Solution generation with qualitative models of preferences

Boi Faltings
Swiss Federal Institute of Technology in Lausanne (EPFL)
Artificial Intelligence Laboratory (LIA)
CH-1015, Lausanne, Switzerland

Marc Torrens
i:FAO's Future Lab. and LIA, EPFL
PSE-C, CH-1015, Lausanne, Switzerland

Pearl Pu
Swiss Federal Institute of Technology in Lausanne (EPFL)
Human Computer Interaction Group
CH-1015, Lausanne, Switzerland

Abstract

We consider automated decision aids that help users select the best solution from a large set of options. For such tools to successfully accomplish their task, eliciting and representing users' decision preferences is a crucial task. It is usually too complex to get a complete and accurate model of their preferences, especially regarding the tradeoffs between different criteria.

We consider decision aid tools where users specify their preferences *qualitatively*: they are only able to state the criteria they consider, but not the precise numerical utility functions. For each criterion, the tool provides a standardized numerical function that is fixed and identical for all users and used to compare solutions. To compensate for the imprecision of this qualitative model, we let the user choose among a *displayed set* of possibilities rather than a single optimal solution.

We consider the probability of finding the most preferred solution as a function of the number of displayed possibilities and the number of preferences. We present a probabilistic analysis, empirical validation on randomly generated configuration problems and a commercial application. We provide mathematical principles for the design of the selection mechanism, guaranteeing that users are able to find the target solution.

1 Introduction

Many real-world applications require people to select the most preferred item, which we call the *target solution*, from a set of options. For example, an electronic catalog system might provide access to millions of products, and the user has to navigate the catalog to find the most preferred one. Decision theory provides algorithms which guarantee to find the target solution given the options and an accurate preference model.

In classical decision theory ([8]), solutions are characterized by a set of criteria. Users express their preferences by the relative weights they give to criteria or combinations of criteria. Such approaches are described for example in [9], and methods for eliciting the weights are discussed in more detail in [6, 7]. An example of its application in a decision aid system is the PC selection tool of IBM described in [17], where users can adjust the weights of different criteria and interactively see how different PC models rank according to these weights.

However, there are many practical situations where the space of possible criteria is so large that stating a numerical preference model for all of them would be cognitively impossible. This is the case for example in travel, where there is a large number of possible criteria involving times, means of transportation, locations, that vary from one user to another. In such a case, it is more feasible to characterize a user's preference *qualitatively* as the specific combination of criteria that apply, without also eliciting a precise numerical model. Such *qualitative* decision theory has recently become the subject of increased interest in the research community ([5, 3, 18]).

One can interpret qualitative decision theory as a form of classical decision theory where the system provides standardized numerical functions that are set to fit an average user. The qualitative model is thus imprecise and will in general not correctly identify the most preferred solution. We consider approaches that compensate for this shortcoming by letting the user pick the most preferred solution from a larger *displayed set* D of k solutions. Such an approach has been taken in numerous practical systems, for example in [10, 11, 4, 14, 20, 12]. The process will be *sound*, i.e. allow the user to find the target solution, only if the displayed set actually contains the target solution. We will show that this property heavily depends on the model used for selecting solutions for display as well as on the number of stated preferences.

Depending on the characteristics of the application, different methods for selecting displayed solutions may be appropriate. In this paper, we consider:

- *dominance filters* that display k non-dominated (Pareto-optimal) solutions;
- *utilitarian filters* that display the k solutions with the lowest sum of preference violations;
- *egalitarian filters* that display the k solutions with the smallest maximum preference violation.

Using both a probabilistic model and empirical examples, we show that none of the filters can guarantee soundness in all cases, and that all have to be carefully tuned to the application domain.

1.1 Assumptions and Definitions

We assume that solutions are characterized by a fixed (but possibly very large) set of *attributes*, and that users formulate *preferences* in the form of penalty functions formulated on the attributes, defined as follows:

Definition 1. *Every configuration solution is characterized by a finite set of n attributes $A = \{a_1, ..., a_n\}$. Every attribute can take values from a fixed domain $\{d_1, ..., d_n\}$.*

A preference $p_i(a_k)$ *is a penalty function $d_k \rightarrow \Re$ from an attribute a_k to a number that gives the penalty of that attribute value to the user. We assume that the smallest values are the most preferred ones.*

Considering that a preference p_i always applies to the same attribute a_j, we simplify the notation and write $p_i(S)$ for $p_i(a_j(S))$.

We assume that there is a predefined and fixed set of parameterized penalty functions that can be used to construct preferences.

Note that from the user's perspective, the preferences are *qualitative* since the numerical details of preference functions and their relative weights cannot be chosen, but are identical for all users.

Example 1: Solution attributes and penalties.
Consider the example of planning a trip consisting of two flights f_1 and f_2 that connect at a transfer airport.

- A solution could be characterized by the following attributes:
  $a_0 =$ `departure_time` (f_1)
  $a_1 =$ `arrival_time` (f_2)
  $a_2 =$ `departure_time` (f_2) `-` `arrival_time` (f_1) : transit time
  $a_3 =$ `transfer_airport` (f_1, f_2)
  $a_4 =$ `arrival_time` (f_2) `-` `departure_time` (f_1) : travel time
- The system could provide the following penalty functions for specifying preferences:
  $latest(time, a_i) = max(0, a_i - time)$
  $earliest(time, a_i) = max(0, time - a_i)$
  $min - duration(length, a_i) = max(0, length - a_i)$
  $max - duration(length, a_i) = max(0, a_i - length)$
  $rule - out(value, a_i) = 1$ if $a_i = value$, 0 otherwise
- The user could use these to specify the following qualitative preferences:
  $latest(a_1, 18:00) = max(0, a_1 - 18:00)$
  $min - duration(2:00, a_2) = max(0, 2:00 - a_2)$
  $rule - out(London, a_3) = 1$ if $a_3 = London$, 0 otherwise

which translate into numeric penalty functions.

We assume that the preferences of the user are given by a set of penalty functions $P^* = \{p_1^*, .., p_d^*\}$. This is expressed qualitatively by a choice of corresponding standard penalty functions $P = \{p_1, .., p_d\}$ which are accessible to the decision aid. We assume throughout this paper that the penalty functions correctly express *ceteris paribus* preferences, i.e. that for any pair of solutions S_1 and S_2 that are identical in all preferences except preference p_i, the user prefers S_1 over S_2 if and only if $p_i(S_1) < p_i(S_2)$.

