Towards a possibilistic logic handling of preferences

Salem Benferhat Didier Dubois Henri Prade
IRTT-CNRS,
University Paul Sabatier, 118 route de Narbonne,
31062 Toulouse Cedex 4,
France

Abstract

The classical way of encoding preferences in decision theory is by means of utility or value functions. However agents are not always able to deliver such a function directly. In this paper, we relate three different ways of specifying preferences, namely by means of a set of particular types of constraints on the utility function, by means of an ordered set of prioritized goals expressed by logical propositions, and by means of an ordered set of subsets of possible candidates reaching the same level of satisfaction. These different expression modes can be handled in a weighted logical setting, here the one of possibilistic logic. The aggregation of preferences pertaining to different criteria can then be handled by fusing sets of prioritized goals. Apart from a better expressivity, the benefits of a logical representation of preferences are to put them in a suitable format for reasoning purposes, or for modifying or revising them.

1. Introduction

In decision analysis, utility or valued functions are supposed to be available for computing the values of global objective functions and for ranking the different possible candidates under consideration.
Artificial Intelligence methods can contribute to a more implicit specification of these functions, for instance in terms of constraints allowing for a more granular expression of preferences. This general line of research has been recently illustrated in various ways by AI researchers (e.g., Boutilier, 1994; Lang 19%; Boutilier et al. 1997). The expected benefit of the logical handling of decision problems is not only to allow for a less abstract, and thus more human-like expression of knowledge and preferences, but also to facilitate explanation capabilities for the candidates proposed by decision support systems.

Logic and decision belong to two different traditions; the first one is concerned with consistency and inference and is oriented towards symbolic processing, while the other deals with trade-offs (and possibly with uncertainty), and is more numerically inclined. However non-classical logics are often using ordering structures, while a need is now expressed in decision analysis for more qualitative evaluations which only require ordinal scales (rather than numerical ones).

Among weighted logics, possibilistic logic based on the conjoint use of classical logic and qualitative possibility theory (Dubois and Prade, 1998; Zadeh, 1978) offers a framework at the meeting point of the two traditions. Besides, possibilistic logic has been already shown to be convenient for handling nonmonotonic reasoning (Benferhat et al., 1992, 1997). More recently, its framework has been used for modelling preferences, as set of prioritized goals in decision under uncertainty (Dubois et al., 1998).

This paper provides a preliminary investigation of the potentials of possibilistic logic in decision analysis, and more specifically in the representation of preferences. Indeed a possibilistic logic base can not only be seen as a set of more or less certain pieces of information (which was the original understanding when possibilistic logic was introduced and then applied to nonmonotonic reasoning). Such a base can also be viewed as a layered set of propositions expressing goals having different levels of priority. The latter view can be connected with the fuzzy set representation of constraints or objective functions proposed by Bellman and Zadeh (1970) a long time ago. Indeed, a utility function can be seen as a membership function of a fuzzy set (the one expressing the more or less acceptable candidates), which gives birth to a weighted set of goals (through the level cuts of the fuzzy set).

Section 2 presents the representation of utility functions over set of possible candidates, either in terms of set prioritized goals or in terms of subsets of possible candidates reaching the same level of satisfaction. Section 3 discusses the symbolic aggregation of utility functions pertaining to different criteria. Section 4 studies how constraint-based specification of preferences can lead to a representation in the previous framework. Section 5 briefly deals with the revision of preferences and goals.

2. Logical handling of prioritized goals

Let U be a finite set of possible candidates. A utility function, associated with some criterion C, is a mapping from U to some valuation scale. In many practical situations, a finite valuation scale is enough, first because the set of candidates is finite, and moreover humans are often only able to differentiate candidates through a rather small number of valuations.

Then a fuzzy set can be equivalently seen as a finite family of nested level cuts, corresponding here to crisp constraints or objectives. The equivalent representation of C as a set of prioritized goal is a direct consequence of the semantics associated to a possibilistic logic base, which is now briefly restated.

