

Logical representation and computation of optimal decisions in a qualitative setting †

Didier Dubois, Daniel Le Berre, Henri Prade, Régis Sabbadin

IRIT - Université Paul Sabatier - 31062 Toulouse Cedex (France)

e-mail: {dubois, leberre, prade, sabbadin}@irit.fr

Abstract

This paper describes a logical machinery for computing decisions based on an ATMS procedure, where the available knowledge on the state of the world is described by a possibilistic propositional logic base (i.e., a collection of logical statements associated with qualitative certainty levels). The preferences of the user are also described by another possibilistic logic base whose formula weights are interpreted in terms of priorities and formulas express goals. Two attitudes are allowed for the decision maker: a pessimistic uncertainty-averse one and an optimistic one. The computed decisions are in agreement with a qualitative counterpart to classical expected utility theory for decision under uncertainty.

Introduction

In classical decision theory under uncertainty, the preferences of the decision maker are directly expressed by means of a utility function, while a probability distribution on the possible states of the world represents the available, uncertain information about the situation under consideration. However, it seems reasonable to allow for a more granular expression of both the preferences and the available knowledge about the world, under the form, e.g., of logical expressions from which it would be possible to build the utility and the uncertainty functions.

Many works are concerned with qualitative decision theory under uncertainty. Some approaches consider only all-or-nothing notions of utility and plausibility (Bonet and Geffner (1996)); others use in addition a preference ordering on consequences (Brafman and Tennenholtz (1997)). Boutilier (1994) also uses a plausibility ordering, but focuses only on the most plausible states. Tan and Pearl (1994) use two integer-valued rankings for preference and plausibility and compare

actions pairwise, according to the relative plausibilities of the subsets in which one action “dominates” the other. This pairwise comparison is not representable by a utility function, and the associated preference relation is not generally transitive. A similar kind of pairwise comparison was also studied in (Dubois, Fargier and Prade (1997)).

In the following, we propose two syntactic approaches based on possibilistic logic, the first one being more cautious than the second, for computing optimal decisions. Here gradual uncertainty and preferences are expressed by means of two distinct possibilistic propositional logic bases (which are stratified bases). Then, the semantics underlying the two syntactic approaches are shown to be in agreement with the two qualitative utility functions advocated in (Dubois and Prade (1995)). Then, we recall some background on the ATMS framework, and it is shown how to encode a decision problem as one of label computation. Then a procedure called MPL is described for computing optimal decisions in terms of labels. It relies on a modified Davis and Putnam (1960) (1962) semantic evaluation algorithm, described in (Castell et al. (1996), (1998)). Two algorithms based on the use of this procedure, are proposed for computing optimistic and pessimistic optimal decisions respectively. A preliminary version of the logical representation of decision problems used in the following was presented in a workshop paper (Dubois et al. (1997a)). A longer version of the present paper, explaining all the computational details and providing all the proofs of the propositions is to be published (Dubois et al. (1998)).

Qualitative decision in stratified propositional bases

In this article, upper case letters (K, D, P, H, \dots) denote sets of propositional formulas that can possibly be literals. For any set A of formulas, A^\wedge denotes the logical conjunction of the formulas in A and A^\vee , the logical disjunction. If $H = \{l_i\}$ is a set of literals, $\sim H = \{\neg l_i, l_i \in H\}$.

A decision problem under uncertainty can be cast in a logical setting in the following way. A vocabulary

† This article is dedicated to the memory of Thierry Castell, our colleague and friend, who accidentally died in August 1997. He significantly contributed to the development of the MPL algorithm.

V of propositional variables contains two kinds of variables: decision variables and state variables. Let V_D be the set of decision variables which are controllable, that is, their value can be fixed by the decision-maker. The propositional variables outside V_D are state variables. Making a decision amounts to fixing the truth value of every decision variable (or possibly just a part of them). On the contrary state variables are fixed by nature, and their value is a matter of knowledge by the decision maker. He/she has no control on them (although he/she may express preference about their values).

