Flexible Constraints and Qualitative Decision in A.I.

(extended abstract)

Didier Dubois and Henri Prade

Institut de Recherche en Informatique de Toulouse (IRIT) – CNRS
Université Paul Sabatier, 118 route de Narbonne
31062 Toulouse Cedex 4, France

1. Flexible constraints

The notion of constraint is basic in operations research and artificial intelligence. A constraint describes what are the potentially acceptable decisions (the solutions to a problem) and what are the absolutely unacceptable ones: it is an all-or-nothing matter. Moreover, no constraint can be violated, i.e., a constraint is classically considered as imperative. Especially the violation of a constraint cannot be compensated by the satisfaction of another one. If a solution violates a single constraint, it is regarded as unfeasible. The idea of flexible constraints is to keep the non-compensatory property of constraints, while introducing intermediary levels between feasibility and non-feasibility as well as levels in the imperativeness of constraints.

A classical hard constraint C is represented by a classical set of solutions, i.e., only using degrees of membership $\mu_C(d) = 1$ if choosing d is feasible and $\mu_C(d) = 0$ if it is not. Introducing intermediary levels of feasibility on a linearly ordered valuation set S (with 0 as a bottom element and 1 as a top element), C is then called a fuzzy, or soft constraint: $\mu_C(d) = 1$ means that a solution d totally satisfies C while $\mu_C(d) = 0$ means that it totally violates C (d is unfeasible). If $0 < \mu_C(d) < 1$, d satisfies C only partially; $\mu_C(d) > \mu_C(d')$ indicates that C is more satisfied by d than by d'. Hence, like an objective function, a fuzzy constraint rank-orders the feasible decisions. However, contrary to an objective function a fuzzy constraint also models a threshold (represented by the bottom level 0) beyond which a solution will be rejected. In fact, a fuzzy constraint can be viewed as the association of a constraint (defining the support of C) and a criterion which rank-orders the solutions satisfying the constraints. In this interpretive framework, a membership function, is similar to a qualitative utility function, or better, a value function. A soft constraint C' will be looser than another one C if and only if $\mu_C \le \mu_{C'}$, that is, if any solution to C is at least as feasible for C'.

In the presence of very loose constraints it is clear that the use of soft constraints can help breaking ties among feasible solutions, just as objective functions do. On the contrary when hard constraints are tight, there is no feasible solution. One way out is to relax the constraints, but automating this process is not so easy and is usually time-consuming. What is usually done is to assign priorities to constraints. Solutions that satisfy all constraints are selected if any. Otherwise the less prioritary constraints are dropped, and the chosen solution satisfies only higher priority constraints.

A priority level ρ for constraint C can be modeled in the framework of possibility theory by a degree of necessity estimating how imperative C is. Indeed, the pair (C, ρ), where C is a crisp constraint, can be represented by a special kind of fuzzy constraint C* (Dubois, Fargier and Prade, 1996a):

$$\mu_{C^*}(d) = 1 \text{ if d satisfies C,}$$
$$= 1 - \rho \text{ if d violates C.}$$

Then it can be checked that N(C) = ρ, where N is a necessity measure defined by N(A) = $\min_d \max(\mu_A(d), 1 - \mu_{C^*}(d))$. Since ρ represents to what extent it is necessary to satisfy C, 1 – ρ indicates to what extent it is possible to violate it. In other words, any potential solution u that violates C satisfies (C,ρ) to a degree equal to 1 – ρ. When possibility degrees are interpreted in terms of feasibility, necessity degrees are thus levels of priority. If C is itself a soft constraint with priority r, it can be modeled by the fuzzy constraint C* defined by:

$$\mu_{C^*}(d) = \max(1 - \rho, \mu_C(d)). \qquad (1)$$

As indicated by (1), compatibility of d with (C, ρ) is 0 only if $\rho = 1$ and if $\mu_C(d) = 0$, i.e., if constraint C is not satisfied at all but is fully important. Moreover if $\rho = 0$ (C has no importance at all), (1) yields the top element 1 of the valuation set for any d, as expected. As shown by (1), even if C is completely violated, the global level of satisfaction w.r.t. all the constraints is only upper bounded by 1 – ρ because of this violation (instead of 0 in case C is imperative). In practice, it may be useful to add a "safeguard" constraint C° in case some decisions among the ones which violate C are really inacceptable (they will be such that $\mu_{C^\circ}(d) = 0$), thus changing (1) into

