COMBINING HYPOTHETICAL REASONING
and
PLAUSIBLE INFERENCE IN POSSIBILISTIC LOGIC*

Didier DUBOIS – Henri PRADE

Institut de Recherche en Informatique de Toulouse (I.R.I.T.)
Université Paul Sabatier, 118 route de Narbonne
31062 Toulouse Cedex, France,
E-mail: {dubois, prade}@irit.irit.fr

ABSTRACT

Possibilistic logic provides an efficient tool for handling uncertain and nonmonotonic reasoning on the one hand, and for hypothetical reasoning on the other hand. In the first situation, logic formulas are simply weighted by lower bounds of a necessity measure. Thus, formulas in a knowledge base are rank-ordered according to their levels of certainty. In the case of hypothetical reasoning, these levels of certainty are made context-dependent. The combination of these reasoning capabilities enables us to perform a case-by-case reasoning. This gives birth to a reasoning machinery where the inferential power of possibilistic logic is increased by reasoning in an hypothetical way. In the paper, the basics of possibilistic logic are first recalled, emphasizing its ability to handling partial inconsistency. Then, it is shown how to handle assumptions under the form of symbolic weights. Possibilistic logic is here shown to provide an inference tool not only for minimizing abnormality (the higher the weight of a formula, the more abnormal are the situations which contradict this formula), but also for letting some formulas vary while others are fixed, as in circumscription-based methods. This method is applied to updating problems where default rules describe the way a system is evolving with time and where the initial state of the system may be incompletely known, thus requiring a case-by-case reasoning in order to compute the following states which can be reached.

---

* A preliminary version of this paper (entitled "Combining hypothetical and nonmonotonic reasoning modes with applications to updating — A possibilistic logic approach") has been presented at the 4th IEEE International Conference on Fuzzy Systems, Yokohama, Japan, March 20-24, 1995.

1. INTRODUCTION

Possibilistic logic provides a simple and efficient tool for handling nonmonotonic reasoning (Dubois and Prade, 1991b; Dubois, Lang and Prade, 1994). Possibilistic knowledge bases are made of classical logic formulas weighted by lower bounds of possibility or necessity measures. In the following we restrict ourselves to propositional formulas associated with lower bounds of necessity measures interpreted as certainty levels. Viewing a formula as expressing a rule with possible exceptions, the higher the certainty level of the formula, the more exceptional the possible exceptions to the rule, or in other words, the more abnormal it is to encounter exceptions.

When the knowledge base is consistent in the classical sense (i.e., without considering the weights of the formulas), possibilistic inference is monotonic. Moreover, the level of certainty of a conclusion is equal to the greatest value attached to an inference path leading to this conclusion, the value of an inference path being the weight associated with the least certain formula(s) used in the path. When the knowledge base is inconsistent, a nonmonotonic behavior may take place. Indeed, a level of partial inconsistency is computed for the knowledge base as the greatest level of certainty attached to the derivation of the contradiction from the knowledge base. Conclusions obtained with levels of certainty strictly higher than the level of partial inconsistency are still valid since they are produced from a consistent subpart of the knowledge base. Thus, when new weighted formulas are added to the knowledge base, its level of inconsistency may become higher, and some previous conclusions with levels of certainty smaller than the new level of inconsistency may be invalidated, while new, possibly opposite, conclusions can be obtained from the new knowledge base.

Besides, some pieces of information might be missing in order to be able to derive a conclusion of interest. Then it is useful to determine under what (minimal) set(s) of hypotheses it is indeed possible to get this conclusion. This computation can be carried out using Assumption-based Truth Maintenance Systems (ATMS for short) as proposed by DeKleer (1986a, b). This method has been extended to possibilistic knowledge bases under the form of a so-called "possibilistic ATMS" (Π-ATMS for short); see Dubois, Lang and Prade (1991), Benferhat et al. (1994). Π-ATMS's take advantage of the fact that a possibilistic logic formula $(\neg p \vee q, \alpha)$ where $\alpha \in [0,1]$, is semantically equivalent to the possibilistic formula $(q, \min(\alpha, v(p)))$ which expresses that q is certain at degree α provided that $v(p) = 1$, i.e., that p is known to be true. Thus, it is possible to incorporate hypothetical parts of formulas into the weights, and to propagate them in the inference process.