2.1. Possibilistic logic semantics

Let $K = \{(p_j, \rho_j); j = 1,n\}$ be a possibilistic logic base where p_j is a classical proposition and ρ_j a level in a linearly ordered valuation set L. The semantics of K is given by the function π_K from the set of interpretations U to L, such that all the interpretations satisfying all the propositions in K get the highest possibility degree, namely 1, and the others are ranked w.r.t. the strongest proposition that they falsify, namely we get:

$$\pi_K(u) = \min_{j=1,n} \max(v_u(p_j), 1 - \rho_j), \quad (1)$$

where $v_u(p_j)$ a 1 if u is a model of p; (i.e., an interpretation which makes it true) and $v_u(pj) = 0$ if u falsifies pj. The valuation scale L may be the unit interval, but a finite linearly ordered set can be often sufficient in practice (then $1 - (\cdot)$ just denotes the order-reversing map of L). The function π_K is a possibility distribution which rank-orders the interpretations u. Associated with π_K are two mappings which evaluate subsets of U, namely

• a possibility measure, which evaluates to what extent p is consistent with interpretations having a high degree of possibility according to π_K:

$$\Pi(p) = \max_{u:\ u \text{ model of } p} \pi_K(u), \quad (2)$$

• a dual necessity measure, which evaluates to what extent all the interpretations having high possibility degrees make p true according to π_K:

$$N(p) = 1 - \Pi(\neg p) = \min_{u:\ u \text{ falsifies } p}(1 - \pi_K(u)) \quad (3)$$

It is worth noticing that (1) results from the application of the minimal specificity principle to the set of constraints (Dubois et al., 1994)

$$N(p_j) \geq \rho_j \quad \text{for } j = 1,n. \quad (4)$$

Indeed, (4) implicitly specifies a set of possibility distributions and the minimal specificity principle consists in choosing the greatest possibility distribution satisfying the constraints. This distribution allocates the greatest possibility level to each interpretation in agreement with the constraints. In the following, the possibility degree $\pi_K(u)$ expresses how satisfactory is the candidate u.

2.2. Logical representation of criteria

Let us consider the case of a unique fuzzy criterion C, the utility function is defined by its membership function μ_C ranging on a finite scale $L = \{\alpha_0 = 0 < \alpha_1 < ... < \alpha_n = 1\}$. C is equivalently represented by the set of constraints $N(C_{a_i}) \geq 1 - \alpha_{i-1}$ for $i = 1,n$, where N is the necessity measure defined $\pi = \mu_C$. C_{α_i} is the α_i-cut of μ_C defined by:

$$C_{\alpha_i} = \{u : \mu_C(u) \geq \alpha_i\},$$

i.e., the set of possible candidates having a degree of satisfaction having degree at least equal to α_i. The greater aj the smaller C_{α_i}. Note that if $C_{\alpha_i} = C_{\alpha_{i+1}}$ then the constraint $N(C_{\alpha_{i+1}}) \geq 1 - \alpha_i$ is redUndant, aAd can be ignored.

This can be also reinterpreted in terms of priority : the goal of picking a candidate in C_{α_i} has priority $1-\alpha_{i-1}$, and the larger the a-cut, the more irtiportant the priority; in particular it is imperative that the chosen u has a non-zero degree of satisfaction, so C_{α_1} has priority 1 ($N(C_{a1})=1$).

This gives birth to a possibilistic knowledge base of the form $K = \{(c_{\alpha_i}, 1 - \alpha_{i-1}), i=1,n\}$ where c_{α_i} denotes the proposition whose set of models is C_{α_i}. Preferences ate thus expressed in terms of sets of crisp (nested) goals having different levels of priority. Clearly, π_K computed by (1) is equal to μ_C. $\pi_K(u)$ is all the smaller as u violates goals with higher priority. Decisions violating goals with priority 1 have a level of acceptability equal to 0.

This representation plays the basic role in the manipulation that we may need to perform on preferences.

2.3. Conjunctive and disjunctive forms

Conversely, a set of crisp goals (not necessarily nested) with different levels of priority can always be represented in terms of a fuzzy set membership function as we are going to see on different examples.

Example 1: Hierarchical requirements.

In operations research, and as well as in the database setting (e.g., Lacroix and Lavency, 1987), requirements of the following form have been considered: "C_1 should be satisfied, and among the solutions to C_1 (if any) the ones satisfying C2 are preferred, and among satisfying both C_1 and C_2, those satisfying C3 are preferred and so on", where C_1, C_2, C_3... are here supposed to be classical constraints (i.e., $\mu_{C_i} = 0$ or 1). Thus, one wishes to express that C_1 should hold (with importance or priority ρ_1), and that if C_1 holds, C_2 holds with priority ρ_2, C3 holds with priority ρ_3 (with $\rho_3 < \rho_2 < \rho_1$). This can be readily expressed by the possibilistic propositional logic base

$$K = \{(c_1, 1); (\neg c_1 \vee c_2, \rho_2); (\neg c_1 \vee \neg c_2 \vee c_3, \rho_3)\}.$$

The semantics of K obtained by applying (1) can be put under the form

$$\pi_K(u) = \min(\mu_{C_1}(u), \max(\mu_{C_2}(u), 1 - \min(\mu_{C_1}(u), \rho_2)), \max(\mu_{C_3}(u), 1 - \min(\mu_{C_1}(u), \mu_{C_2}(u), \rho_3)). \quad (5)$$

is a weigntea mm aggregation wmen renects me iaea that we are completely satisfied C_1 and C_2 C3 ($\pi_K(u) = 1$) are completely satisfied, we are less satisfied ($\pi_K(u) = 1 - \rho_3$) if C_1 and C_2 only are satisfied, and we are even less satisfied if only C_1 ($\pi_K(u) = 1 - \rho_2$) is satisfied.