$$\mu_{C^*}(d) = \min(\mu_{C^\circ}(d), \max(1 - \rho, \mu_C(d)). \qquad (2)$$

There is a different way of relaxing a constraint by considering that the constraint is sufficiently satisfied if the level of satisfaction for u reaches some threshold θ, i.e., $\mu_C(d) \geq \theta$. Then, we shall speak of a discounted constraint. Then $\mu_C(d)$ will be changed into an evaluation $\mu_{C'}(d) = 1$. If $\mu_C(d) < \theta$ we may consider that the constraint is satisfied at the level which is actually reached, i.e., $\mu_{C'}(d) = \mu_C(d)$ (the use of a *relative* level of satisfaction $\mu_C(d) / \theta$ would require a numerical valuation set like [0,1] and not a simple linearly ordered valuation set). Then we have

$$\mu_{C'}(d) = \theta \rightarrow \mu_C(d) \qquad (3)$$

where $a \to b$ is Gödel implication ($a \to b = 1$ if $a \le b$, $a \to b = b$ if $a > b$), while (1) exhibits Dienes' implication $a \to b = \max(1 - a, b)$.

When using finite valuation scales, which is enough in many practical situations, a fuzzy set can be equivalently seen as a finite family of nested level cuts, corresponding here to crisp constraints or objectives. In the simple case of a unique fuzzy constraint C, defined by its membership function μ_C ranging on the finite scale $\{\alpha_0 = 0 < \alpha_1 < \ldots < \alpha_n = 1\}$, C is equivalently represented by the collection of its α_i-cuts $C_{\alpha_i} = \{d \text{ s.t. } \mu_C(d) \ge \alpha_i\}$.This is still equivalent to the set of constraints $N(C_{\alpha_i}) \ge 1 - \alpha_{i-1}$ for i=1,n. In other words, the goal of finding a solution in C_{α_i} has priority $1 - \alpha_{i-1}$, and the larger the α_i-cut, the more important the priority; in particular it is imperative to find a solution in the support of C (i.e., $\{d \text{ s.t. } \mu_C(d) > 0\}$). This gives birth to a possibilistic logic knowledge base (Dubois, Lang and Prade, 1994), of the form $(c_{\alpha_i}, 1 - \alpha_{i-1})$ for i=1,n where c_{α_i} denotes the proposition whose set of models is C_{α_i}. Note that $C_{\alpha_i} \supset C_{\alpha_{i+1}}$ for i=0,n-1. Preferences are thus expressed in terms of sets of crisp (nested) goals having different levels of priority. Moreover we have

$$\mu_C(d) = \min_i \max(\mu_{C_{\alpha_i}}(d), \alpha_{i-1}). \quad (4)$$

The modification of a constraint by a level of importance ρ, or by a threshold θ, is then easy to handle:

- the importance weighting operation $\max(\mu_C(d), 1 - \rho)$ translates into the suppression of the most prioritary constraints $(c_{\alpha_i}, 1 - \alpha_{i-1})$ such that $1 - \alpha_{i-1} > \rho$.
- the thresholding operation $\theta \to \mu_C(d) = 1$ if $\mu_C(d) \ge \theta$ and $\theta \to \mu_C(d) = \mu_C(d)$ if $\mu_C(d) < \theta$ translates into the suppression of the least prioritary constraints $(c_{\alpha_i}, 1 - \alpha_{i-1})$ such that $\alpha_i > \theta$.

More generally, the conjunctive aggregation of fuzzy (prioritized or thresholded) constraints 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. Conversely, a set of crisp goals with different levels of priority can be always represented in terms of a fuzzy set membership function.