When hypothetical reasoning and partial inconsistency take place at the same time, it may happen that the level of inconsistency of the knowledge base is a function of hypothetical variables. Then depending on the actual value of the level of inconsistency, and thus on the fact that some hypotheses hold or not, some (nonmonotonic) conclusions can be obtained or not. Then reasoning case-by-case, according to a set of mutually exclusive hypotheses, it is possible to obtain non-trivial conclusions, that, in possibilistic logic, differ from the ones that would be obtained otherwise. This is the main topic of this paper.

Section 2 gives the necessary background on possibilistic logic. Section 3 shows how the certainty ordering between two formulas (induced by their weights) can depend on a given context and can be reversed in another context. Section 4 provides a preliminary investigation of the extended inference capabilities obtained by combining possibilistic inference with hypothetical reasoning. Section 5 applies this technique to updating problems where possibilistic logic formulas describe how a system evolves over time. Concluding remarks point out lines for further research.

2. BACKGROUND ON POSSIBILISTIC LOGIC

Possibilistic logic is a particular multiple-valued logic (Rescher, 1969) that is not truth-functional. It is akin to probabilistic logic (as in Adams and Levine (1975), for instance) but not so much to what is usually called multiple-valued logic, where truth-tables for connectives are supplied (for instance Lukasiewicz logic) and that do not rely on a Boolean algebraic structure. Possibilistic logic is built on top of a Boolean algebra of propositions. The syntax of possibilistic logic contains weights attached to formulas, as in several multiple-valued systems (e.g., Pavelka, 1979), but the truth-functionality assumption is often taken for granted in the latter, while in possibilistic logic only the conjunction is truth-functional. The weights that appear in a possibilistic theory should not be interpreted as degrees of "truth" stricto sensu, but as levels of certainty of truth of the corresponding (Boolean) propositions, as in probabilistic logic. However possibilistic logic differs from probabilistic logic in at least two respects: probabilistic logic is truth-functional with respect to negation only, and this is not so for possibilistic logic; moreover, in possibilistic logic the weights need not be numbers in the unit interval but can be taken as elements in a total ordering. The syntactic part of possibilistic logic is actually very close to Rescher(1976)'s method of plausible reasoning.

2.1-Necessity-weighted formulas and their semantics

Let us first recall that a necessity measure, denoted by N, is a function from a logical (propositional or first-order) language $\mathcal{L}$ to [0,1], such that $N(\top) = 1$ and $N(\bot) = 0$, where $\top$ (resp. $\bot$) denotes any tautology (resp. any contradiction), and obeying the following axiom (Dubois and Prade, 1988):

$$\forall p, \forall q, N(p \wedge q) = \min(N(p), N(q)). \quad (1)$$

The unit interval is a matter of convention and can be changed into a bounded totally ordered set S, with top 1 and bottom 0. A necessity measure N is the dual of a possibility measure Π, such that $\forall p$, $N(p) = 1 - \Pi(\neg p)$, where Π obeys the characteristic axiom (Zadeh, 1978)

$$\forall p, \forall q, \Pi(p \vee q) = \max(\Pi(p),\Pi(q)).$$

Again, the function $1 - (\cdot)$ can be turned into an order-reversing map in S. A possibility measure Π can be expressed in terms of a possibility distribution π on the set Ω of interpretations of the language by Zadeh(1978):

$$\forall p, \Pi(p) = \max \{\pi(\omega), \omega \in M(p)\}. \quad (2)$$

This is in agreement with the characteristic axiom of possibility measures, and requires that the set Ω be finite. This is assumed here.

A possibilistic logic formula is a pair (p,α) where p is a classical first-order or propositional logic formula and α an element of S, here a number belonging to the semi-open real interval (0,1], which estimates to what extent it is certain that p is true considering the available, possibly incomplete information about the world. More formally (p,α) is a syntactic way of encoding the semantic constraint $N(p) \geq \alpha$. Notice the similarity with probabilistic logic where (p,α) can encode a semantic constraint such as $\mathrm{Prob}(p) = \alpha$, for instance.