In some analyses, we assume furthermore that the user follows a particular model for combining preferences:

- in the utilitarian model, the user combines the preferences by forming a linear sum, i.e. the cost $C(S) = \sum_{p_i^* \in P^*} p_i^*(S)$. Solution S_1 is preferred over solution S_2 whenever it has a lower cost, i.e. $C(S_1) < C(S_2)$.
- in the egalitarian model, the user combines the preferences by considering the maximum penalty, i.e. the function $F(S) = max_{p_i^* \in P^*} p_i^*(S)$. Again, Solution S_1 is preferred over solution S_2 whenever it has a lower maximum penalty, i.e. $C(S_1) < C(S_2)$.

In both of these preference models, we assume that the standardized preference functions are normalized for average users so that the inaccuracy of p_i with respect to p_i^* is bounded by ϵ:

$$(1-\epsilon)p_i \leq p_i^* \leq (1+\epsilon)p_i$$

This inaccuracy reflects different degrees of importance that users attach to violations of a preference. Note that we assume that each normalized penalty function still correctly reflects the qualitative ordering of solutions when all other criteria are equal.

We call the best solutions for the user's true preference model P^* his *target solution* S_t.

2 Decision support with qualitative preference models

As the qualitative preference model is imprecise, we cannot be sure that it will select the target solution. To compensate for this shortcoming, we propose an approach where we generate a *displayed set* D of k target solutions and let the user pick himself from this set.

For this approach to be practical and successful, the following conditions must be satisfied:

1. it must be possible to limit the computed set D to a size of at most k displayed solutions so that it can be displayed in a consistent manner.
2. solutions that are Pareto-optimal within the set D must also be Pareto-optimal with respect to the set of all feasible solutions. This is important

to keep the user from unknowingly picking a dominated solution as the final choice.

3. when the user has stated all of his preferences, the target (user's most preferred) solution must be included in the displayed solutions.

We now examine in detail the three candidate techniques: dominance, utilitarian and egalitarian filters. In order to allow a theoretical comparison, we assume that preferences p_i are independent, with real-numbered values in the interval [0..1], and that m solutions are distributed uniformly in the $|P^*|$-dimensional space of preference combinations. While in reality both assumptions are not likely to hold perfectly, comparing the theoretical results to measurements on real-world configuration problems shows a very good match.

3 Dominance filter

In the dominance filter, we choose as the displayed set a set of k *Pareto-optimal* solutions, defined as follows:

Definition 2. *A solution S_d is* dominated *with respect to P iff there is another solution S' such that for all $p_i \in P$, $p_i(S_d) \geq p_i(S')$ and at least one $p_j \in P$, $p_j(S_d) > p_i(S')$. We write $S_d \prec S'$.*

A solution S_p is Pareto-optimal *iff it is not dominated.*

In Figure 1, the Pareto-optimal set is $\{1, 3, 4, 6\}$, as solution 7 is dominated by 4 and 6, 5 is dominated by 3 and 4, and 2 and 8 are dominated by 1.

From the point of view of decision theory, the dominance filter is the most appropriate, since this does not require any assumption about the weights that are not precisely known.

As shown in Figure 1, a solution S_j with preference values $p_1 = c_1, ..., p_d = c_d$ is dominated by any solution that falls within the subspace $p_1 \in [0..c_1], ..., p_d \in [0..c_d]$, which is a hypercube. Given a uniform distribution of solutions, the average probability that a solution S_j is dominated by another solution S_l is thus equal to the probability that S_l falls into that subspace. This probability is given by the proportion of the subspace with respect to the entire space. Since we assume that all attributes vary between 0 and 1, the size of the entire space is $1^d = 1$. Thus, the probability is just the size of the subspace, i.e.:

$$Pr(S_j \prec S_l) = \prod_{p_i \in P} p_i(S_j)$$

and the probability that a solution S_j is Pareto-optimal is the probability that in the $m-1$ other solutions, not a single one dominates S_j:

$$Pr(S_j \text{ is Pareto-optimal}) = \left(1 - \prod_{p_i \in P} p_i(S_j)\right)^{m-1}$$

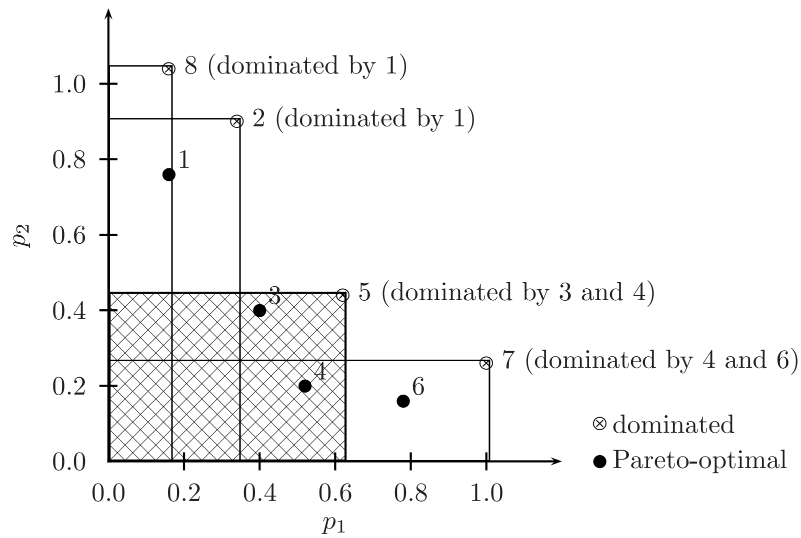

Fig. 1: Example of solutions with two penalties. The two coordinates show the values of preferences p_1 (horizontal) and p_2 (vertical). Rectangles show where dominating solutions must lie. For example, solutions dominating solution 5 must fall into the hatched rectangle.

The expected number of Pareto-optimal solutions is given by integrating this probability over the possible value combinations of the preferences, thus computing the expected size of the set $O_p = \{S_j | S_j \text{ is Pareto-optimal}\}$:

$$E[|O_p|] = m \int_0^1 \cdots \int_0^1 \left(1 - \prod_{p_i \in P} p_i\right)^{m-1} dp_1 \cdots dp_d$$

where $d = |P|$ is the number of preferences in the set P.

Figure 2 shows graphs of this number for two example scenarios that match experiments we made with randomly generated configuration examples[1] shown in Figure 3. The scenarios involve either $m = 778$ and $m = 6,444$, and the number of preferences d ranges from 3 to 12. In the experiments, we can observe that the randomly generated preferences are not completely independent of one another: for example, the number for 12 preferences in the experiments are already reached with 8 preferences in the theoretical analysis, and similarly 8 preferences in the experiments correspond to about 5 preferences in the theory. However, besides this effect, there is a very good match, showing that the model is quite realistic.