Let K be a knowledge base (here in propositional logic) describing what is known about the world including constraints relating the decision variables. Let P be another propositional base describing goals delimiting the preferred states of the world. K , and P are assumed to be finite, as is the logical propositional language L under consideration. To get a flavour of the decision procedures, assume first that K and P are classical logic bases, and preferences are all-or-nothing. The aim of the decision problem, described in the logical setting, is to try to make all formulas in the goal set P true by acting on the truth-value of decision variables which control the models of K and P . A good decision d^\wedge (from a pessimistic point of view) is a conjunction of decision literals that entails the satisfaction of every goal in P , when formulas in K are assumed to be true. Therefore, d should satisfy

$$K^\wedge \wedge d^\wedge \vdash P^\wedge. \quad (1)$$

Moreover, the allowed decisions must be such that $K^\wedge \wedge d^\wedge$ be consistent, for if it is not the case, (1) is trivially satisfied. Under an optimistic point of view, we may just look for decisions d which are consistent with the knowledge base and the goals,

$$K^\wedge \wedge d^\wedge \wedge P^\wedge \neq \perp. \quad (2)$$

Note that (2) is verified as soon as (1) is. (2) which is over-optimistic should be used only if there is no solution to (1).

In the logical form of decision problems, the knowledge base may be pervaded with uncertainty, and the goals may not have equal priority. Let us enrich our logical view of the decision problem, by assigning levels of certainty to formulas in the knowledge base, and levels of priority to the goals. Thus we obtain two stratified logical bases that model gradual knowledge and preferences. It has been shown (e.g., Dubois et al. (1994)) that a possibility distribution encodes the semantics of a possibilistic logic base, i.e., a stratified base whose formulas are gathered into several layers according to their levels of certainty or priority. In the following we focus on how a decision problem can be stated, expressing knowledge and preferences in terms of stratified bases. Then we will show that the corresponding semantics can be represented by the qualitative utility introduced in (Dubois and Prade (1995)).

In the whole paper it is assumed that certainty degrees and priority degrees are commensurate, and

assessed on the same (*finite*, as is the language under consideration) linearly ordered scale¹ S . The top element of S will be denoted $\mathbb{1}$, and the bottom element, $\mathbb{0}$. Knowledge and preferences are stored in two distinct possibilistic bases. The knowledge base is $K = \{(\phi_i, \alpha_i)\}$ where $\alpha_i \in S$ ($\alpha_i > \mathbb{0}$) denotes a level of certainty, and the ϕ_i 's are formulas in L where decision literals may appear. The base expressing preferences or goals is $P = \{(\psi_i, \beta_i)\}$, where $\beta_i \in S$ ($\beta_i > \mathbb{0}$) is a level of priority, and the ψ_i are formulas of L (where decision literals may also appear).

Let K_α (resp. P_β) denote the set of formulas with certainty at least equal to α (resp. the formulas with priority at least equal to β). Note that we only consider layers of K (or P) such that $\alpha > \mathbb{0}$ and $\beta > \mathbb{0}$ since $K_{\mathbb{0}} = P_{\mathbb{0}} = L$. In the following we also use the notations $K_{\bar{\alpha}}$ and $P_{\bar{\beta}}$ (with $\alpha < \mathbb{1}$ and $\beta < \mathbb{1}$), for denoting the set of formulas with certainty or priority *strictly* greater than α or β respectively. In particular $K_{\bar{\mathbb{0}}}$ and $P_{\bar{\mathbb{0}}}$ are the sets of formulas in K and P respectively, without their certainty levels. Since the scale S is finite, $K_{\bar{\alpha}} = K_{\alpha'}$, where α' is the level of S just above α (the same holds for P).

Making a decision amounts to choosing a subset d of the decision set $D = \{l_i\}$ where the l_i are distinguished positive literals of the language L . Our objective is to rank-order decisions, which will be done by using a utility function $U : 2^D \rightarrow S$ such that d is not preferred to d' iff $U(d) \leq U(d')$. In the following, we will use two different functions: U_* which agrees with a pessimistic view, and U^* which agrees with an optimistic one.

