

Multi-Objective Decision Making under Uncertainty and Incomplete Knowledge of Designer Preferences

2011-01-1080

Published
04/12/2011

Vijitashwa Pandey
Oakland Univ.

Efstratios Nikolaidis
Univ. of Toledo

Zissimos Mourelatos
Oakland Univ.

Copyright © 2011 SAE International
doi:10.4271/2011-01-1080

ABSTRACT

Multi-attribute decision making and multi-objective optimization complement each other. Often, while making design decisions involving multiple attributes, a Pareto front is generated using a multi-objective optimizer. The end user then chooses the optimal design from the Pareto front based on his/her preferences. This seemingly simple methodology requires sufficient modification if uncertainty is present. We explore two kinds of uncertainties in this paper: uncertainty in the decision variables which we call *inherent design problem (IDP) uncertainty* and that in knowledge of the preferences of the decision maker which we refer to as *preference assessment (PA) uncertainty*. From a purely utility theory perspective a rational decision maker maximizes his or her expected multi attribute utility. We show how this is inherently inconsistent with providing the decision maker with alternatives on the Pareto Front unless the decision maker trades off attributes or some function thereof linearly. In this paper we propose a methodology, rooted in a set of axioms that can be used in conjunction with a modified Pareto Front to select the best design. We present our findings on a simple optimization problem involving two attributes.

Definition of terms

Inherent Design Problem (IDP) Uncertainty- Inherent design problem uncertainty because of variability in decision variables and/or noise in objective function evaluation.

Preference Assessment (PA) Uncertainty- Uncertainty in the knowledge of the decision maker's utility function because of imperfect assessment.

1. INTRODUCTION

Engineering decisions routinely involve tradeoff between multiple attributes and uncertainty. To make decisions consistent with mathematics and decision maker preferences, many methods have been proposed. Decision analysis, a normative field, aims to build techniques that can help make decisions consistent with decision maker preferences (Howard, 1988). It is argued that if the preferences are captured perfectly, decision analysis can provide the best decision to make. In an engineering design paradigm this involves knowing the attributes of concern of the decision maker, their tradeoff behavior over these attributes and also the risk attitude so that the process of finding the optimal design can be automated.

Von Neumann and Morgenstern (1947) proposed that a rational decision maker maximizes his or her expected utility while making a decision under uncertainty. This translates to expected multiattribute utility (MAU) when more than one attribute are present. Since engineering design essentially involves a sequence of decisions, this notion can be applied directly. Thurston (1991) proposed that engineering design should be looked at as a decision making process where the designer maximizes the end user multi attribute utility. In recent years this approach to engineering design has seen a




Figure 1. Design decision making with possible use of Pareto front.

lot support. Hazelrigg (1998) provided a framework for decision based design. Lewis et al (2006) provides a comprehensive background of decision making in engineering design.

Once the preferences of the end user/decision maker are known in mathematical terms i.e. an explicit utility function, an optimization problem can be formulated to maximize this utility of the decision maker. If the utility function is not known *a priori* a Pareto front can be generated (which is a set of non-dominated solutions) and the decision maker is asked to choose one or more options on this front. There are many algorithms developed in optimization that help find a set of non-dominated solutions from which a decision maker can choose one option. Numerous engineering design problems have been solved using this approach (Hula et al. 2003, Jolly et al. 2006, Deb and Jain, 2003).

Uncertainty presents itself in all facets of engineering design. Figure 1 schematically shows an engineering system that takes in inputs (design variables) and converts them into objectives that the end user cares about. The utility function of the decision maker combines these objective values into an abstract notion of satisfaction. It is also shown how uncertainty affects different aspects of design decision making. There are two kinds of uncertainty we identify for this work:

1. Inherent Design Problem (IDP) Uncertainty: This is the uncertainty that is present because of variability in the realization of the design variables or the objectives as shown in figure 1. Since the design variables are non-deterministic, handling this kind of uncertainty translates to controlling the statistics of the design variables, the most common one being the mean. Since the objective function becomes probabilistic as well, only its statistics can be controlled. The objective

function evaluation itself being noisy simply modifies the probability density of the objectives.