Let M(p) be the set of models of p. The semantics of the weighted formula (p,α) is represented by the fuzzy set of models $M(p,\alpha)$ defined (e.g., Dubois, Lang and Prade, 1994) by

$$\mu_{M(p,\alpha)}(\omega) = 1 \text{ if } \omega \in M(p);$$
$$\mu_{M(p,\alpha)}(\omega) = 1 - \alpha \text{ if } \omega \notin M(p); \quad (3)$$

where ω denotes an interpretation of the language. In other words, the interpretations compatible with (p,α) are restricted by the above fuzzy set $M(p,\alpha)$ whose membership function is interpreted as a possibility distribution. The interpretations in M(p), where p is true, are considered as fully possible while the interpretations outside M(p) (where p is false) are all the more possible as α is smaller, i.e., the piece of knowledge is less certain. Any possibility distribution π satisfying the constraint $N(p) \geq \alpha$ is such that

$$\forall\omega, \pi(\omega) \leq \mu_{M(p,\alpha)}(\omega)$$

since $\Pi(\neg p) \leq 1 - \alpha \Leftrightarrow N(p) \geq \alpha$. Hence (3) is justified by the minimal specificity principle (for instance, Yager, 1992) since it yields the least specific possibility distribution $\pi = \mu_{M(p,\alpha)}$, i.e., the one with the greatest possibility degrees, compatible with (p,α).

In case of several pieces of knowledge (p_i,α_i), for i=1,n, forming a knowledge base $\mathcal{K}$, we combine them in accordance with the minimal specificity principle. Viewing a possibility distribution as an elastic restriction on possible interpretations, this principle allocates the greatest possibility degrees allowed by the constraints to each interpretation. Thus the constraints are not made more restrictive than they actually are and the information is not made more precise than what is available. Consequently, a possibility distribution $\pi_{\mathcal{K}}$ is attached to the set of possibilistic formulas $\mathcal{K}$, and built by applying the largest conjunction operation (which is also the only idempotent one) to the membership functions $\mu_{M(p_i,\alpha_i)}$, for i =1,n, namely

$$\pi_{\mathcal{K}}(\omega) = \min_{i=1,n} \mu_{M(p_i,\alpha_i)}(\omega) \tag{4}$$

It can be checked that the necessity measure $N_{\mathcal{K}}$ induced from $\pi_{\mathcal{K}}$ by

$$N_{\mathcal{K}}(q) = 1 - \sup\{\pi_{\mathcal{K}}(\omega), \omega \models \neg q\},$$

where $\omega \models q$ means $\omega \in M(q)$, is the smallest necessity measure satisfying the constraints $N(p_i) \geq \alpha_i$ for i=1,n. Thus propositions are not made more certain than they are. This is nothing but another formulation of the minimal specificity principle in terms of certainty. Generally, a strict inequality $N_{\mathcal{K}}(p_i) > \alpha_i$ may be obtained for some i. Indeed, for instance, the set of constraints $\{N(p) \geq \alpha, N(\neg p \vee q) \geq \alpha, N(q) \geq \beta\}$ with $\alpha > \beta$, entails $N_{\mathcal{K}}(q) \geq \alpha > \beta$.

Let us take an illustrative example:

$$\mathcal{K} = \{(\neg p \vee q, 1), (\neg p \vee r, 0.7), (\neg q \vee r, 0.4), (p, 0.5), (q, 0.8)\}$$

It induces the constraints,

$$\begin{array}{lcl}
\pi(\omega) \leq \mu_{M(\neg p \vee q,\ 1)}(\omega) & & \forall\omega \models p \wedge \neg q, \pi(\omega) = 0 \\
\pi(\omega) \leq \mu_{M(\neg p \vee r,\ 0.7)}(\omega) & & \forall\omega \models p \wedge \neg r, \pi(\omega) \leq 0.3 \\
\pi(\omega) \leq \mu_{M(\neg q \vee r,\ 0.4)}(\omega) & \Leftrightarrow & \forall\omega \models q \wedge \neg r, \pi(\omega) \leq 0.6 \\
\pi(\omega) \leq \mu_{M(p,\ 0.5)}(\omega) & & \forall\omega \models \neg p, \pi(\omega) \leq 0.5 \\
\pi(\omega) \leq \mu_{M(q,\ 0.8)}(\omega) & & \forall\omega \models \neg q, \pi(\omega) \leq 0.2
\end{array}$$