In the first case (pessimistic view), we are interested in finding a decision d (if it exists) such that

$$K_\alpha^\wedge \wedge d^\wedge \vdash P_\beta^\wedge \quad (3)$$

with α high and β low, i.e., such that the decision d together with the most certain part of K entails the satisfaction of the goals, even those with low priority. d is implicitly assumed to be included in the most certain part of $K \cup d$ (certainty level equal to $\mathbb{1}$). Moreover, $K_\alpha^\wedge \wedge d^\wedge$ should be consistent for the α 's satisfying (3). One way of guaranteeing this consistency requirement is to assume $K_{\bar{\mathbb{0}}}^\wedge \wedge d^\wedge$ is consistent. By convention, utility $\mathbb{0}$ is assigned to every decision d that is not consistent with K . Besides, observe that the values of the β satisfying (3) are necessarily such that $\beta > \mathbb{0}$ (since $P_{\mathbb{0}}^\wedge = L$ is inconsistent). Let n be the order reversing map of scale S . Namely if S is $\mathbb{0} = \alpha_0 < \dots < \alpha_n = \mathbb{1}$ then $n(\alpha_i) = \alpha_{n-i}$.

¹An attempt to relax this assumption has been made in (Dubois, Fargier and Prade (1997)). These authors point out that working without the commensurability assumption leads to a decision method close to rational inference machinery in non-monotonic reasoning. Unfortunately, that method also proves to be either very little decisive or to lead to very risky decisions.

Ideally, d , along with the most certain part of K only ($K_{\mathbb{1}}$), should entail every goal in P , even the least preferred ones ($P_{\mathbb{0}}$). Such a decision should have a maximal utility ($\mathbb{1}$). The worst case would be when a decision is unable, even with the whole knowledge ($K_{\mathbb{0}}$) to entail even only the most preferred formulas of P ($P_{\mathbb{1}}$). Such a decision should have a utility of $\mathbb{0}$.

It can be proved that the solution of the problem of maximizing α and minimizing β in (3) satisfies $\beta = \overline{n(\alpha)}$. Thus, the pessimistic utility of decision d , defined at the syntactic level, shall take the form:

Definition 1

$U_*(d) = \max\{\alpha / K_{\alpha}^{\wedge} \wedge d^{\wedge} \vdash P_{\overline{n(\alpha)}}^{\wedge}, K_{\alpha}^{\wedge} \wedge d^{\wedge} \neq \perp\}$ and if $\{\alpha > \mathbb{0}, K_{\alpha}^{\wedge} \wedge d^{\wedge} \vdash P_{\overline{n(\alpha)}}^{\wedge} \text{ and } K_{\alpha}^{\wedge} \wedge d^{\wedge} \neq \perp\}$ is empty, then $U_*(d) = \mathbb{0}$.

If now we consider the optimistic case, we are interested in finding a decision d such that:

$$K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\beta}^{\wedge} \neq \perp \quad (4)$$

with α and β as low as as possible: The preferred states are among the most plausible ones and are also consistent with the decision. The optimistic utility of d is thus given by

Definition 2

$U^*(d) = \max\{n(\alpha) / K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\alpha}^{\wedge} \neq \perp\}$ and $U^*(d) = \mathbb{0}$ if $\{\alpha < \mathbb{1}, K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\alpha}^{\wedge} \neq \perp\} = \emptyset$.

Observe that $U^*(d) = \mathbb{1}$ iff $K_{\mathbb{0}}^{\wedge} \wedge d^{\wedge} \wedge P_{\mathbb{0}}^{\wedge} \neq \perp$, that is if the decision is consistent with every goal and piece of knowledge. This is of course over-optimistic in the sense that it assumes that goals will be attained as soon as their negation cannot be proved: however (4) can be useful to discriminate solutions d, d' to (3) such that $U_*(d) = U_*(d')$.

Possibilistic semantics of decision in stratified bases

Let us present the semantics underlying the logical expression of decision problems we have adopted. Interpreting the α_i 's (which are attached to the layers of K) as the degrees of necessity of the formulas in the corresponding layers of $K \cup d$, we compute a possibility distribution π_{K_d} over Ω (the set of all the interpretations of the language L), expressing the semantics of $K \cup d$ (see, e.g., (Dubois et al. (1994))): $\forall \omega \in \Omega$, $\pi_{K_d}(\omega) = \min_{(\phi_i, \alpha_i) \in K/\omega \models \neg \phi_i} n(\alpha_i)$ if $\omega \models d^{\wedge}$, and $\pi_{K_d}(\omega) = \mathbb{1}$ if $\{\phi_i/\omega \models \neg \phi_i\} = \emptyset$ and $\omega \models d^{\wedge}$, and $\pi_{K_d}(\omega) = \mathbb{0}$ if $\omega \not\models d^{\wedge}$.