A semantically equivalent form for K can be obtained by applying the possibilistic logic resolution rule (Dubois et al., 1994), $(\neg a \vee b, \alpha), (a \vee c, \beta) \vdash (b \vee c, \min(\alpha, \beta))$, namely $K = \{(c_1, 1); (c_2, \rho_2); (c_3, \rho_3)\}$. Indeed (5) can be put under the form (1) with n = 3 and $v_u(c_j = UC_j(U)$ Thus the priorities can directly reflect a hierarchy in possibilistic logic.

Expressions (1) or (5) correspond to conjunctive normal forms (i.e., it is a min of max). They can be turned into disjunctive normal forms (max of min) and then provide a description of the different classes of candidates ranked according to their level of preference, as seen in the example below (add the candidate in a class reaching the same level of satisfaction).

Example 2:

Let us consider the following three criteria-based evaluation: "if u satisfies A and B, u is completely satisfactory, if A is not satisfied, less satisfactory solutions should at least satisfy C". Such an evaluation function can be encountered in multiple criteria problems for handling "special" cases (here situations where A is not satisfied) coexist with a normal case (here situation where both A and B can be satisfied). It can be directly represented by the disjunctive form:

$$\mu_D(u) = \max(\min(\mu_A(u), \mu_B(u)), \min(\mu_C(u), 1-\mu_A(u), 1-\rho))$$

with $\rho < 1$. The reading of this expression is easy. Either the candidate satisfies both A and B, or if it falsifies A, it satisfies C, which is less satisfactory.

This expression of μ_D obtained as the weighted union of the different classes of more or less acceptable solutions can be transformed into an equivalent conjunctive form like (1); it can be checked that this conjunctive form corresponds to the base $K = \{(a \vee c, 1); (\neg a \vee b, 1); (a, \rho); (b, \rho)\}$, where A, B, C are the sets of models of a, b and c respectively; this provides a logical, equivalent description of the evaluation process in terms of prioritized requirements to be satisfied by acceptable solutions.

It is worth noticing that the clausal form corresponding to the possibilistic logic base may be sometimes less natural for expressing the goals than the normal disjunctive form as shown by Example 2 above. Example 1 illustrates the converse situation.

The normal disjunctive form provides a logical description of the different subsets of solutions each with its level of acceptability. On the contrary, a possibilistic logic base which can always be put under the form of a conjunction of possibilistic clauses corresponds to a prioritized set of goals.

2.4. Basic modifications of a set of goals

Discounting and thresholding are two elementary operations that can be performed on a preference profile. Indeed the discounting of a preference profile μ_C, associated with a criterion C by a level of importance $w \in L$ in a qualitative setting amounts to modifying $\mu_C(u)$ into $\max(\mu_C(u), 1-w)$ for each u. It expresses that even if the candidate u is not at all satisfactory w.r.t. the initial criterion C ($\mu_C(u) = 0$) the candidate is no longer completely rejected w.r.t. the discounted criterion, and receives a value which is all the greater as the level of importance of w is smaller, i.e., as the discounting is stronger. Thresholding a preference profile by θ in a qualitative setting amounts to modifying $\mu_C(u)$ into $\theta \to \mu_C(u) = 1$ if $\mu_C(u) \geq \theta$ and $\theta \to \mu_C(u) = \mu_C(u)$ if $\mu_C(u) < \theta$. In other words, as soon as the cndidate u reaches the satisfactory level θ w.r.t. μ_C, it is regarded as fully satisfactory w.r.t. to the thresholded criterion, otherwise the satisfactory level remains unchanged.

These two operations are easy to perform on the representation in terms of prioritized goals. Indeed,

• the importance weighting operation $\max(\mu_C(u), 1-w)$ translates into the suppression of the most prioritary goals $(c_{\alpha_i}, 1-\alpha_{i-1})$ such that $1-\alpha_{i-1} > w$. When w=1 no modification occurs, while when w=0 all the goals disappear.

• the thresholding operation defined by $\theta \to \mu_C(u)$ translates into the suppression of the least prioritary goals $(c_{\alpha_i}, 1-\alpha_{i-1})$ such that $\alpha_i > \theta$. As in the previous case, if 0 = 1 'ho modification occurs, while when 0 = 0 all the goals disappear.