The corresponding possibility distribution $\pi_{\mathcal{K}}$, given by (4), is

$$\pi_{\mathcal{K}}(pqr) = 1;\ \pi_{\mathcal{K}}(pq\neg r) = 0.3;$$
$$\pi_{\mathcal{K}}(p\neg qr) = \pi_{\mathcal{K}}(p\neg q\neg r) = 0;$$
$$\pi_{\mathcal{K}}(\neg pqr) = \pi_{\mathcal{K}}(\neg pq\neg r) = 0.5;$$
$$\pi_{\mathcal{K}}(\neg p\neg qr) = \pi_{\mathcal{K}}(\neg p\neg q\neg r) = 0.2;$$

where $\omega = p\neg qr$ denotes the interpretation where p is true, q is false, r is true, and so on.

The semantic entailment of (p, α) from a possibilistic knowledge base $\mathcal{K}$ is defined by

$$\mathcal{K} \vDash (p, \alpha) \Leftrightarrow \forall\omega, \pi_{\mathcal{K}}(\omega) \leq \mu_{M(p,\ \alpha)}(\omega).$$

This agrees with Zadeh (1979)'s view of entailment in approximate reasoning. It can be checked that

$$\begin{aligned}\mathcal{K} \vDash (p, \alpha) &\Leftrightarrow \forall\omega \vDash \neg p, \pi_{\mathcal{K}}(\omega) \leq 1 - \alpha\\ &\Leftrightarrow \Pi_{\mathcal{K}}(\neg p) \leq 1 - \alpha\\ &\Leftrightarrow N_{\mathcal{K}}(p) \geq \alpha.\end{aligned}$$

In the following we omit the index $\mathcal{K}$, when writing Π and N.

2.2 - From partial inconsistency to nonmonotonic reasoning

The possibility distribution in the example above is normalized, i.e., $\exists\omega, \pi_{\mathcal{K}}(\omega) = 1$ (in the general case, there may exist several ω such that $\pi_{\mathcal{K}}(\omega) = 1$). It means that $\mathcal{K}$ is fully consistent since there is at least one interpretation in agreement with $\mathcal{K}$ which is completely possible. More generally we define the degree of inconsistency of $\mathcal{K}$ by

$$\mathrm{Inc}(\mathcal{K}) = 1 - \sup_{\omega} \pi_{\mathcal{K}}(\omega). \tag{5}$$

It can be also established that $\mathrm{Inc}(\mathcal{K}) > 0 \Leftrightarrow \mathcal{K}^*$ is inconsistent, where $\mathcal{K}^*$ is the classical knowledge base obtained from $\mathcal{K}$ by deleting the weights. In other words, the consistency of $\mathcal{K}^*$ ensures that $\pi_{\mathcal{K}}$ is normalised and conversely.

When $\mathcal{K}$ is partially inconsistent ($\mathrm{Inc}(\mathcal{K}) > 0$), the semantic entailment can be refined by considering as legitimally entailed only those propositions to which a degree of certainty, higher than the level of inconsistency, can be attached. Namely:

$$\begin{aligned}\exists\, \alpha > \mathrm{Inc}(\mathcal{K}),\ \mathcal{K} \vDash (p,\alpha) &\Leftrightarrow \forall\omega, \pi_{\mathcal{K}}(\omega) \leq \mu_{M(p,\ \alpha)}(\omega)\\ &\Leftrightarrow \Pi(\neg p) \leq 1 - \alpha < 1 - \mathrm{Inc}(\mathcal{K}) = \Pi(p)\\ &\Leftrightarrow \Pi(\neg p) < \Pi(p).\end{aligned}$$

since $1 - \mathrm{Inc}(\mathcal{K}) = \max(\Pi(p), \Pi(\neg p))$. This type of entailment can be studied independently of $\mathcal{K}$ as follows.