The possibility distribution π_{K_d} rank-orders the interpretations according to their level of possibility/plausibility induced by the levels of certainty of the formulas in K . This semantics agrees with the idea that an interpretation ω is all the less possible as it violates formulas with an higher level of certainty. Note that since $K_{\mathbb{0}}^{\wedge} \wedge d^{\wedge}$ is consistent, π_{K_d} is normalized,

i.e., there exists at least an interpretation ω with degree $\pi_{K_d}(\omega) = \mathbb{1}$.

From P , interpreting the β_i attached to the layers of P as degrees of priority of the formulas in P , we build a utility function μ over Ω in a similar way (ω is all the more satisfactory as it violates no goal with a high priority):

$$\mu(\omega) = \min_{(\psi_j, \beta_j) \in P/\omega \models \neg \psi_j} n(\beta_j)$$

$$\text{and } \mu(\omega) = \mathbb{1} \text{ if } \{\psi_j/\omega \models \neg \psi_j\} = \emptyset.$$

The two utility functions U_* and U^* defined prece-
dently can be expressed in terms of the possibility distribution π_{K_d} and the utility function μ :

Proposition 1 *Semantic expressions of the utilities.*

$$U_*(d) = \max_{\alpha / K_{\alpha}^{\wedge} \wedge d^{\wedge} \vdash P_{\overline{n(\alpha)}}^{\wedge}} \alpha = \min_{\omega \in \Omega} \max(n(\pi_{K_d}(\omega)), \mu(\omega)).$$

$$U^*(d) = \max_{\alpha / K_{\alpha}^{\wedge} \wedge d^{\wedge} \wedge P_{\alpha}^{\wedge} \neq \perp} n(\alpha) = \max_{\omega \in \Omega} \min(\pi_{K_d}(\omega), \mu(\omega)).$$

The semantic expression of $U_*(d)$ is exactly the qualitative utility function introduced in (Dubois and Prade (1995)). These utility functions have been also justified in a Savage-like setting in (Dubois et al. (1997b)). Note that S is an ordinal scale, and decisions computed as above are robust since only \min , \max and the order reversing function of S are used. The ranking of decisions is insensitive to any bijective monotonic transformation of S .

Maximizing $U_*(d)$ means finding a decision d whose highly plausible consequences are among the most preferred ones. $U_*(d)$ is small as soon as it exists a possible consequence which is both highly plausible and bad with respect to preferences. This is clearly an uncertainty-averse and thus a pessimistic attitude. When π_d is the characteristic function of a set A , $U_*(d)$ reduces to: $U_*(d) = \min_{\omega \in A} \mu(\omega)$, which is the Wald criterion, that evaluates the worth of a decision as the worst-case utility. The other utility function $U^*(d)$ corresponds to an *optimistic* attitude since $U^*(d)$ is high as soon as it exists a possible consequence of d which is both highly plausible and highly prized. The two utility functions may not be opposed one against the other; the optimistic utility function should be used to refine the pessimistic one, when the latter proves not to be decisive.

A question may be raised as to the meaning of the different levels of preference or certainty that are assigned to each sentence. It is clear that the *preference* ordering can be directly given by the decision maker. The *uncertainty* ordering may be assessed by a unique agent classifying the sentences into layers of different levels of certainty. In case the knowledge is given by multiple sources, we can suppose that they have levels of reliability (which may be different), and thus rank the sentences according to the levels of reliability of the sources which provide them (all the information given by a source having the same reliability). On the contrary if the sources are equally reliable, but each of them has its own ordering, we have to suppose that there exists a common agreement on the meaning of the layers

of each source, so as to be able to merge the layers of the different sources. Besides, System Z (Pearl (1990)) may also help to rank order pieces of generic conditional knowledge by taking the specificity of formulas into account (Benferhat et al. (1997)).

Computation of decisions

The similarity is striking between the two modes of decision under uncertainty and the two modes of diagnostic reasoning, namely abductive and consistency-based diagnosis solutions (e.g., Console, de Kleer (1992)). It is then tempting to encode a logical decision problem under uncertainty by means of techniques coming from the theory of assumption-based truth maintenance systems (ATMS) initiated by (De Kleer (1986)).

In this section, we give some algorithms based on the use of the MPL procedure (which stands for *Modèles Préférés et Littéraux* in French) described in (Castell et al. (1996)) to solve qualitative possibilistic decision problems.