More generally, the conjunctive aggregation of fuzzy (discounted or thresholded) preference profiles can be interpreted in terms of conjunctions of crisp goals having different levels of priority, thus providing an expression of preferences in a possibilistic logic form.

3. Logical aggregation of fuzzy constraints

The pointwise aggregation of two fuzzy preference profiles C_1 and C_2 defined by means of the min operation can be easily interpreted in the prioritized goals framework. It corresponds to the union of the two sets of possibilistic logic formulas $\{(c_{1\alpha_i}, 1-\alpha_{i-1})\}$ and $\{(c_{2\alpha_i}, 1-\alpha_{i-1})\}$ This is a particular case of the syntactic fusion of possibilistic pieces of information (Benferhat et al. 1997b).

Aggregation operations other than min can be also accommodated in a symbolic manner. Indeed reinforcement and compensation operators, such as the product and the average respectively, can also be interpreted in terms of operations on prioritized goals. Let (a, α) and (b, β) be two crisp constraints with priorities a and β and * be an increasing aggregation operator. The aggregation of (a, α) and (b, β) is expressed pointwisely at the semantical level by:

$$\max(\mu_A(u), 1-\alpha) * \max(\mu_B(u), 1-\beta)$$

which can be easily interpreted in terms of prioritized goals. As it can be checked, this aggregation symbolically denoted by $(a, \alpha) * (b, \beta)$ is equivalent to the (min) conjunction of the prioritized goals:

$$(a, 1-((1-\alpha) * 1)), \quad (b, 1-(1 * (1-\beta))),$$

$$(a \vee b, 1-(1-\alpha) * (1-\beta)).$$

provided that $1 * 1 = 1$. Note that the combination amounts to adding the goal $a \vee b$ to a level of priority higher than the ones of a and b. Indeed, provided that * is an increasing operation, $1-((1-\alpha) * (1-\beta))$ is greater or equal to $1-((1-\alpha) * 1)$ and $1-(1 * (1-\beta))$. If * = min, the third weighted clause is redundant w.r.t. the two others.

This can be generalized to fuzzy preference profiles A and B. It can be shown that A * B is equivalent to the conjunction of the following sets of prioritized goals:

$$\{(a_{\alpha_i} \vee b_{\beta_k}, 1-(\alpha_{i-1} * \beta_{k-1})) \text{ for all } (i, k)\},$$

$$\{(a_{\alpha_i}, 1-(\alpha_{i-1} * 1)) \text{ for all } i\} \text{ and}$$

$$\{(b_{\beta_k}, 1-(1 * \beta_{k-1})) \text{ for all } k\}.$$

It should be emphasized that the translation of aggregation * into a possibilistic propositional logic base is done at the expense of the introduction of new levels in the scale. Indeed * is not closed on the finite scale $\{\alpha_0 = 0 < \alpha_1 < \ldots < \alpha_n = 1\}$ generally. Moreover, note that the symmetry of * is not required. The goal base K in Example 2 can be retrieved by combination of its syntactic components $\{(a, 1), (b, 1)\}$ and $\{(c, 1), (\neg a, 1), (\bot, \rho)\}$ for * = max.

Example 3 : (Moura-Pires and Prade, 1998)

Consider three preference profiles A, B and C where A and B are fuzzy and C is discounted by p. A is supposed to be thresholded by 0. Moreover (C, p) and B are supposed to be aggregated by a compensatory operation, here the arithmetic mean $(s * t = (s + t)/2)$. This can be formally written as : $(\theta \to A) \wedge ((C, \rho) * B)$ where $\wedge$ stands for the min aggregation. In the example we use the satisfaction scale $\{\alpha_0, \alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5\} = \{0, 0.2, 0.4,$

0.6, 0.8, 1} for *A* and B and we take 0 *m* 0.8, p = 0.6. The problem can be then translated under the form of a stratified possibilistic base. Namely, let $A_j = A_{\alpha_{i+1}}$, eg., $A_1 = A_{\alpha_2}$, (here the 0.4-cut). Thus, the level cuts of the fuzzy set $(\theta \rightarrow A) \wedge ((C, \rho) * B)$ lead to the possibilistic base;

$\{(a_0, 1), (a_1, 0.8), (b_0 \vee c, 0.8), (b_1 \vee c, 0.7), (a_2, 0.6), (b_2 \vee c, 0.6), (b_0, 0.5), (b_3 \vee c, 0.5), (a_3, 0.4), (b_1, 0.4), (b_4 \vee c, 0.4), (b_2, 0.3), (c, 0.3), (b_3, 0.2), (b_4, 0.1)\}$.