A possibility distribution is a way of encoding a preference ordering among interpretations, i.e., the kind of relation used by Shoham (1988) for providing non-monotonic logics with a semantics. An interpretation of p is usually called a model of p. When a possibility distribution π is available on Ω, it leads to the notion of a preferred model of a formula. A model of p is a preferred model with respect to π, denoted $\omega \models_\pi p$, if and only if $\pi(\omega) = \Pi(p) > 0$. In other terms, $\omega \models_\pi p \Leftrightarrow \omega \models p, \Pi(p) > 0$ and $\nexists \omega', \omega' \models p$ and $\pi(\omega) < \pi(\omega')$. The preferential satisfiability $\models_\pi$ can be extended to a notion of preferential entailment as shown in (Dubois and Prade, 1991) and is defined by

$$p \models_\pi q \Leftrightarrow (\exists \omega, \omega \models_\pi p \text{ and } \forall\, \omega, \omega \models_\pi p \Rightarrow \omega \models q)$$

This definition is in complete agreement with nonmonotonic consequence relations obeying the axiomatics of the so-called rational closure system proposed by Lehmann and Magidor (1992). The preferential entailment $p \models_\pi q$ means that all the preferred models of p (as induced by π), i.e., all those which maximize $\Pi(p)$ defined by (2), are models of q. Any rational nonmonotonic consequence relation $\mid\!\sim$ that satisfies the rational monotonicity axiom (if $p \mid\!\sim r$ and $\neg(p \mid\!\sim \neg q)$ then $p \wedge q \mid\!\sim r$) and a few others that ensure that nonmonotonic consequences are deductively closed, can be modelled as a possibilistic preferential entailment (Benferhat S., Dubois D. and Prade H. 1992).

Then, it has been established (Dubois and Prade, 1991) that the above inconsistency-free entailment from $\mathcal{K}$ can be expressed as follows:

$$p \models_\pi q \Leftrightarrow \exists\, \alpha > \mathrm{Inc}(\mathcal{K} \cup \{(p, 1)\}),\ \mathcal{K} \cup \{(p, 1)\} \models (q, \alpha)$$

It can be also shown that the preferential entailment is closely related to the notion of conditional possibility (Dubois and Prade, 1988). Let π be a possibility distribution on Ω. the conditional possibility distribution $\pi(\cdot \mid p)$ is generated by Hisdal (1978)'s equation

$$\Pi(p \wedge q) = \min(\Pi(q \mid p), \Pi(p))$$

where the conditional possibility $\Pi(q \mid p)$ is defined, when $\Pi(p) > 0$, as the greatest (least specific) solution to this equation (Dubois and Prade, 1988), namely

$$\begin{aligned}\Pi(q \mid p) &= 1 \text{ if } \Pi(p \wedge q) = \Pi(p)\\ &= \Pi(p \wedge q) \text{ otherwise.}\end{aligned}$$

If $\Pi(p) = 0$, $\Pi(q|p) = 1$, $\forall\, q \neq \bot$. The conditional possibility distribution is then for $\omega \models p$ ($\pi(\omega|p) = 0$ if $\omega \models \neg p$):

$$\begin{aligned}\pi(\omega \mid p) &= 1 \text{ if } \pi(\omega) = \Pi(p),\\ &= \pi(\omega) \text{ otherwise.}\end{aligned}$$

The conditional necessity is $N(q \mid p) = 1 - \Pi(\neg q \mid p)$. This is a purely ordinal notion of conditioning, similar to the one of probability theory. The equivalences hold:

$$\begin{aligned}p \models_{\pi} q &\Leftrightarrow \Pi(q \mid p) > \Pi(\neg q \mid p)\\ &\Leftrightarrow N(q \mid p) > 0\\ &\Leftrightarrow \Pi(p \wedge q) > \Pi(p \wedge \neg q)\end{aligned}$$

which agrees with our intuition that p entails q plausibly, if given the information $\mathcal{K}$, $p \wedge q$ is found to be strictly more possible than $p \wedge \neg q$.

2.3- Possibilistic resolution principle

In this section we suppose that weighted formulas are put under the form of weighted *clauses*; this can be done without loss of expressivity, due to the conjunctive compositionality of necessity measures. The following resolution rule (Dubois and Prade, 1987)

$$(p \vee q, \alpha)\ ;\ (p \vee r, \beta) \vdash (q \vee r, \min(\alpha, \beta)) \qquad (6)$$

where $\vdash$ still denotes the syntactic inference in possibilistic logic, has been proved sound and complete for refutation with respect to the semantics recalled in sections 2.1-2.2, (Dubois et al., 1994).