The MPL procedure

The MPL procedure introduced in (Castell et al. (1996)) does the following: given a logical formula ϕ in conjunctive normal form, involving two types of literals, it computes its projection by restricting to one type of literals; and this projection is the most informative such consequence of ϕ , expressed in disjunctive normal form. It is shown that nogoods in an ATMS are easily obtained by means of this procedure.

Principle of the MPL algorithm A (Davis and Putnam (1960)) algorithm enumerates the interpretations of a knowledge base K until it finds a model (consistent case) if any (the inconsistent case is when it finds no model). So doing, it is obvious that searching for models is closely related to finding a disjunctive normal form for a knowledge base, since it is easy to exhibit models of a DNF. Let H be a *consistent* set of literals ($\forall l_i \in H, \neg l_i$ does not belong to H). $V_H \subseteq V$ is the set of variables involved in H . Davis and Putnam's algorithm builds a binary search tree over V , starting with the instantiation of the variables in V_H . At each node it branches on the truth value of a variable. An interpretation I over V is then a path from the root node to a leaf of the tree. Equivalently, it is a set of literals. For any interpretation I we can define its restriction over the set H by: $R_H(I) = I \cap H$. If I is a model of K , $R_H(I)$ is called H-restricted model of K .

The MPL algorithm tries to find (if it exists) a H-restricted model of K , whereas the Davis and Putnam's one searches for a complete one (or a V-restricted one, stated differently). Moreover if it finds one, it is minimal with respect to set inclusion. It is so because the algorithm goes *depth-first* through the binary tree starting with literals in $\sim H$, building the current interpretation I by adding to it literals of $\sim H$ and checking its consistency with K . If $\{l_i, l_j \in \sim H\}$ is consistent

with K , then $R_H(I) = \sim H \cap H = \emptyset$ is of course minimal. If it is not the case, it is at least guaranteed that the first model I found is such that $R_H(I)$ is minimal (because of the depth-first aspect of the algorithm and the fact that the $\sim H$ part of it is explored first). Moreover, the MPL algorithm aims at computing *every* H-restricted model of K . In order to perform this computation, it does not stop after the first model is found, but instead it goes through the whole tree. As soon as a (minimal) H-restricted model $R_H(I)$ is found, the clause $C = \bigvee_{l \in R_H(I)} \neg l$ is added to K in order to eliminate every other H-restricted models containing $R_H(I)$ (since they are not minimal), and the exploration of the tree goes on from the same point. In this way, the whole set of minimal H-restricted models of K , denoted $MP_H(K)$, is found.

The computation of $MP_H(K^\wedge)$ is performed by a call to the procedure MPL (see Castell et al. (1996)). MPL has three arguments: K , KA and H , KA being the set of clauses that are added during the run of MPL (it contains the C clauses described above). We get that $MP_H(K^\wedge) = MPL(K, KA = \{\}, H)$, and KA contains exactly $\neg MP_H(K^\wedge)$ which is a CNF form of the nogoods after a call to $MPL(K, KA = \{\}, H)$. Note that KA is generally initially empty except, as we will see, when MPL is used for performing a label computation.

Application to ATMS Indeed, the basic elements of an ATMS, labels, nogoods, will be efficiently computed by the MPL() algorithm, without any minimization step contrary to De Kleer's original one, due to the following properties (first proposed in (Castell et al. (1996))).

Proposition 2 *The set of nogoods of a knowledge base K with respect to a set of hypotheses H is exactly $MP_H(\neg MP_{\sim H}(K^\wedge))$.*

Proof: E is a nogood iff $K^\wedge \vdash (\sim E)^\vee$ and E is minimal – $MP_{\sim H}(K^\wedge) \vdash (\sim E)^\vee$ – $E^\wedge \vdash \neg MP_{\sim H}(K^\wedge)$. So, the DNF form of the nogoods satisfies the property.

Proposition 3 *Let K be a set of clauses and H a set of hypotheses. Let ϕ be a formula. The label of ϕ exactly contains the elements of $MP_H(\neg MP_{\sim H}(K^\wedge \wedge \neg \phi))$ that are not among the nogoods of K .*

Proof: Similar to the one above, as $label(\phi) = nogoods(K^\wedge \wedge \neg \phi)$.