2. Preference Assessment (PA) Uncertainty: There are various reasons identified in literature that impede modeling the preference of the decision maker perfectly. This could be because - (a) preference structure has not formed in the decision maker's mind (Macdonald et al. 2009), (b) preferences are acquired using a survey (stated preferences which can differ from actual decisions) (c) preferences are acquired by observing the decision maker's actions (revealed preference which need to be extrapolated if the decision problem changes). Each of these adds uncertainty in some way or the other to our knowledge of decision maker preferences (i.e. MAU function).

IDP uncertainty in engineering design is usually handled by designing around the mean values of the decision variables, for example in Reliability Based Design Optimization. Mean values of the objectives are optimized while constraints become probabilistic i.e. a pre-specified low probability of their violation is contended with. This method has the serious shortcoming that the risk attitude of the end user is not explicitly taken into account. On the other hand, in a decision based design framework, if the multiattribute utility function of the end user is known, then its expected value is maximized. Constraint handling is usually accomplished by defining the value of the MAU function when a constraint is violated. For example, if a design is infeasible, the MAU value is assumed to be zero. The risk attitude of the end user is explicitly incorporated into the MAU function.

Post-optimal tradeoffs will by definition be precluded if the MAU function is perfectly known *a priori*. Since this is hard to accomplish, based on our arguments above, we need a method that still allows us to make decisions. The authors




Figure 2. Isopreference curves over two attributes for a decision maker

believe that Post-Optimal tradeoffs have value even when the end user knows his/her MAU. The reason is that Post-Optimal tradeoffs allows for the decision maker to choose an option among many that more or less meet the criteria. Before we proceed we create a set of four axioms that a design decision making method proposed should satisfy:

1. It should be consistent with multiattribute utility theory in the limiting case when there is no uncertainty in the knowledge about decision maker preferences
2. It should account for the risk attitude of the end user. The optimal solution should change with the decision maker or his/her risk attitude
3. It should handle both Inherent Design Problem as well as Preference Assessment uncertainties
4. It allows for Post-optimal tradeoffs, i.e. at the end of the optimization process enables the end user to choose between many alternatives.

This paper is organized as follows. In the next section we discuss the basics of Isopreference curves, Pareto front and Multi attribute utility functions. We also discuss the difficulties encountered in assessment of end user preferences. In section 3 we explore idea of Pareto fronts under uncertainty and prove that the only way to make Pareto fronts consistent with multi attribute utility theory is if the decision maker trades the utilities from individual attributes linearly. In section 4 we present a methodology that satisfies the axioms presented before and also does not require separability of the MAU function into single attribute utility functions. Finally we demonstrate our approach using a short mathematical example in section 5 while section 6 concludes.

2. ISOPREFERENCE CURVES AND THE PARETO FRONT

In the attribute space, an isopreference curve is a locus of points along which a decision maker is indifferent to different

choices. This captures his/her tradeoff behavior for the attributes in question. [Figure 2](#) presents a set of isopreference curves for a hypothetical decision maker. There are two attributes in question and the decision maker is trying to minimize them. Two isopreference curves are shown, curve *a* and curve *b*. Any point on *a* has the same value or utility to the decision maker as any other point. For example, Options 1 and 2 have the same utility. All points on curve *b* have the same utility. However, by moving from one curve to another the utility value changes. In the figure, since origin is the best possible option, any point on curve *a* is better than any point on curve *b*, since it is closer to the origin.