Indeed A is encoded by $\{(a_0, 1), (a_1, 0.8), (a_2, 0.6), (a_3, 0.4)\ (a_4, 0.2)\}$, $\{(a_0, 1), (a_1, 0.8), (a_2, 0.6), (a_3, 0.4)\}$ **and the base associated to** $(c, 0.6) * B$ **is** $K_{(c, .6) * B} = \{(c, 0.3), (b_0, 0.5), (b_1, 0.4), (b_2, 0.3), (b_3, 0.2), (b_4, 0.1), (b_0 \vee c, 0.8), (b_1 \vee c, 0.7), (b_2 \vee c, 0.6), (b_3 \vee c, 0.5), (b_4 \vee c, 0.4)\}$.

Viewing $(\theta \rightarrow A) \wedge ((C, \rho) * B)$ as a fuzzy constraint satisfaction problem, this can be exploited for relaxing it into crisp problems corresponding to the different level cuts of the above possibilistic logic base; see (Moura-Pires and Prade, 1998). See also (Moura-Pires et al., 1998).

4. Constraint-based specification of preferences

The possibilistic framework can be, useful in qualitative preference profile elicitation from a set of constraints specifying them in a granular way. For instance, a preference in favor of a binary property q w.r.t. not q can be expressed by a constraint of the form

$$\Pi(q) > \Pi(\neg q) \qquad (6),$$

which is equivalent to say that there exists at least one decision value in the set of models of q which is better than all the decision values in the set of models of - q .

This is rather a weak manner for expressing the preference about q. Indeed, due to the definition of a possibility measure, (6) expresses that the most satisfactory candidate satisfying p, is preferred to the most satisfactory candidate, (hence to all candidates), not satisfying p.

Such a constraint can be easily made context dependent: the requirement if p is satisfied, q is preferred to $\neg q$, can be expressed by the constraint

$$\Pi(p \wedge q) > \Pi(p \wedge \neg q). \qquad (7)$$

More generally, a collection of such requirements gives birth to possibilistic constraints, whose greatest solution π^* (in the sense that $\pi^* \geq \pi$ for any solution π) can be computed and represents a preference profile agreeing with the requirements. The minimal specificity principle expresses that any candidate is satisfactory inasmuch it complies with the constraints. However, there may exist other worth-considering selection procedures of a particular possibility distribution satisfying the set of constraints; this is open to discussion.

This approach is formally the same as the possibilistic treatment of default rules. Indeed, a default rule "if p then generally q" is translated into the constraint $\Pi(p \wedge q) > \Pi(p \wedge \neg q)$ which expresses that p and q true is strictly more plausible than p true and q false. A set of consistent default rules of the form "if p_i then generally q_i" is thus represented by a set of constraints like (7) which implicitly defines a set of possibility measures. The greatest solution Π^* of this set always exists. Then, applying the minimal specificity principle (which here amounts to keep the level of normality for each possible state of the world as great as permitted by the available knowledge), it induces a plausibility ordering on the interpretations encoded by the associated possibility distribution π^* . This ordering can be encoded at the formula level by constraints of the form $N^*(\neg p \vee q) \geq \alpha$, i.e., by a possibilistic logic base, where N^* is the dual measure associated with Π^*. See (Benferhat et al., 1992, 1997a) for details.

Example 1 (continued)

In Section 2.3., the set of stratified goals C_1, C_2, C3 was directly

However, such a stratification can be related to the possibility distribution which can be selected from a set of constraints of the form (6)-(7). instance if an agent expresses that he wants coffee, and if coffee is not available he would like tea. This corresponds to the possibilistic base $K = \{(c_1, 1), (\neg c_1 \vee c_2, \alpha)\}$ with $\alpha < 1$, or equivalently $K = \{(c_1, 1), (c_2, \alpha)\}$, which corresponds the least specific solution of the two following constraints $\Pi(c_1) > \Pi(\neg c_1)$, $\Pi(c_1 \wedge c_2) > \Pi(c_1 \wedge \neg c_2)$, where c_1=coffee and C_2=tea.

Other types of constraints can be introduced, for instance for expressing contextual indifference as $\Pi(p \wedge q) = \Pi(p \wedge \neg q)$, or for expressing forms of independence as in (Dubois et al., 1997). Besides, a stronger counterpart of (6) is:

$$\Delta(q) > \Pi(\neg q) \qquad (8)$$

where $\Delta(q) = \min_{u:\, u \text{ model of } p} \pi(u)$ is the guaranteed possibility function (Dubois and Prade, 1998). (8) expresses that any candidate satisfying q is preferred to any candidate satisfying $\neg q$. This is the ceteris-paribus principle. See (Boutilier et al., 1997) for a similar approach to preference modelling, although not expressed in the possibility theory framework.