A consequence of this rapid rise in the number of Pareto-optimal solutions is that when there are few preferences, there may not be enough solutions to fill the displayed set, while with too many preferences, there will be too many

[1] The generation of the randomly generated configuration problems is described in Section 4.

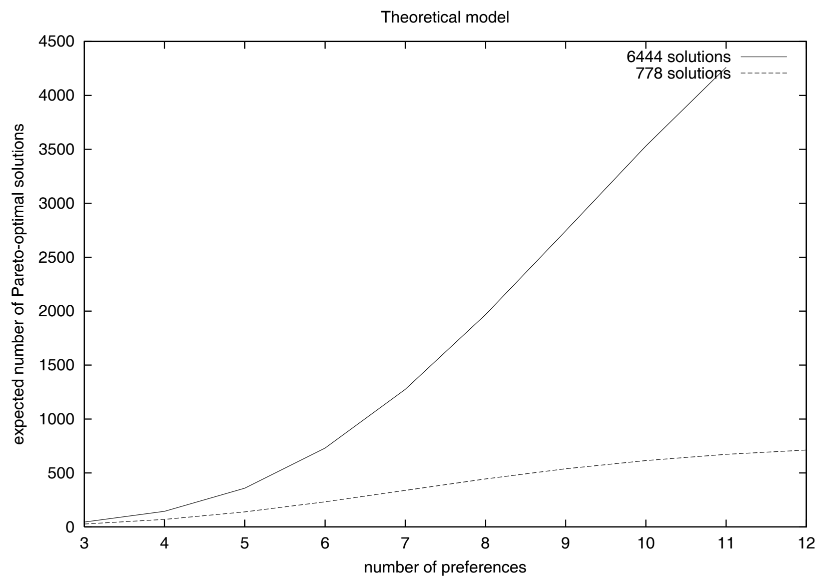

Fig. 2: Number of Pareto-optimal solutions expected from the probabilistic analysis for 2 example scenarios.

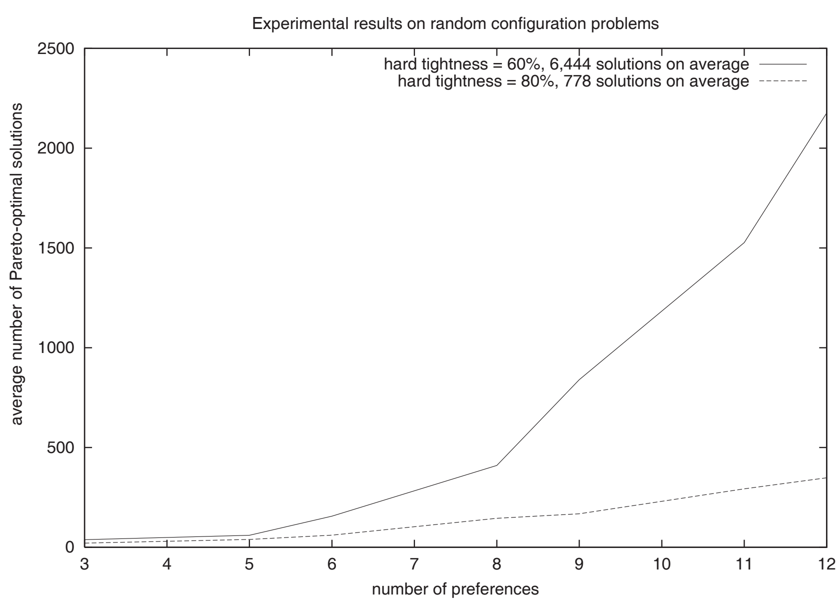

Fig. 3: Average number of Pareto-optimal solutions actually observed on randomly generated configuration examples in the two example scenarios.

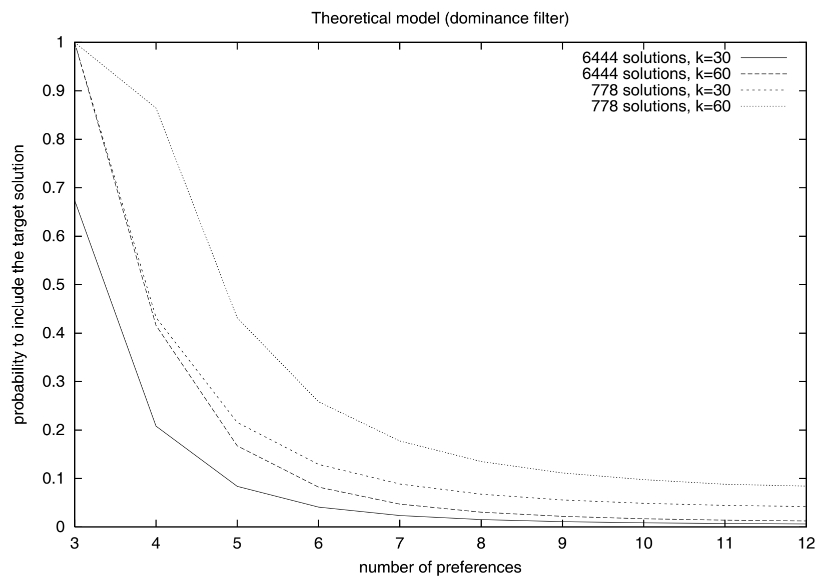

Fig. 4: Probability that the target solution is actually in the displayed set for different numbers of preferences.

solutions. Thus, in order to satisfy condition (1), the filter has to perform random sampling when the number of preferences gets too high. The sampling may include a bias to reflect a-priori knowledge about where the most preferred solution may lie.

Condition (2) is trivially satisfied as only solutions that are Pareto-optimal in the entire set of feasible solutions are shown.

For verifying condition (3), we assume that the target solution S_t is Pareto-optimal. However, as already mentioned, it will often be necessary to make a random selection of solutions to actually display. Thus, the probability of it being included in the k displayed solutions can be estimated as:

$$\begin{aligned} Pr(S_t \in \text{k-best}) &= \frac{k}{max(k, E[|O_p|])} \\ &= \frac{k}{max(k, (m-1)\int_0^1 \cdots \int_0^1 \left(1 - \prod_{p_i \in P} p_i\right)^{m-1} dp_1 \cdots dp_{|P|})} \end{aligned}$$

Figure 4 shows a plot of this probability vs. $|P| = 3, \ldots, 12$ for $m = 778, k = 30$ and $m = 6,444, k = 60$. We can see that the probability of including the target solution rapidly decreases as the overall set of Pareto-optimal solutions becomes too large and the target often is no longer selected, even for a relatively small number of preferences. Thus, the method does not satisfy condition (3) very well.