A Pareto front captures the constraint posed by the decision problem itself. It is a collection of points for which simultaneous improvement in all attributes is not possible for the given decision problem. Any two points on the Pareto front show a different tradeoff between attributes as imposed by the decision problem itself. Pareto fronts are generated using multiobjective optimization techniques. Some of the most common methods employ genetic algorithms, for example the Non-dominated Sorting Genetic Algorithms ([Deb and Srinivasan, 1995](#) and [Deb et al. 2002](#)). [Figure 3](#) shows an example of Pareto front. If we take any two points on the Pareto front, one of them improves one attribute over the other point by sacrificing the other attribute. Since a Pareto front is inherent to the decision problem itself, it is not possible to realize a feasible point to the left of or below any point on the Pareto front. When no PA uncertainty is present, Pareto fronts are not needed because we can simply maximize the utility or the expected utility. In the case of no IDP uncertainty Pareto fronts are useful because they help make post-optimal tradeoffs.




Figure 3. Pareto Front over two attributes for a decision maker

2.1. MULTIATTRIBUTE UTILITY FUNCTIONS AND A RATIONAL DECISION MAKER

MAU functions model the tradeoff behavior of the decision maker over the various attributes in addition to the decision maker's risk attitude. In other words, when comparing two options a decision maker maximizes his or her MAU or the expected MAU. There are many ways one can construct MAU functions, the most common method uses single attribute utility functions. These are combined into a MAU function if certain independence conditions are met ([Keeney and Raiffa, 1994](#)). If the independence conditions are not met, a value function can first be created using indifferent points and a MAU function defined over the value function to take into account the risk attitude. The reader is referred to [Keeney and Raiffa \(1994\)](#) or [Clemen and Reilly \(1997\)](#) for a complete treatment of this subject.

THE DILEMMA OF CHOICE: STATED VS. REVEALED PREFERENCES

Some researchers argue that preference structure is not even formed in a decision maker's mind until a choice is presented ([MacDonald et al., 2009](#)). There is also a lot of literature in engineering and the sciences that addresses the issue of differences in the utility functions of decision makers elicited through questions and surveys (stated preferences) and those acquired by looking at actual decisions they made (revealed preferences) ([Wassenaar et al., 2005](#), [Mark and Swait, 2004](#)). Each method has its advantages and disadvantages. Stated preferences are relatively easier to acquire and can be used to find utility of outcomes that may or may not exist. There is no guarantee, however, that the decision maker will be consistent with his or her stated preferences when making actual decisions. Revealed preferences need to be acquired by

collecting and analyzing data from past decisions. These reflect true decision making behavior but cannot be used in situations where decision scenarios change significantly. An example where revealed preferences are of limited use is when a radically new design is to be evaluated for which no prior preference data exists.

It is clear that regardless of the method used to determine the MAU function there is going to be significant uncertainty (PA) regarding how it truly reflects the decision maker preferences. It is therefore important to come up with a design decision making methodology that incorporates this uncertainty.

PROPOSED DESIGN METHODOLOGY THAT ACCOUNTS FOR BOTH IDP AND PA UNCERTAINTIES

As discussed before, if the multiattribute utility function of the decision maker is known, maximizing its expected value yields the optimal design. If the decision problem is deterministic, Pareto fronts can be used to offer the decision maker alternatives without explicitly modeling their utility function. What should we do when neither is true?

Since we cannot always know the decision maker's utility function completely, they should be presented with different design options and not just one design. We also argue that offering different designs has a lot of value even when there is no PA uncertainty. In the next subsection we explore the concept of Pareto front under uncertainty.

PARETO FRONT UNDER UNCERTAINTY (BOTH IDP AND PA)

Defining the Pareto front under uncertainty is not an easy task. This is because different realizations of designs, may occasionally force some points to not lie on the Pareto front. On the other hand, if a representative Pareto front is used, it should take into account the uncertainty, while still allowing the decision maker to have options and thus make tradeoff decisions between nearly optimal options. If the axes on a Pareto Front are the attributes themselves then we would have to plot some statistic(s) of them, which would imply a certain risk attitude which may differ from the decision maker's true attitude. For example, if the mean values of the attributes are plotted then we are assuming a risk neutral decision maker. One way to incorporate the true risk attitude is to plot single attribute utility functions on the Pareto front¹. A decision maker always prefers higher value of an attribute's utility over a lower value, regardless of the value of the other objective and its corresponding utility; i.e. the preference order does not change, implying that we can have a Pareto front. One can plot the expected utilities as shown in [figure 4](#).