We remove the nogoods from the label because every formula is a logical consequence of an inconsistent one. This is done by initializing KA with the set of nogoods, in the MPL algorithm.

Proposition 4 *Let K be a set of clauses. Let ϕ and ψ two formulas. The label of $\phi \wedge \psi$ is exactly $MP_H(\neg MP_{\sim H}(K^\wedge \wedge \neg \phi) \wedge \neg MP_{\sim H}(K^\wedge \wedge \neg \psi))$ (except nogoods of K).*

Proof: $MP_H(\neg MP_{\sim H}(K^\wedge \wedge \neg(\phi \wedge \psi))) = MP_H(\neg MP_{\sim H}((K^\wedge \wedge \neg \phi) \vee (K^\wedge \wedge \neg \psi))) =$

$$MP_H(\neg(MP_{\sim H}(K^\wedge \wedge \neg\phi) \vee (K^\wedge \wedge \neg\psi))) = MP_H(\neg MP_{\sim H}(K^\wedge \wedge \neg\phi) \wedge \neg MP_{\sim H}(K^\wedge \wedge \neg\psi)).$$

Owing to Prop. 4, an MPL-based ATMS is able to compute the label of a phrase (a conjunction of literals) or a clause. So we can compute the label of each preference clause in a simple way. Let us point out the fact that *only the function MPL()* is used to compute labels and nogoods. Difficult operations like subsumption are not explicitly performed for these computations.

The main advantage of the MPL-technique is its ability to compute the label of a unique literal without computing the labels of the other literals as with de Kleer's technique. Moreover, an MPL-based ATMS can be applied on any set of clauses (CNF formula) and can compute in the same way the label of a literal, a disjunction or a conjunction of literals. The label computation presupposes a computation of the nogoods, in order to remove from the label the inconsistent environments. Nogoods and labels are computed in the same way, from the knowledge base (K) for nogoods, and from the knowledge base augmented with the negation of the formula, ($K^\wedge \wedge \neg\phi$), for the label of this formula.

Computation of optimal decisions via MPL

Optimistic decisions The use of MPL to solve an optimistic decision problem is easy. Assuming that K and P are CNF representations of knowledge and preference bases² of the decision problem, *good* decisions d can be obtained by a call to MPL:

Proposition 5 $K^\wedge \wedge d^\wedge \wedge P^\wedge \neq \perp$ iff $\exists E \in MPL(K \cup P, \{\}, D)$ s.t. $E \subseteq d$.

A good (optimistic) decision is then a consistent superset of an element from $MPL(K \cup P, \{\}, D)$.

Let $K^\wedge = \phi_1 \wedge \phi_2 \dots \wedge \phi_n$ and $P^\wedge = \psi_1 \wedge \psi_2 \wedge \dots \wedge \psi_m$. Finding d maximizing $U^*(d) = n(\alpha)$ such that: $K^\wedge_\alpha \wedge d^\wedge \wedge P^\wedge_\alpha \neq \perp$ (cf Definition 2) is equivalent to finding $MPL(K^\wedge_\alpha \cup P^\wedge_\alpha, \{\}, D) \neq \{\}$ minimizing α .

This method has one requirement: $K^\wedge \wedge P^\wedge$ must be a CNF formula, so P^\wedge must be a CNF formula, that is P must contain only clauses.

Pessimistic optimal decisions We propose to translate the *pessimistic* decision problem into a problem tractable by an ATMS. Let us define the set of assumption symbols $\mathcal{H} = D$. Then, assume that K is the knowledge base of the decision problem in conjunctive normal form and consider the goal base P as a formula P^\wedge . Using the symbols in \mathcal{H} , a decision d is a subset of \mathcal{H} . For any decision d such that $K^\wedge \wedge d^\wedge \vdash P^\wedge$ and $K^\wedge \wedge d^\wedge \neq \perp$ there is at least one element E of $label_K(P)$ according to the assumption set \mathcal{H} such that $E \subseteq d$.

²A stratified possibilistic knowledge base can always be put in an equivalent base of weighted clauses (Dubois et al. (1994)), since necessity measures are min-decomposable for conjunction.

Proposition 6 $K^\wedge \wedge d^\wedge \vdash P^\wedge$ and $K^\wedge \wedge d^\wedge \neq \perp$ iff $\exists E \in label_K(P)$ s.t. $E \subseteq d$.