4 Utilitarian filter

This model is tailored to the case where the user's true preference model is quasi-linear, i.e. by minimizing a cost function $C^*(S) = \sum_{p_i^* \in P^*} p_i^*(S)$.

In the utilitarian filter, we approximate this by the predefined normalized penalty functions and choose as displayed solutions the k best solutions according to the unweighted sum of the penalties:

$$C(S) = \sum_{p_i \in P} p_i(S)$$

This set can be generated efficiently using branch-and-bound or other optimization algorithms (see [8]); in fact it can often be integrated with the generation of the feasible set itself.

The method obviously satisfies condition (1).

The following theorem shows that it also satisfies condition (2):

Theorem 1
Given a set of m solutions $\mathcal{S} = \{S_1, \ldots, S_m\}$ and a set of d preference penalties $\{p_1, \ldots, p_d\}$. Let $\mathcal{S}' = \{S_{i_1}, \ldots, S_{i_k}\} \subseteq \mathcal{S}$ be the best k solutions according to the utilitarian filter: $\forall S' \in \mathcal{S}', \forall S \notin \mathcal{S}' : C(S') \leq C(S)$.

If $S_x \in \mathcal{S}'$ and S_x is not dominated by any other solution $S_y \in \mathcal{S}'$, then S_x is Pareto-optimal in $\mathcal{S}$.

Proof. Assume that S_x is not Pareto-optimal in $\mathcal{S}$. Then, there is a solution $S_z \notin \mathcal{S}'$ which dominates solution S_x, and by definition:

$$\begin{aligned} \forall\, p_i,\, p_i(S_z) &\leq p_i(S_x) \text{ and} \\ \exists\, p_j,\, p_j(S_z) &< p_j(S_x) \end{aligned}$$

As a consequence, we also have:

$$\sum_{i=1}^{d} p_i(S_z) < \sum_{i=1}^{d} p_i(S_x)$$

and therefore $C(S_z) < C(S_x)$. But this contradicts the fact that $S_x \in \mathcal{S}'$ and $S_z \notin \mathcal{S}'$. □

Thus, the method also satisfies condition (2).

To understand how far the method satisfies condition (3), consider the relation between the sums of penalties for the best solutions and the number of preferences.

Let S_j be the j-th best solution according to the sum of the penalties $C(S) = \sum_{p_i \in P} p_i(S)$, and let S_t be the target solution. Since the target solution is

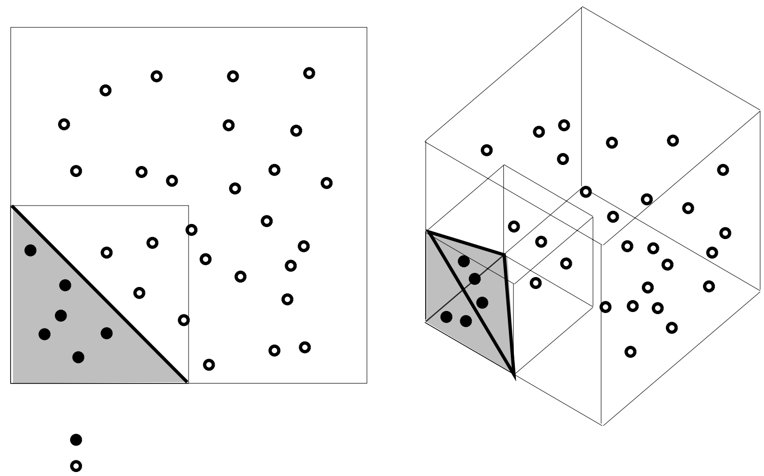

Fig. 5: The volume containing the k solutions with the smallest sum of penalties is bounded by the surface C(S) = Ck. It delimits a fraction $\beta(d)$ of the volume of a cube whose edge length is equal to Ck.

optimal with the true preferences P^* and a quasi-linear preference model, we have:

$$C^*(S_t) = \sum_{p_i^* \in P^*} p_i^*(S_t) \leq \sum_{p_i^* \in P^*} p_i^*(S_j)$$

As we assume that the standardized preferences model the true ones closely:

$$(1-\epsilon)p_i \leq p_i^* \leq (1+\epsilon)p_i$$

We can rewrite this as:

$$\begin{aligned} \sum_{p_i \in P} (1-\epsilon) \cdot p_i(S_t) &\leq \sum_{p_i \in P} (1+\epsilon) \cdot p_i(S_j) \ ; \ t \neq j \\ (1-\epsilon) \cdot C(S_t) &\leq (1+\epsilon) \cdot C(S_j) \\ C(S_t) &\leq \frac{1+\epsilon}{1-\epsilon} C(S_1) \end{aligned}$$

In the following, we write C_i for $C(S_i)$.

In order to ensure that S_t is among the k best solutions with probability 1 once all preferences are specified, we need to choose k sufficiently large so that:

$$C_k / C_1 \geq \frac{1+\epsilon}{1-\epsilon}$$

For example, if we assume a maximum error in the weights of $\epsilon = 0.2$, then $C_k \geq 1.5 C_1$.

We now consider how the ratio C_k/C_1 depends on the number of stated preferences. The k best solutions are characterized by $C(S) \leq C_k$. Assuming

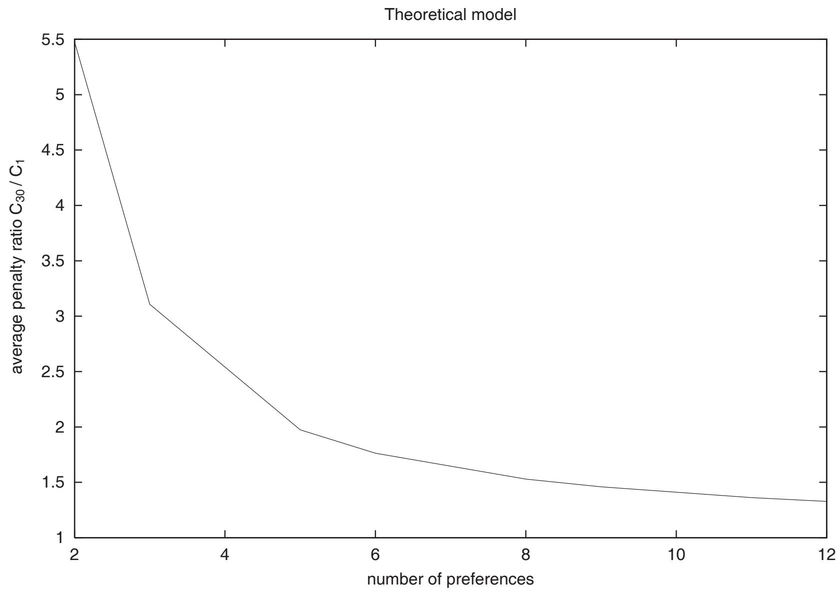

Fig. 6: Theoretically expected ratio of the summed penalty for the k-th best and the best solution as a function of the number of criteria.