Figure 4. Pareto front over individual utilities. Negative signs are used to make it a minimization problem.

If we argue that the tradeoff behavior of the decision maker can be fully represented on the plots, i.e. the decision maker can choose the best point on the Pareto front given above, then we are implying that the multiattribute utility function, U_m is separable in U_1 and U_2 , because for each point, the decision maker is able to find his/her multi attribute utility. We have in this case,

$$U_m(f_1(x), f_2(x)) = G(U_1(f_1(x)), U_2(f_2(x))) \quad (1)$$

where G is a function that maps single attribute utility functions to the multi attribute utility. Therefore we have under uncertainty

$$E[U_m(f_1(x), f_2(x))] = E[G(U_1(f_1(x)), U_2(f_2(x)))] \quad (2)$$

But if uncertainty has been taken into account in [figure 4](#) by the representative Pareto front, the decision maker should be able to make a decision (choose a design) on the front. A rational decision maker maximizes his expected multiattribute utility to make decisions. This means that the *expected* MAU can be found for each point on the Pareto front as well; we have:

$$E[G(U_1(f_1(x)), U_2(f_2(x)))] = K(E[U_1(f_1(x))], E[U_2(f_2(x))]) \quad (3)$$

where K is a function of the expected utilities of attributes $f_1(x)$ and $f_2(x)$. It is straightforward to show that K is identical to G . Using Taylor expansion of G about the mean values of the two attributes we get,

$$\begin{aligned} G(U_1(f_1(x)), U_2(f_2(x))) &= G(E[U_1(f_1(x))], E[U_2(f_2(x))]) + \\ &\quad \frac{\partial G[U_1(f_1(x))]}{\partial U_1(f_1(x))} [U_1(f_1(x)) - E[U_1(f_1(x))]] \\ &\quad + \frac{\partial G[U_2(f_2(x))]}{\partial U_2(f_2(x))} [U_2(f_2(x)) - E[U_2(f_2(x))]] + \\ &\quad \frac{1}{2} \cdot \frac{\partial^2 G[U_1(f_1(x))]}{\partial U_1(f_1(x))^2} [U_1(f_1(x)) - E[U_1(f_1(x))]]^2 + \dots \end{aligned} \quad (4)$$

where the derivatives are evaluated at the mean values of the utilities of the attributes. The LHS of [Equation 4](#) should depend only on the expected utilities of the two attributes, because of the RHS. Taking the expectation of both sides of [Equation \(4\)](#) and substituting in [Equation \(3\)](#), we get:

$$\begin{aligned} G(E[U_1(f_1(x))], E[U_2(f_2(x))]) &+ \frac{1}{2} \cdot \\ \frac{\partial^2 G[U_1(f_1(x))]}{\partial U_1(f_1(x))^2} Var(U_1(f_1(x))) &+ \dots = K(E[U_1(f_1(x))], E[U_2(f_2(x))]) \end{aligned} \quad (5)$$

Since the RHS of [Equation \(5\)](#) necessitates that the LHS be independent of the higher moments of the U_1 and U_2 , it implies that the derivatives higher than one are zero. Thus, we have:

$$E[G(U_1(f_1(x)), U_2(f_2(x)))] = G(E[U_1(f_1(x))], E[U_2(f_2(x))]) = K(E[U_1(f_1(x))], E[U_2(f_2(x))]) \quad (6)$$

which means that K is the same as G . The left equality means that the mean value of the function G is equal to its value

¹We could also plot certainty equivalents in the attribute space as the Pareto front. The treatment of the topic will not change.