2. Min-based aggregation and optimisation

The notion of fuzzy set as viewed by Bellman and Zadeh (1970) means to represent constraints by fuzzy subsets C of possible decisions. If $C_1, \ldots, C_m$ denote m fuzzy constraints, the fuzzy decision set is then defined by

$$\mu_D(d) = \min_{i=1,m} \mu_{C_i}(d). \quad (5)$$

However, this definition of a decision set is not really in accordance with the usual paradigm of multiple-criteria decision-making since an optimal solution in the sense of Bellman and Zadeh does not make a trade-off between the membership values $\mu_{C_i}(d)$. On the contrary an optimal solution is one that least violates the most violated constraint. Hence, Bellman and Zadeh's proposal is a constraint-directed view of problem-solving, where constraint are flexible: if there is a solution u that completely satisfies all constraints, this solution is optimal. Otherwise $\lambda^* = \sup_d \mu_D(d) < 1$, and this indicates that the constraints are partially contradictory. To accept a solution u such that $\mu_D(d) = \lambda^*$ means to partially relax some of the constraints C_i (with respect to the constraint corresponding to the core of C_i, i.e., $\{d, \mu_{C_i}(d) = 1\}$). In that sense constraints C_i are flexible.

This notion of optimality is not new in the literature of decision making. It is known as the maximin stategy and is well-known in game theory as the most cautious strategy for players. It is also well-known in decision-making under uncertainty, where again it is a cautious strategy. Despite these formal analogies, the problem tackled by flexible constraint satisfaction is very different, since index i pertains neither to an opposing player nor to a state of the world. The paradigm most akin to Bellman and Zadeh's is to be found in social choice theory (Moulin, 1988) and $\min_{i=1,m} \mu_{C_i}(d)$ is then an egalitarist social welfare function, $\mu_{C_i}(d)$ being the welfare index of an individual i. Maximizing $\min_{i=1,m} \mu_{C_i}(d)$ tends to select solutions which equally satisfy all the constraints.

Like averaging operations, min operations (as well as the DeMorgan dual disjunctive connective 'max') can be weighted. In the qualitative setting, thresholding constraints in order to acknowledge their respective levels of importance may also help breaking ties between possible decisions (especially when $\forall d, \min_i \mu_{C_i}(d) = 0$). Importance is here an attribute attached to each constraint, either in terms of priority or of discount as seen above. Thus, relaxing constraints by assuming that all the constraints have not the same importance can be done in different ways in the context of conjunctive aggregation.

An example of thresholded min aggregation is thus obtained as the min-based conjunction of prioritized constraints:

$$\mu_D(d) = \min_i \max(1 - \rho_i, \mu_{C_i}(d)). \qquad (6)$$

Note that now $\mu_D(d) \geq \min_{i=1,m} \mu_{C_i}(d)$, which indeed expresses a relaxation. The convention $\max_{i=1,m} \rho_i = 1$ is assumed in order to have an appropriate scaling of the levels of importance. Denoting by W the fuzzy set of important constraints (i.e., $\mu_W(C_i) = \rho_i$), (6) expresses a strong type of inclusion of W into the fuzzy set of criteria satisfied by choice d, S_d (where $\mu_{S_d}(C_i) = \mu_{C_i}(d)$), based on Dienes implication $a \to b = \max(1 - a, b)$. Indeed $\mu_D(d) = 1 \Leftrightarrow \text{support}(W) \subseteq \text{core}(S_d)$, i.e., if all the constraints with a non-zero importance are among the constraints fully satisfied by d.

With (3) instead of (1) in (6), we still have $\mu_D(d) \geq \min_{i=1,m} \mu_{C_i}(d)$. However, it contrasts with the use of Dienes implication in (6), and $\mu_D(d) = 1 \Leftrightarrow \forall i, \mu_{C_i}(d) \geq \theta_i$ expresses an ordinary fuzzy set inclusion of W into S_d (with now $\mu_W(C_i) = \theta_i$).

The maxmin ordering of solutions is coarse but can be refined on a qualitative basis either by focusing on the least satisfied discriminating constraint (the "discrimin" ordering), or using lexicographical techniques such as the leximin ordering (Dubois et al., 1996b).

The Pareto ordering corresponds to the notion of fuzzy set inclusion: $\vec{s} \geq_{\text{Pareto}} \vec{t} \Leftrightarrow S_e \subseteq S_d$ where $\vec{s} = (s_1,...,s_m)$ with $s_i = \mu_{C_i}(d)$ and $\vec{t}$ is similarly defined changing d into e. Thus, it is not a refinement of the min-ordering ($>_{\min}$). Moreover the discrimination power of the min-ordering is low as already said, while the Pareto-ordering leads to incomparabilities. Keeping in mind the idea of inclusion, a refinement of these two orderings can be defined by comparing the decisions on the basis of level-cut inclusion. Hence, we consider a new ordering, $>_{\text{Discrimin}}$, defined in the finite case by:

$\vec{s} >_{\text{Discrimin}} \vec{t}$ iff $\exists\ \alpha \in S$ s.t.