a uniform distribution of solutions, C_k is the height of the volume between the hyperplane $C = C_k$ and the point $(\forall i) p_i = 0$ such that the volume is equal to k/m. As shown in Figure 5, the hyperplane $C = C_k$ cuts out a portion $\beta(d)$ of a d-dimensional hypercube H of size C_k, where $\beta(d)$ is a function of the dimension d. H thus has a volume of $\beta(d)C_k^d$. As $\beta(d)$ of the cube should contain a fraction k/m of the solutions, the volume of the displayed set is equal to:

$$\beta(d)C_k^d = k/m$$

and thus

$$C_k = \left(\frac{k}{m\beta(d)}\right)^{1/d}$$

and

$$C_k/C_1 = k^{1/d}$$

which somewhat surprisingly is independent of the total number of solutions m.

We again compare the predictions of the theory with the one observed in practical experiments. Figure 6 shows the theoretically predicted penalty ratio C_k/C_1 obtained for random configuration problems for $k = 30$ displayed solutions. Figure 7 shows the corresponding curves generated with random configuration problems, again with $k = 30$ displayed solutions. For the case of $m = 6444$ solutions, the curves match exactly, while for $m = 778$ solutions, we observe a significant discrepancy when the number of preferences is low. This is due to the fact that many solutions to the configuration problem have identical

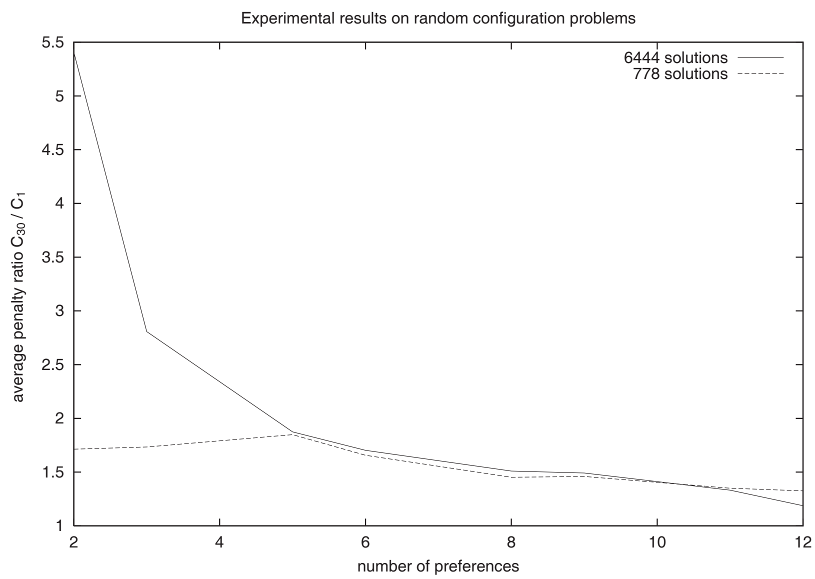

Fig. 7: Experimentally observed ratio of the summed penalty for the k-th best and the best solution as a function of the number of criteria.

quality and thus often score the same when only a few preferences are stated. In general, we can again conclude that the theoretical model is a good match of reality.

Given the error margin ϵ as described above, we can compute the expected maximal rank t of the target solution in terms of the summed penalty metric by the fact that:

$$C_t = C_1 \frac{1+\epsilon}{1-\epsilon} = \left(\frac{t}{m\beta(d)}\right)^{1/d}$$

from which we obtain:

$$\begin{aligned} t &= m\beta(d)\left(\frac{C_1(1+\epsilon)}{1-\epsilon}\right)^d \\ &= \left(\frac{1+\epsilon}{1-\epsilon}\right)^d \end{aligned}$$

Using the fact that:

$$C_1 = \left(\frac{1}{m\beta(d)}\right)^{1/d} \Rightarrow C_1^d m\beta(d) = 1$$

This function can be used to calculate the number of solutions that have to be displayed for a given number of preferences and a given tolerance ϵ for user

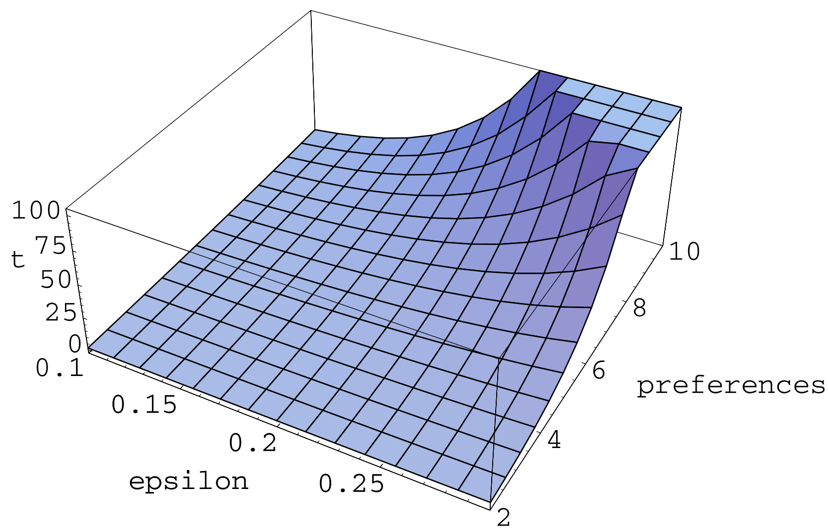

Fig. 8: Size of the displayed set required to ensure that the target solution is displayed, for different values of $|P|$ and ϵ.

diversity. This can be used in the design of a decision support system, or to adaptively adjust the number of displayed solutions during a decision process. Figure 8 shows a plot of this function for $|P| = 2, \cdots, 10$ and ϵ varying from 0.1 to 0.3.

With the weight errors at the maximum admissible limits, the probability that the target solution lies within the displayed set is given by:

$$pr(S_t \in \text{k-best}) = \begin{cases} 1 & \text{if t} \leq \text{k} \\ k/t & \text{if t} > \text{k (random selection)} \end{cases}$$

Figure 9 shows a plot of this probability for $|P| = 2, \ldots, 12$ and combinations of $\epsilon = 0.2, 0.3$ and $k = 30, 60$.

We can see that in comparison to the dominance filter, we can often accommodate significantly more preferences. However, this depends strongly on the weight diversity and the number of displayed solutions. While the utilitarian filter in principle can be expected to perform better than dominance filtering, careful tuning is still required for each application.