A good (pessimistic) decision is then a superset of an element from $label_K(P)$. In the following we will only look for decisions which are minimal for set inclusion.

Let $K^\wedge = \phi_1 \wedge \phi_2 \dots \wedge \phi_n$ and $P^\wedge = \psi_1 \wedge \psi_2 \wedge \dots \wedge \psi_m$. Finding all decisions d maximizing α such that: $K^\wedge_\alpha \wedge d^\wedge \vdash P^\wedge_{n(\alpha)}$ and $K^\wedge_\alpha \wedge d^\wedge \neq \perp$ is equivalent to finding $label_{K_\alpha}(P_{n(\alpha)}) \neq \emptyset$ maximizing α .

Algorithm 1: COMPUTE_PESSIMISTIC_DECISION

Data: K the knowledge base, P the preference base, and D the set of decision symbols.

Result: Utility of the best pessimistic decisions, set of the best pessimistic decisions.

begin

```

 $\alpha \leftarrow 1$  % we consider the most certain layer ;
 $S \leftarrow \{\}$  ;
while  $S = \{\}$  and  $\alpha > 0$  do
    % we must first compute the nogoods of  $K$  ;
     $NG' \leftarrow MPL(K_\alpha, \{\}, \sim D)$  ;
     $NG \leftarrow \sim MPL(\neg NG', \{\}, D)$  ;
    % NG contains the nogoods of  $K_\alpha$  ;
     $S' \leftarrow MPL(K_\alpha \cup \sim P_1, \{\}, \sim D)$  ;
     $S \leftarrow MPL(\neg S', \sim NG, D)$  ;
    % S contains the label of  $P_1$  ;
     $\beta \leftarrow 1$  ;
    while  $\beta > n(\alpha)$  and  $S \neq \{\}$  do
         $S' \leftarrow S' \wedge MPL(K_\alpha \cup \sim P_\beta, \{\}, \sim D)$  ;
         $S \leftarrow MPL(\neg S', \sim NG, D)$  ;
        % S contains the label of  $P_\beta$  ;
         $\beta \leftarrow \beta$  ;
    if  $S = \{\}$  then  $\alpha \leftarrow \max(\alpha, n(\beta))$  ;
return  $\langle \alpha, S \rangle$  ;

```

end

where $P_{(\beta)} = P_\beta - P_{\bar{\beta}}$ is the set of goals with priority level β , α denotes the level (either of certainty or of priority) of the next non-empty layer below α .

We need to compute the nogoods, and then the required label. The restriction here is to have $K^\wedge \wedge (\sim P)^\wedge$ as a CNF formula, so P^\vee being a DNF formula. Since P is a CNF, this procedure will accept only P as a single clause or a single phrase (both are CNF and DNF form). Thus, we have to use a particularity of MPL to compute the label of a conjunction of formulas: the label of a conjunction $\psi \wedge \phi$ can be performed from the two first steps needed to compute the label of both ψ and ϕ (Prop. 4). This approach allows to stop label computation as soon as the intermediate label is empty.

One of the major advantages of our approach is that we only need to implement the MPL algorithm. An efficient implementation of MPL entails an efficient implementation of the decision algorithm. Thanks to the relation between the MPL algorithm and the Davis and Putnam algorithm, some improvements on the lat-

ter can be used in the former (heuristics for instance). The anytime aspect of the MPL algorithm can also be pointed out here. If you stop the algorithm before its normal end, you can obtain a subset of the set of optimal decisions. This can be used for instance if we only need a single optimal decision, or the utility of the optimal decision(s).

Concluding remarks

The main contribution of this paper has been to describe a logical machinery for decision-making, implementing the qualitative possibilistic utility theory, in the framework of possibilistic logic. A link between this logical machinery and the ATMS framework has been pointed out, which allowed to adapt some efficient algorithms proposed in this framework to possibilistic qualitative decision making.

Besides, in (Le Berre and Sabbadin (1997)), another logical machinery has been presented, in the diagnosis and repair framework. There, probabilities are assigned to assumptions, and numerical rewards to goals, leading to a variant of the expected-utility criterion, based on belief functions.