As suggested by the following example, a similar approach can be useful for completing orderings which are implicitly specified through both examples and general principles.

Example 4: Let us consider the following situation with three criteria, namely the levels in mathematics (M), in physics (P), and in literature (L), and three candidates A, B and C rated on the 6 level scale $a > b > c > d > e > f$:

	M	P	L
A	a	b	f
B	f	e	a
C	d	c	c

where M and P are supposed to have the same importance, greater than the one of L, while the result of the global aggregation of the three criteria should be such that the candidate C is preferred to A and A is preferred to B1. This can be expressed by the following sets of constraints:

1 This example has been recently used by Michel Grabisch and Marc Roubens (with a=18, b=16, c=15, d=14, e=12, MO) for illustrating the case where no weighted average aggregation function can agree with both the proposed orderings between the candidates and the respective importance of the criteria, while a Choquet integral can represent the situation.

i) $\pi(dcc) > \pi(abf) > \pi(fea)$ (C is preferred to A and A is preferred to B)
ii) $\pi(xyz) = \pi(yxz)$ for all x, y and z (M and P have the same importance)
iil) $\pi(xyz) > \pi(xzy)$ for all x if $y > z$ (P is more important than L)
tv) $\pi(xyz) > \pi(zyx)$ for ally if $x > z$ (M is more important than L) non-decreasing
y) π is decreasing w.r.t. x, y and z (the greater the grades, the better the candidate).

where $\pi(xyz)$ denotes the level of acceptability of having grade x in M, y in P and z in L (using an encoding of the grades x, y, z into the 6 level scale $a > b > c > d > e > f$). Please note that π just encodes a ranking and $\pi(xyz)$ is not an absolute value. Such a family of constraints defines a family of π-rankings compatible with the constraints. This family is non-empty if the constraints are consistent (which is the case in the example). Note that the constraint (i) reflects the example provided the user, while the others express general principles which should be applied to any tuples of grade in M, P and L. Then we may think of using a selection principle such as putting each triple xyz at the higher possible rank compatible with the constraints. Let π^* be the selected ranking; because of (v), it is equivalent to a stratified set of propositions expressing prioritized goals of the form " the grade in M is at least θ, the grade in P is at least θ', and the grade in L is at least θ'' ". This provides another reading of the preferences implicitly specified by means of the ranking "C is preferred to A and A is preferred to B" and the relative assessment of the importance of the three criteria. This new expression of the preferences can be presented to the user for verification.

Note that the approach only looks for a ranking between triples of grades in M, P and L, without trying to get this ordering by means of some aggregation function to be determined in a given family (e. g., choquet integral (Grabish et al., 1995)), as classical approaches do. The development of such an approach raises computational issues which are not addressed here.

Let us emphasize that the interest of such an approach would be to obtain a ranking of the situations without having to identify an aggregation function for the criteria grades. It also enables us to check the consistency of the user requirements and to restate the preferences as a set of stratified goals (which may be checked by the user).

5. Revising preferences

This section briefly points out another issue where the possibilistic logic representation of preferences can be useful: the modelling of the dynamics of preference.

The revision of a possibilistic belief base K by an input *a* (whose set of models is A), held as certain, amounts at the semantical level at defining the conditional possibility distribution (e.g., (Dubois and Prade, 1998))

$$\pi_K(u|A) = 1 \text{ if } \pi_K(u) = \Pi_K(A),\ u \in A,$$
$$= 0 \text{ if } u \notin A,$$
$$= \pi_K(u) \text{ if } \pi_K(u) < \Pi_K(A),\ u \in A.$$

where Π_K is the possibility measure associated with π_K. Let $(K)_\alpha$ be the set of classical propositions in the possibilistic logic base K with a level strictly greater than a, and let a be the smallest degree such that $(K)_\alpha \cup \{a\}$ is consistent, then a is called the level of inconsistency of $K \cup \{(a,1)\}$; the formulas in the layers of K with a degree smaller than a are forgotten in the revision process (even if they are not involved in the inconsistency). Indeed at the syntactical level, K is changed into $(K)_\alpha \cup \{(a,1)\}$ when K is revised by the input *a.*