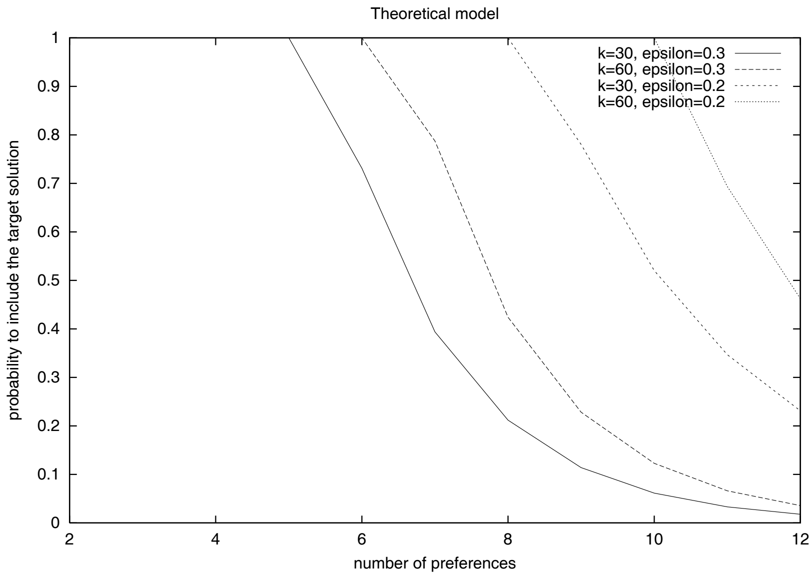

Fig. 9: Probability that the target solution is within the displayed set.

5 Egalitarian filter

The egalitarian filter is designed for the case where the user's true preference model is the egalitarian model, i.e. he minimizes the maximum penalty $F^*(S) = max_{p_i^* \in P^*} p_i^*(S)$. The egalitarian filter approximates this by applying the same criterion on the standardized penalty functions:

$$F(S) = max_{p_i \in P} p_i(S)$$

where ties among several solutions with the same maximal value are broken by applying the same criterion to the remaining preferences (this is called lexicographic). The method corresponds to the combination rules used in fuzzy logic, and has the advantage that it can be efficiently implemented using constraint satisfaction techniques such as MAX-CSP ([15]) or Fuzzy CSP ([13]).

Similarly to the utilitarian filter, we assume that the penalty functions are normalized for a standard user, but that for a particular user they can deviate by no more than a factor of ϵ in either direction.

This method obviously satisfies condition (1), as it generates exactly k solutions.

It also satisfies condition (2) by the following theorem:

Theorem 2
Given a set of m solutions $\mathcal{S} = \{S_1, \dots, S_m\}$ and a set of d preference penalties $\{p_1, \dots, p_d\}$. Let $\mathcal{S}' = \{S_{i_1}, \dots, S_{i_k}\} \subseteq \mathcal{S}$ be the best k solutions according to the egalitarian filter.

If $S_x \in \mathcal{S}'$ and S_x is not dominated by any other solution $S_y \in \mathcal{S}'$, then S_x is Pareto-optimal in $\mathcal{S}$.

Proof. Assume that S_x is not Pareto-optimal in $\mathcal{S}$. Then, there is a solution $S_z \notin \mathcal{S}'$ which dominates solution S_x, and by definition:

$$\begin{aligned} \forall\, p_i,\, p_i(S_z) &\leq p_i(S_x) \text{ and} \\ \exists\, p_j,\, p_j(S_z) &< p_j(S_x) \end{aligned}$$

Let p_m be a preference with the highest penalty in solution S_x:

$$\forall\, p_i,\, p_i(S_x) \leq p_m(S_x)$$

Then since S_z dominates S_x, $p_m(S_z) \leq p_m(S_x)$ and $\forall p_i, p_i(S_z) \leq p_m(S_x)$ and thus $F(S_z) \leq F(S_x)$.

When $F(S_z) < F(S_x)$, this contradicts the fact that $S_x \in \mathcal{S}'$ and $S_z \notin \mathcal{S}'$.

When $F(S_z) = F(S_x)$, we have a tie and the same argument applies to the set of preferences with p_m removed. As S_z dominates S_x, there must eventually be a preference p_e such that $p_e(S_z) < p_e(S_x)$, leading to the contradiction as shown above. □

The degree to which the egalitarian filter satisfies condition (3) depends on how much the penalties for the actual user deviate from the normalized values. We first determine the expected value F_t of the optimization function for the target solution S_t. For S_t, let $p_i^*(S_t) = \alpha_i \cdot p_i(S_t)$, and let $F_t^* = p_m^*(S_t)$ be the value of the maximum penalty function p_m^*. In the user's preference model, the target solution S_t can be found on a volume of size $\prod_{i=1}^{d} p_m^* = (p_m^*)^d$. The dimension p_m^* can be found by the fact that this volume is expected to contain exactly one solution (the optimal one). As we assume the solutions to be uniformly distributed in the space of standardized preference functions, this means that the volume in the original formulation must be equal to $1/m$:

$$p_m^d \prod_{i=1}^{d} \alpha_i = 1/m \Rightarrow p_m = \left(m \prod_{i=1}^{d} \alpha_i \right)^{-1/d}$$

The egalitarian filter selects all solutions in a cube with equal sides $l(k)$ whose volume is expected to contain k solutions:

$$l^d(k) = k/m$$

We can guarantee that the cube can be expected to include a solution that is optimal under the skewed metric if it is large enough so that:

$$l(k) \geq p_m \alpha_{max}$$

where $\alpha_{max} = max_{\alpha_i}$. This can be achieved if we choose:

$$\begin{aligned} k/m = l^d(k) &\geq 1/m\left(\frac{1+\epsilon}{1-\epsilon}\right)^{(d-1)} \\ &\geq \frac{(\alpha_{max})^d}{m\prod_{i=1}^{d}\alpha_i} \\ &= (p_m\alpha_{max})^d \end{aligned}$$

where the second inequality has been obtained by choosing $\alpha_{max} = 1+\epsilon$ and for all the remaining α_i the minimum $1-\epsilon$.

This gives us in fact a similar behavior as for the utilitarian filter, but the egalitarian filter allows one more preference with the same number of displayed solutions.

6 Robustness against violated assumptions

In our analysis of the utilitarian and egalitarian filters, we have assumed that the user's preferences actually follow this model. However, it may be that this assumption is false, and that the only thing that can be assumed of a particular user is that he will not prefer a dominated solution.

The performance of each filter in this case can be characterized by the number k of best solutions required to include the target solution S_t, which can be any Pareto-optimal solution.