(i) $\forall\ \beta \in S$ s.a. $\beta < \alpha$, $(S_d)_\beta = (S_e)_\beta$;

(ii) $(S)_\alpha \supset (S_e)_\alpha$,

where $(S_f)_\gamma = \{C_i\ /\ \mu_{C_i}(f) \geq \gamma\}$ is a level-cut of S_f. The relation $>_{\text{Discrimin}}$ is irreflexive and transitive: it is a strict partial ordering. It can be verified that it is a refinement of $>_{\min}$ and of the Pareto ordering. The use of $>_{\text{Discrimin}}$ consists in practice in comparing pointwise satisfaction degrees like in $>_{\text{Pareto}}$, and focusing on the lowest satisfaction degrees among the constraints satisfied *at different degrees* by the competing decisions: decisions are compared on the basis of the least satisfied discriminating constraints. This seems to be in accordance with the intuition, since only the constraints making a difference between two alternatives can help make a final decision. Let us note $D(\vec{s}, \vec{t}) = \{C_i\ /\ s_i \neq t_i\}$ the set of constraints which are satisfied by $\vec{s}$ and $\vec{t}$ to a different extent. An equivalent characterization of $>_{\text{Discrimin}}$ is:

$$\vec{s} >_{\text{Discrimin}} \vec{t} \Leftrightarrow \min_{C_i \in D(\vec{s},\vec{t})} s_i > \min_{C_i \in D(\vec{s},\vec{t})} t_i \ . \qquad (7)$$

The "discri-min" refinement is based on the idea that the constraints on which two decisions receive the same evaluation, have no importance when comparing the decisions.

Another refinement of the egalitarist maximin ordering has been proposed (e.g., Moulin, 1988): the leximin ordering. The idea is to keep all the information pertaining to S_d, instead of summarizing it by a collective utility function as with $>_{\min}$, by considering a ranked multi-set of satisfaction degrees. Let $\vec{s}$ be a vector of satisfaction degrees, and denote by $\vec{s}^*$ another such vector such that $s^*_i = s_{\sigma(i)}$ where σ is a permutation such that $s_{\sigma(1)} \leq s_{\sigma(2)} \leq \ldots \leq s_{\sigma(n)}$ obtained by reordering the components of $\vec{s}$. Indeed, two multi-sets built on an ordered referential can always be compared using the leximin ordering $>_{\text{Lex}}$:

- $\vec{s} >_{\text{Lex}} \vec{t}$ iff $\exists\ k \leq m$ such that $\forall\ i < k$, $s^*_i = t^*_i$ and $s^*_k > t^*_k$, (8) where s_i* (resp. t_i*) stands for the i-th component of $\vec{s^*}$ (resp. $\vec{t^*}$). The two possible decisions are indifferent if the corresponding multi-sets are the same.

The leximin-ordering is a refinement of both the Pareto-ordering and the min-ordering:

• $\vec{s} >_{min} \vec{t} \Rightarrow \vec{s} >_{Lex} \vec{t}$; • $\vec{s} >_{Pareto} \vec{t} \Rightarrow \vec{s} >_{Lex} \vec{t}$.

Leximin optimal decisions are always discrimin optimal decisions, and thus indeed min-optimal and Pareto-optimal: $>_{Lex}$ is the most selective among these preference relations. Reciprocal implications are not verified. The leximin ordering can discriminate among equivalent decisions according to the min-preordering ($\min_i s_i = \min_i t_i$) and among incomparable ones with respect to $>_{Pareto}$ and $>_{Discrimin}$. The leximin-ordering tends to favor solutions that violate as *few* fuzzy constraints as possible (where 'few' refers to a fuzzy cardinality evaluation; see (Dubois et al., 1996b)).