References

- Benferhat, S.; Dubois, D.; and Prade, H. 1997. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence* 92:259–276.
- Bonet, B., and Geffner, H. 1996. Arguing for decisions : a qualitative model of decision making. In Horwitz, E., and Jensen, F., eds., *Proc. of the 12th Conf. on Uncertainty in Artificial Intelligence (UAI'96)*, 98–105. Portland, Oregon: Morgan Kaufman.
- Boutillier, C. 1994. Toward a logic for qualitative decision theory. In Doyle, J.; Sandewall, E.; and Torasso, P., eds., *Proc. of the Fourth Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'94)*, 75–86. Bonn, Germany: Morgan Kaufman.
- Brafman, R., and Tennenholtz, M. 1997. On the axiomatization of qualitative decision criteria. In *Proc. of the 14th National Conf. on Artificial Intelligence (AAAI'97)*, 76–81. Providence, R. H.: AAAI Press / MIT Press.
- Castell, T.; Cayrol, C.; Cayrol, M.; and Le Berre, D. 1996. Efficient computation of preferred models with Davis and Putnam procedure. In *Proc. of the 12th European Conf. on Artificial Intelligence (ECAI'96)*, 354–358. Budapest: Wiley.
- Castell, T.; Cayrol, C.; Cayrol, M.; and LeBerre, D. 1998. Modèles p-restreints. applications à l'inférence propositionnelle. In *Proc. 11eme Congrès Reconnaissance des Formes et Intelligence Artificielle (RFIA'98)*, pp. 205–214. Clermont-Ferrand, France.
- Console, L., and de Kleer, J., eds. 1992. *Readings in Model-Based Diagnosis*. San Mateo CA: Morgan Kaufmann.
- Davis, H., and Putnam, L. 1960. A computing procedure for quantification theory. *J. of the Assoc. Comp. Mach.* 7:201–215.
- Davis, H.; Logemann, G.; and Loveland, D. 1962. A machine program for theorem proving. *Commun. Assoc. Comp. Mach.* 5:394–397.
- De Kleer, J. 1986. An assumption based truth maintenance system, and extending the atms. *Artificial Intelligence* 28:127–162, 163–196.
- Dubois, D., and Prade, H. 1995. Possibility theory as a basis for qualitative decision theory. In *Proc. of the 14th Int. Joint Conf. in Artif. Intelligence (IJCAI'95)*, 1925–1930. Montreal, Canada: Morgan Kaufman.
- Dubois, D.; Le Berre, D.; Prade, H.; and Sabbadin, R. 1998. Using possibilistic logic for modeling qualitative decision. *Fundamenta Informaticae*. To appear.
- Dubois, D.; Fargier, H.; and Prade, H. 1997. Decision-making under ordinal preferences and uncertainty. In *Proc. of the 13th Conf. on Uncertainty in Artificial Intelligence (UAI'97)*, 157–164. Providence R. H.: Morgan Kaufman.
- Dubois, D.; Lang, J.; and Prade, H. 1994. Automated reasoning using possibilistic logic: Semantics, belief revision, and variable certainty weights. *IEEE Trans. on Knowledge and Data Engineering* 6(1):64–69.
- Dubois, D.; Prade, H.; and Sabbadin, R. 1997a. A possibilistic logic machinery for qualitative decision. In *AAAI 1997 Spring Symposium Series (Qualitative Preferences in Deliberation and Practical Reasoning)*, 47–54.
- Dubois, D.; Prade, H.; and Sabbadin, R. 1997b. Towards axiomatic foundations for decision under qualitative uncertainty. In *Proc. of the 7th World Congress of the Inter. Fuzzy Systems Association (IFSA'97)*, 441–446. Prague, Czech Republic: Academia Verlag.
- Le Berre, D., and Sabbadin, R. 1997. Decision-theoretic diagnosis and repair: representational and computational issues. In *Proc. of the 8th International Workshop on Principles of Diagnosis (DX'97)*, 141–145. Le Mont-Saint-Michel, France: M.O. Cordier, ed.
- Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proc. of the 3rd Conf. on Theoretical Aspects of Reasoning About Knowledge (TARK'90)*, 121–135. Morgan Kaufman.
- Tan, S. W., and Pearl, J. 1994. Qualitative decision theory. In *Proc. of the 12th National Conf. on Artificial Intelligence (AAAI'94)*, 928–933. Seattle WA: AAAI Press / MIT Press.