Recall that the number of expected number of Pareto-optimal solutions $E[|O_p|]$ is given by the integral:

$$E[|O_p|] = m\int_0^1 \cdots \int_0^1 \left(1 - \prod_{p_i \in P} p_i\right)^{m-1} dp_1 \cdots dp_d$$

where $d = |P|$ is the number of preferences in the set P.

Assuming a uniform distribution of the solutions, the same analysis can be applied to the volumes covered by the utilitarian and egalitarian filters.

In particular, the expected number of Pareto-optimal solutions displayed by the utilitarian filter is:

$$E[|O_s|] = m)\int_0^{C_k} dp_1 \int_0^{C_k - p_1} dp_2 \cdots \int_0^{C_k - p_1 - \ldots - p_{d-1}} dp_d \left(1 - \prod_{p_i \in P} p_i\right)^{m-1}$$

where $C_k = \left(\frac{k}{\beta(d)m}\right)^{\frac{1}{d}}$ is the utilitarian threshold for having k solutions. As long as C_k is not greater than 1, the volume is simple and $\beta(d) = 1/d!$[2], and so for the purposes of this calculation, we approximate $C_k = \left(\frac{kd!}{m}\right)^{\frac{1}{d}}$.

[2] When $C_k > 1$, there are corners that need to be excluded from the volume and $\beta(d)$ is a more complex expression.

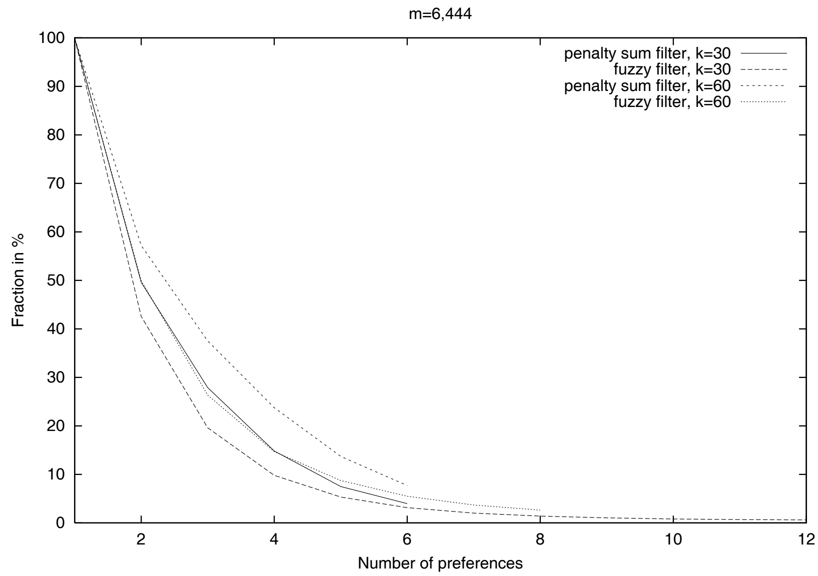

Fig. 10: Fraction of Pareto-optimal solutions shown within the k best ones according to utilitarian and egalitarian filters.

The expected number of Pareto-optimal solutions displayed by the egalitarian filter is:

$$E[|O_f|] = (m-1)\int_0^{F_k}\cdots\int_0^{F_k}\left(1-\prod_{p_i\in P} p_i\right)^{m-1} dp_1\cdots dp_d$$

where $F_k = \left(\frac{k}{m}\right)^{\frac{1}{d}}$ is the threshold for an expected number of k solutions.

When the number of Pareto-optimal solutions is larger than the set generated by these filters, they will compute a random sample and the probability that the target solution S_t is included is given by the fractions $E[|O_s|]/E[|O_p|]$ and $E[|O_f|]/E[|O_p|]$.

Unfortunately, both integrals cannot be solved in closed form. However, we can compare numerical solutions. Figure 10 shows the probability of including a Pareto-optimal solution in a problem with $m = 6,444$ solutions when $k = 30$ and $k = 60$ solutions are shown for different numbers of preferences. We can see that in both cases the probability decreases quite dramatically; however, utilitarian filter is in general slightly better than the egalitarian one.

When comparing with Figure 4, the performance of utilitarian and egalitarian filters is almost as good as that of a dominance filter. Thus, even when the user's true preference model is unknown, a pure dominance filter seems to be of little use. This is particularly true considering that methods for computing Pareto-optimal solutions that form the basis for dominance filters are computationally significantly more expensive than optimization methods for utilitarian and egalitarian filters ([8]).

In a similar analysis that we omit for reasons of space, we can show that the utilitarian filter behaves extremely poorly when the true preference model is an egalitarian one, and the egalitarian filter behaves even worse when the true preference model is utilitarian.

7 Experimental results on randomly generated problems

Throughout the paper, we have validated the results obtained through theoretical analysis on randomly generated configuration problems. We model configuration problems with preferences as valued constraint satisfaction problems (VCSP [16, 2]) as described in [19].

A random valued CSP is defined by: $< n, m, hc, ht, sc, st >$, where n is the number of variables in the problem and m the size of their domains. hc is the graph density in percentage for unary and binary hard constraints. ht is the tightness in percentage for disallowed tuples in unary and binary hard constraints. sc and st are the graph density and tightness in percentage for unary and binary soft constraints. Valuations for soft constraints range from 0 to 1. For simplicity, hard and soft constraints are separated and we are not considering mixed constraints, therefore $hc + sc \leq 100$.

For building random instances of valued CSPs, we choose the variables for each constraint following a uniform probabilistic distribution. In the same way we choose the tuples in constraints. Valuations for soft tuples are randomly generated between 0 and 1 and valuations for hard tuples are represented by a maximum valuation (∞).

The algorithms have been tested with different sets of problems of valued CSPs with 5 variables and 10 values for each variable. Hard unary/binary constraint density hc has been varied from 20% to 80% in steps of 20, and the tightness for hard constraints ht varies also from 20% to 80% in steps of 20. Soft unary/binary constraint density sc has been varied from 20% to 80% in steps of 10, resulting in 3, 5, 8, 9, 11 and 12 soft constraints with tightness $st = 100$. In total, there could be $5 + (5 * 4)/2 = 15$ constraints (5 unary constraints and 10 binary constraints).

For every different class of problems, 40 different instances were generated. For the different problem topologies the average of the results for each instance are evaluated.

For simplicity and easing the readability of the graphs presented in this paper, only the graphs for problems with the topology $< n = 5, m = 10, hc = 20, ht = \{80, 60\}, sc = \{20, \ldots, 80\}, st = 100 >$ are shown. The total number of solutions for these two problem topologies ($ht = 60$ and $ht = 80$) on average are 778 and $6,444$.

