of Multi-Criteria Decision Analysis 5(4), 301–309 (1996)
23. Stewart, T.J.: Evaluation and refinement of aspiration-based methods in MCDM. European Journal of Operational Research 113(3), 643–652 (1999)
24. Stewart, T.J.: Goal programming and cognitive biases in decision-making. Journal of the Operational Research Society 56(10), 1166–1175 (2005)
25. Viappiani, P., Faltings, B., Pu, P.: Preference-based search using example-critiquing with suggestions. Journal of Artificial Intelligence Research 27, 465–503 (2006)
26. Viappiani, P., Pu, P., Faltings, B.: Preference-based search with adaptive recommendations. AI Communications 21(2), 155–175 (2008)
27. Wallenius, J.: Comparative evaluation of some interactive approaches to multicriterion optimization. Management Science 21(12), 1387–1396 (1975)
28. Wang, R., Purshouse, R.C., Fleming, P.J.: Preference-inspired coevolutionary algorithms for many-objective optimization. IEEE Transactions on Evolutionary Computation 17(4), 474–494 (2013)