3. Decision under uncertainty & multiple criteria decision

In decision under uncertainty, possibilistic logic can be used for modelling the available information about the world on the one hand, and the preferences on the other hand. This section summarizes a recent proposal by (Dubois, Prade and Sabbadin, 1997). We distinguish between two possibilistic logic bases. The first one $K = \{(p_j, \rho_j); j = 1,n\}$ represents the available knowledge about the world. The presence of (p_j, ρ_j) in K means that we are certain at least at level ρ_j that p_j is true, i.e., $N(P_j) \geq \rho_j$. The second one $P = \{(q_i, \lambda_i); i = 1,m\}$ represents the preferences of the decision-maker under the form of a prioritized set of goals: the goal of having q_i true has priority λ_i.

Each potential decision d is represented by a formula (d,1) to be added to K if the decision is chosen. Let $K_d = K \cup \{(d,1)\}$ be the description of what is known about the world when d is applied. Asssociated with K_d is the possibility distribution πK_d defined by the fuzzy set,

$$\pi K_d(u) = \min_{j=1,n} \max(\mu_{P_j}(u), 1 - \rho_j), \quad (8)$$

where $\mu_{P_j}(u) = 1$ if u is a model of p_j and $\mu_{P_j}(u) = 0$ otherwise. This is formally similar to (4). πK_d rank-orders the more or less plausible states of the world when d is chosen. Associated with the layered set of goals P is the ordinal utility function

$$\mu_P(u) = \min_{i=1,m} \max(\mu_{Q_i}(u), 1 - \lambda_i), \quad (9)$$

where $\mu_{Q_i}(u) = 1$ if u is a model of q_i and $\mu_{Q_i}(u) = 0$ otherwise. μ_P rank-orders the different states accorting to their acceptability.

From π_{K_d} and μ_P a qualitative counterpart to expected utility can be computed as

$$E*(d) = \min_u \max(\mu_P(u), 1 - \pi_{K_d}(u)). \qquad (10)$$

$E*(d)$ is all the greater as all the plausible states u according to K_d are among the most preferred states according to μ_P. $E*(d) = 1$ if and only if $\{u \text{ s.t. } \pi_{K_d}(u) > 0\} \subseteq \{u \text{ s.t. } \mu_P(u) = 1\}$, and $E*(d) = 0$ if $\exists$ u s.t. $\pi_{K_d}(u) = 1$ and $\mu_P(u) = 0$. Thus, $E*(d) = 0$ as soon as there exists a state u , i.e., a consequence of d, which is fully plausible and which is abssolutely not acceptable. A von Neumann and Morgenstern-like justification for (10) is given in (Dubois and Prade, 1995). Note that (10) presupposes a commensurability assumption between the uncertainty and the preference scales.

It is possible to compute $E*(d)$ by only using a classical logic machinery on α-level cuts. Indeed it has been shown in (Dubois, Prade and Sabbadin, 1997) that $E*(d)$ can be computed as the maximum of the α such that

$$(K_d)_\alpha \text{ entails } (P)_{\underline{n(\alpha)}} \qquad (11)$$

where $(B)_\alpha$, resp. $(B)_{\underline{\alpha}}$ is the set of classical propositions in a possibilistic logic base B with a level greater or equal to α, resp. strictly greater than α. As it can be seen in (11), $E*(d)$ is equal to 1 (α=1) if the completely certain part of K_d entails the satisfaction of all the goals, even the ones with low priorities, since $P_{\underline{0}}$ is just the set of all the propositions with a non-zero priority level. In (Dubois et al., 1998) a computation procedure using an Assumption-based Truth Maintenance Systems is proposed for computing the best decision in the sense of (10). More generally, stratified logical frameworks such as possibilistic logic, enable the design of anytime algorithms (Williams, 1997), since sub-optimal solutions can be easily obtained at each step of the computation procedure.

The evaluation index (10) is a pessimistic one since it is very demanding. Its optimistic counterpart reads

$$E^*(d) = \max_u \min(\mu_P(u), \pi_{K_d}(u)). \qquad (12)$$

and is equal to 1 as soon as one fully acceptable choice u (i.e., such that $\mu_P(u)=1$) is also completely plausible. In logical terms, $E^*(d)$ is equal to the greatest a such that $(K_d)_{\underline{\alpha}}$ and $(P)_{\underline{\alpha}}$ is logically consistent.