Viewing a possibilistic logic base as a set of goals with their level of priority, rather than as a set of uncertain pieces of information as it is usual in belief revision* it enables us to express the dynamics of the preferences when a new imperative goal *a* is added to K. More generally, we may think of applying rules for belief revision under uncertain inputs for modelling changes in preference as suggested by Ryan and Williams (1997). The role of the uncertain input is then played by a new preference profile (under the form of a prioritized set of goals). As emphasized in (Dubois and Prade, 1997), an uncertain input information $(A;\alpha)$ is not understood in the same way whether it is a constraint or an unreliable input. In the first case, it forces the revised profile (modelled by Π') to satisfy α (i.e., $\Pi'(A) = 1$ and $\Pi'(\bar{A}) = 1 - \alpha$), where A is the complement of A; the following revision rule respects these constraints

$$\pi(u\ A;\alpha)) = \max(\pi(u|A), \min((1-\alpha), \pi(u \mid \bar{A}))) \quad (9)$$

Note that when $\alpha = 1$, $\pi(u \mid (A;\alpha)) = \pi(u \mid A)$, but when a = 0, we obtain a possibility distribution less specific than π, such that $N(A) = N(\bar{A}) = 0$ (where N is associated with $\pi(\cdot \mid (A;\alpha))$, which corresponds to a complete lack of priority concerning A.

When $\alpha > 0$, rule (9) expresses that the most satisfactory candidates in A become fully satisfactory, the most satisfactory candidates in A are forced to level $1 - \alpha$ and all other candidates that were originally more satisfactory than $1 - \alpha$, if any, are forced to level $1 - \alpha$ as well. This operation minimizes changes of the satisfaction levels of candidates so as to accommodate the constraint N'(A) $= \alpha$. Only firmly entrenched preferences are left untouched. Rule (9) can be extended to a set of input constraints $\{\Pi(A_i) = \lambda_i,\ i = 1,n\}$ where $\{A_i,\ i = 1,n\}$ forms a partition of Ω, such that $\max_{i=1,n} \lambda_i = 1$ (normalisation). It gives the following rule

$$\pi(u \mid \{(A_i;\lambda_i)\}) = \max_i \min(\lambda_i, \pi(u \mid A_i)). \quad (10)$$

Rules (9) and (10) are qualitative counterpart of rules proposed by Spohn (1988), for revising uncertain information; they have been also proposed more recently by Williams (1994) under a different but equivalent form. In the second case (A,α) is viewed as an unreliable input, represented by the weighted *nested* pair of subsets F = $\{(A,1), (U, 1 - \alpha)\}$ where the weights denote degrees of possibility and U is the whole set of candidates. The revised profile $\pi(\cdot \mid F)$ is defined by formal analogy with (9):

$$\pi(u \mid F) = \max(\pi(u \mid A), \min(\pi(u), (1 - \alpha))) \quad (11)$$

Note the difference with (9): there is no conditioning on A $(\pi(u) = \pi(u \mid U))$. However, contrary to (911, the equality $N(A \mid F) = \alpha$ is not warranted since $N(A \mid F) =$ N(A) whenever $N(A) > \alpha$. Lastly, $\pi(u \mid F) = \pi(u)$ if $\alpha = 0$ since then $F = U$: no revision takes place. This behavior is very different from the case when the uncertain input is taken as a constraint. Besides, these two types of revision can be performed directly on the corresponding possibilistic logic bases; see (Dubois and Prade, 1997).

However, the transposition of belief revision techniques, especially under uncertain inputs, to preference dynamics is not completely straightforward. Indeed, the above rules (9) and (10) suppose that the uncertain input is defined on a partition, while there is no reason for the preference input to have this particular structure. Rule (11) might be preferred for preference dynamics. However if we want to give strong priority to the contents of the input, it would be advisable to revise the input by the previous preference profile, an unusual procedure in belief revision! This is open to discussions.

6. Concluding remarks

This paper has been advocating the use of possibilistic logic in various aspects of decision analysis where representation issues are important. The qualitative handling of preferences, their symmetric combination as well as their revision have been addressed.

In connection with the new approach proposed in this paper, one may think of other lines of research. First, the logical framework does not only provide a convenient representation tool, but also provides a basis for generating explanations of interest for the user.

Besides, another worth investigating issue, where a layered logic framework may be useful, is the analysis of conflict between preferences. Suppose that different preference profiles, expressing different points of view, are to be combined symmetrically. Taking these different preference profiles together very often creates inconsistencies (see, e.g., Felix, 1992). The problem is then to determine what goals can be relaxed or put at smaller levels of priority, taking advantage of the stratification of the preferences. Methods developed for reasoning from stratified inconsistent prepositional logic bases may be very useful for that purposes: these methods are based on the selection of particular consistent subbases, or on the research of arguments pro and cons (Benferhat, Dubois, and Prade, 1996), or on the exploitation of minimally inconsistent subsets (Benferhat and Garcia, 1997).