8 reality: a commercial application for decision support in business travel

reality is an application commercialized by i:FAO[3] which significantly extends the reach of electronic commerce in travel (for more details about this application, refer to [20]). In particular, *reality* addresses the challenge of modeling customers' personal preferences and providing tailored solutions. In contrast to earlier technology, which optimizes only a small and predefined set of preferences, our tool can accurately model the preferences of different customers. *reality* thus allows customers to more quickly find solutions of better quality than existing tools.

The idea underlying *reality* is to replace the travel agent while keeping the same interaction model, *i.e.* keeping a kind of dialog between the customer and the system by means of a mixed-initiative system.

Following the analysis presented in this paper, the solution generation adopted for *reality* is the utilitarian filter presented in Section 4. We have chosen $k = 30$ (the number of solutions to be displayed) as the optimal number since we found that users typically state between 3 and 6 preferences for a trip and that $\epsilon = 0.3$ is a good upper bound on their diversity.

Figure 11 shows how *reality* displays the best 30 solutions out of the generation process to the user. A scatterplot (inspired by [1]) shows how solutions are positioned with respect to two criteria chosen by the user. Below the scatterplot, details of the current selected solution are displayed by a standard table of flight attributes. Preferences can be posted in any attribute of the shown solution. In the example of Figure 11, the user is posting a preference on the transit time for the itinerary segment from Geneva to San Francisco. The scatterplot shows the 30 best solutions with respect to the criteria *satisfaction* and total travel time. The *satisfaction* criterion represents the penalty sum of all the preferences, *i.e.* the quality of the solutions.

9 Conclusions

In this paper, we have considered the problems posed by the use of a qualitative preference model in decision aid systems. We have considered compensating for the inaccuracy of a qualitative model by displaying not a single optimum, but letting the user choose among a displayed set of k alternatives. A key design issue in such a system is how to select the displayed set to ensure that the target solution can indeed be found.

The main result is that when the user's preference model is not known beyond the basic relations of dominance, it is very difficult to guarantee that the target solution can be found at all. This is due to the strong growth in the number of Pareto-optimal solutions, and occurs for all three kinds of filters: dominance, utilitarian and egalitarian filters. It can be argued that even in this case using a pure dominance filter has only marginally better performance

[3] `http://www.ifao.net`

Fig. 11: The user analyzes the solutions shown in a starfield where two criteria are compared (in this case satisfaction vs. total flying time).

than either the utilitarian or egalitarian filter so that its usefulness in practice is questionable.

On the other hand, when the user's preference model follows a utilitarian or egalitarian model, it is possible to guarantee that the optimal solution is found as long as the number of preferences does not become excessive (in the case of travel planning, 5-6 preferences).

This makes it possible to show that a decision aid system using the utilitarian or egalitarian filter is sound, i.e. finds the target solution, if the size of the displayed set k is sufficiently large for the ϵ and the number of preferences that the users poses, and if the user correctly and completely states his preferences. This allows computing the right number of displayed solutions to display based on experiments that determine the number of preferences that a user is likely to state.

Acknowledgments

We would like to sincerely thank all the people that contributes, or contributed, to the development of *reality*.

References

[1] C. Ahlberg and B. Shneiderman: “Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays,” Human Factors in Computer Systems: Proceedings of the CHI’94 Conference, New York, pp. 313-317, 1994

[2] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex and G. Verfaillie: “Semiring-based CSPs and Valued CSPs: Basic Properties and Comparison,” CONSTRAINTS **4**(3), Kluwer Academic Publishers, 1999

[3] C. Boutilier, R. Brafman, C. Geib, and D. Poole: “A Constraint-Based Approach to Preference Elicitation and Decision Making,” AAAI Spring Symposium on Qualitative Decision Theory, Stanford, 1997

[4] R. Burke, K. Hammond, and B. Young: “The findme approach to assisted browsing,” IEEE Expert **12**(4), pp. 32-40, 1997

[5] J. Doyle and R. Thomason: “Background to qualitative decision theory,” AI Magazine **20**(2), Summer 1999

[6] V. Ha and P. Haddawy: “Problem-focused incremental elicitation of multi-attribute utility models,” Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, pp. 215–222, August 1997

[7] V. Ha and P. Haddawy: “Toward Case-Based Preference Elicitation: Similarity Measures on Preference Structures,” in Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence. 1998

[8] R.Keeney and H. Raiffa: “Decision with multiple objectives: Preferences and value tradeoffs,” Cambridge University Press, 1976

[9] R. Keeney: “Value-Focused Thinking: A Path to Creative Decision Making,” Harvard University Press, 1992

[10] G. Linden, S. Hanks, and N. Lesh: “Interactive assessment of user preference models: The automated travel assistant,” In Proceedings of User Modeling ’97, 1997

[11] P. Pu and B. Faltings: “Enriching Buyers’ experiences: the SmartClient Approach,” Proceedings of The ACM SIGCHI Conference on Human Factors in Computing Systems, 2000

[12] P. Pu, B. Faltings and M. Torrens: “User-Involved Preference Elicitation,” Workshop on Configuration, International Joint Conference on Artificial Intelligence (IJCAI), Acapulco, August 2003

[13] Z. Ruttkay: “Fuzzy Constraint Satisfaction,” Proceedings of the 1st IEEE Conference on Evolutionary Computing, Orlando, pp. 542-547, 1994

[14] S. Shearin and H. Lieberman: "Intelligent Profiling by Example," Proceedings of the International Conference on Intelligent User Interfaces (IUI 2001), pp. 145-152, 2001

[15] T. Schiex: "Possibilistic constraint satisfaction problems or 'How to handle soft constraints?' ", *Proceedings of the 8th International Conference on Uncertainty in AI*, Stanford, 1992

[16] T. Schiex, H. Fargier and G. Verfaillie: "Valued Constraint Satisfaction Problems: Hard and Easy Problems," Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) pp. 631-637, 1995

[17] M. Stolze: "Soft Navigation in electronic product catalogs," International Journal on Digital Libraries 3(1), pp. 60-66, 2000

[18] M. Torrens: "Scalable Intelligent Electronic Catalogs," Ph. D. Thesis No. 2690, Swiss Federal Institute of Technology, Lausanne (EPFL), 2002

[19] M. Torrens and B. Faltings: "Multi-criteria Optimization in Constraint Satisfaction Problems," AAAI-02 workshop on preferences in AI and CP: symbolic approaches, 2002

[20] M. Torrens, P. Hertzog, L. Samson and B. Faltings: "*reality*: a Scalable Intelligent Travel Planer, Decision Support for the Business Traveler," Proceedings of the Eighteenth Annual ACM Symposium on Applied Computing, Melbourne, 2003