The similarity between (6) and (10) is striking, for a fixed decision d. In (6) criterion i plays the role of state u in (10), $\mu_{C_i}(d)$ the role of $\mu_P(u)$, ρ_i the role of $\pi_{K_d}(u)$. On the one hand, $\mu_D(u) = \min_i \max(1 - \rho_i, \mu_{C_i}(d))$ is a degree of inclusion of the set of important criteria in the set of criteria satisfied by d; on the other hand, $E*(d) = \min_u \max(\mu_P(u), 1 - \pi_{K_d}(u))$ is a degree of inclusion of the set of plausible states reached when decision d is taken in the set of valuable states. See (Dubois et al., 1997) for a general discussion of formal similarities between approaches in multiple criteria decision making (MCDM) and in decision under uncertainty.

This analogy might suggest to consider the following alternative to (10)

$$E'*(d) = \min_u \pi_{K_d}(u) \rightarrow \mu_P(u),$$

where $\rightarrow$ is Gödel implication. Then $E*(d) = 1$ as soon as in each state u, the level of satisfaction to be in state u is at least as large as the plausibility of reaching u by d. This corresponds to a much weaker degree of inclusion of the set of plausible states (reached when decision d is taken) in the set of valuable states, than with (10). As such $E'*(d)$ appears more debatable.

The similarity between MCDM and decision under uncertainty may also suggest to apply the possibilitic logic view expressed by (11) to MCDM. However this is not straightforward, first since in (6) and (10) decision d does not appear at the same place. Although the fuzzy sets of criteria W and S_d introduced in Section 2 can be easily interpreted in terms of α-level cuts, namely the set of criteria whose importance is greater or equal to α and the set of criteria which are satisfied at least at the level α for decision d, the interpretation of the α-level cuts of W and S_d in propositional terms, meaningful for a user, is not obvious in practice.

4. Concluding remark

This extended abstract aims at providing a brief overview of recent works by the authors and their colleagues pertaining to the qualitative representation of flexible and prioritized constraints, preference and uncertainty. Qualitative representations are often more in agreement with the quality of the available information and lead to simple computations. For more details, the reader is referred to overview papers by the authors (Dubois and Prade, 1997,1998).

References

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

Dubois D., Fargier H. and Prade H. 1996a. Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty. *Applied Intelligence*, 6, 287-309.

Dubois D., Fargier H. and Prade H. 1996b. Refinements of the maximin approach to decision-making in fuzzy environment. *Fuzzy Sets and Systems*, 81, 103-122.

Dubois D., Grabisch M., Modave F., Prade H. 1997. Relating decision under uncertainty and MCDM models. *Working Notes of the AAAI'97 Fall Symp. Series on Frontiers in Soft Computing and Decision Systems*, Boston, November, 8-10, 1997 (L.Medsker, ed.), 6-15.

Dubois D., Lang J. and 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 Berre D., Prade H., Sabbadin R. 1998. Logical representation and computation of optimal decisions in a qualitative setting. *Proc. 15th National Conf. on Artif. Intellig.(AAAI-98)*, Madison, Wi,, July 26-30, 1998

Dubois D. and Prade H. 1995. Possibility theory as a basis for qualitative decision theory. *Proc. of the 14th Inter. Joint Conf. on Artificial Intelligence (IJCAI'95)*, Montréal, Canada, Aug. 20-25, 1924-1930.

Dubois D., Prade H. 1997. Constraint satisfaction and decision under uncertainty based on qualitative possibility theory. *Proc. of the 6th Int. Conf. on Fuzzy Sytems (FUZZ-IEEE'97)*, 23-30.

Dubois D., Prade H. 1998. Possibilistic logic in decision. *Proc of the ECAI-98 Workshop "Decision theory meets artificial intelligence - Qualitative and quantitative approaches",* Brighton, Aug. 25th, 1998

Dubois D., Prade H. and Sabbadin R. 1997. A possibilistic machinery for qualitative decision.*Working Notes of the AAAI'97 Spring Symp. Series on Qualitative Preferences in Deliberation and Practical Reasoning*, Stanford, CA, March 24-26, 47-54.

Moulin H. 1988. *Axioms of Cooperative Decision-Making*. Cambridge University Press, Cambridge, UK.

Williams M.A. 1997. Anytime belief revision. *Proc. 15th Int. Joint Conf. on Artificial Intelligence (IJCAI-97)*, Aug. 23-29, 1997, 74-79.

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