References

Bellman R., Zadeh L.A. (1970) Decision-making in a fuzzy environment. Management Sciences, 17,141-164.

Benferhat S., Dubois D., Prade H. (1992) Representing default rules in possibilistic logic. Proc. KR'92,673-684.

Benferhat S„ Dubois D., Prade H. (1996) Reasoning in inconsistent stratified knowledge bases. Proc. of the 26 Inter. Symp. on Multiple-Valued Logic (ISMVL'96), Santiago de Compostela, Spain, 29-31 May, 184-189.

Benferhat S., Dubois D., Prade H. (1997a) Nonmonotonic reasoning, conditional objects and possibility theory. Artificial Intelligence, 92,259-276.

Benferhat S., Dubois D., Prade H. (1997b) From semantic to syntactic approaches to information combination in possibilistic logic. In: Aggregation and Fusion of Imperfect Information, Studies in Fuzziness and Soft Computing Series, (B.Bouchon-Meunier, ed.), Physica. Verlag, 141-161.

Benferhat S., Garcia, L. (1997) Dealing' with locally-prioritized inconsistent knowledge bases and its application to default reasoning. In Applications of Uncertainty Formalisms (T. Hume, and S. Parsons eds.), LNA11455, Springer.

Boutilier C.(1994) Toward a logic for qualitative decision theory Proc. of the 4th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR-94), Bonn, (J.Doyle, E.Sandewall, P.Torasso, eds.), Morgan Kaufmann, 75-86.

Boutiiier C, Brafman R., Gcib G, Poole D. (1997) A constraint-based approach to preference eiicitation and decision making. Working Notes of the AAAI'97 Spring Symp. Series on Qualitative Preferences in Deliberation and Practical Reasoning, Stanford, C A, Mar.24-26,19-28.

Dubois D., Farinas L., Herzig A., Prade H., (1997) Qualitative relevance and independence: a roadmap. Proc. UCAI-97,62-67.

Dubois D., Lang J., Prade H. (1994) Automated reasoning using possibilistic logic: semantics, belief revision and variable certainty weights. IEEE Trans, on Data and Knowledge Engineering, 6(1), 64-71.

Dubois D., Le Bene D., Prade H., Sabbadin R. (1998) Logical representation and computation of optimal decisions in a qualitative setting. Proc. AAAI-98,588-593.

Dubois D., Prade H. (1997) A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. Int.J. Approx. Reasoning, 17,295-324.

Dubois D., Prade H. (1998) Possibility theory: qaulitative and quantitative aspects. In Handbook of defeasible reasoning and uncertainty management systems. Vol 1, PP. 169-226, Kluwer Academic Press.

Felix R. (1992) Towards a goal-oriented application of aggregation operators in fuzzy decision-making. Proc. of the Int. Conf. on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU-92), Mallorca, July 6-10,585-588.

Grabish M, H. T. Nguyen, and E. A. Walker (1995) Fundamentals of uncertainty calculi, with applications to fuzzy inference. Kluwer Academic.

Lacroix M., Lavency P. (1987) Preferences: Putting more knowledge into queries. Proc. of the 13rd Inter. Conf. on Very Large Data Bases, Brighton, UK, 215-225.

Lang J. Conditional desires and utilities - an alternative logical framework for qualitative decision theory. Proc. 12th European Conf. on Artif. Intellig.(ECAI-96), Budapest, Wiley, U.K., 318-322.

Moura Pires J., Prade H. (1998) Logical analysis of fuzzy constraint satisfaction problems. Proc. of the 1998 IEEE Int. Conf. on Fuzzy Systems (FUZZ-IEEE'98), Anchorage, Alaska, May 4-9,1998,857-862.

Moura-Pires J., Dubois D., Prade H. (1998) Fuzzy constraint problems with general aggregation operations under possibilistic logic form. Proc. 6th Europ. Cong, on Inteilig. Techniques & Soft Comput., Aachen, Germany, Sept. 7-10,1998, pp. 535-539.

Ryan J., Williams, M.-A. (1997) Modelling changes in preference: an implementation. ISRR-027-1997, Dept. of Management, Univ. of Newcastle, NSW, Australia

Spohn W. (1988) Ordinal conditional functions: a dynamic theory of epistemic states. In: Causation in Decicion, Belief Change and Statistics, Vol. 2 , (W.L.Harpcr and B.Skyrms, eds.). Reidel, Dordrecht, 105-134.

Williams, M.-A. (1994) Transmutations of knowledges systems. Proc. KR-94,619-629.

Zadeh L.A. (1978) Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems, 1,3-28.