In order to compute the maximal certainty degree which can be attached to a formula according to the constraints expressed by a knowledge base $\mathcal{K}$, for instance r in the above example of section 2.1, just add to $\mathcal{K}$ the clause(s) obtained by refuting the proposition to evaluate, with a necessity degree equal to 1; here add $(\neg r, 1)$. Then it can be shown that any lower bound obtained on $\bot$, by resolution, is a lower bound of the necessity of the proposition to evaluate. See (Dubois et al., 1987) for an ordered search method which guarantees that we obtain the greatest derivable lower bound on $\bot$. It can be shown (Dubois, Lang, and Prade, 1994; Lang et al., 1991), that this greatest derivable lower bound on $\bot$ is nothing but the inconsistency degree $\mathrm{Inc}(\mathcal{K} \cup \{(\neg r, 1)\})$ where r is the proposition to establish. In the example, we have the following derivation

$(\neg p \vee r, 0.7)$ $(\neg r, 1)$

$(\neg p, 0.7)$ $(p, 0.5)$

$(\bot, 0.5)$

i.e., $N(r) \geq 0.5$ and indeed it can be checked that using the possibility distribution of section 2.1, we have $\Pi(\neg r) = \sup\{\pi(\omega), \omega \vDash \neg r\} = 0.5$, in fact $\Pi(\neg r) \leq 0.5$ since π is the greatest possibility distribution compatible with $\mathcal{K}$, and then $N(r) = 1 - \Pi(\neg r) \geq 0.5$.

In case of partial inconsistency of $\mathcal{K}$, a refutation carried out in a situation where $\mathrm{Inc}(\mathcal{K} \cup \{(\neg r, 1)\}) = \alpha > \mathrm{Inc}(\mathcal{K})$ yields the non-trivial conclusion (r, α), only using formulas whose degree of certainty is strictly greater than the level of inconsistency of the base (and particularly, at least equal to α).

It should be pointed out that the two following particular inference rules are always valid:

. if $p \vdash q$ (in the classical sense) then $(p,\alpha) \vdash (q,\alpha)$ (7)

for any α due to the monotonicity of necessity measures, since $N(q) \geq N(p) \geq \alpha$,

. if $\beta \leq \alpha$ then $(p,\alpha) \vdash (p,\beta)$ (8)

due to $N(p) \geq \alpha \geq \beta$. These two rules express subsumption in possibilistic logic. Thus, adding (q,β) to a knowledge base which already includes (p,α) such that $p \vdash q$ and $\beta \leq \alpha$ has no effect. Moreover formulas of the form $(p,0)$ do not appear in practice since the level 0 encodes the trivial fact that $N(p) \geq 0$. Moreover since $N(p) \geq \alpha$ and $N(p) \geq \beta$ entail $N(p) \geq \max(\alpha, \beta)$, we have the following noticeable inference rule.

. $(p,\alpha); (p,\beta) \vdash (p,\max(\alpha,\beta))$. (9)

2.4 - Formulas labelled by contexts

As pointed out in (Dubois et al., 1991, 1994), the weighted clause $(\neg p \vee q, \alpha)$ is semantically equivalent to the weighted clause $(q, \min(\alpha, v(p)))$ where $v(p)$ is the truth value of p, i.e., $v(p) = 1$ if p is true and $v(p) = 0$ if p is false. Indeed, for any uncertain proposition (p,α) we can write $\mu_{M(p,\alpha)}(\omega)$ under the form $\max(v_\omega(p), 1 - \alpha)$, where $v_\omega(p)$ is the truth-value assigned to p by interpretation ω. Then obviously:

$$\begin{aligned}
\forall\omega, \mu_{M(\neg p \vee q, \alpha)}(\omega) &= \max(v_\omega(\neg p \vee q), 1 - \alpha) \\
&= \max(1 - v_\omega(p), v_\omega(q), 1 - \alpha) \\
&= \max(v_\omega(q), 1 - \min(v_\omega(p), \alpha)) \\
&= \mu_{M(q,\ \min(v_\omega(p),\ \alpha))}(\omega).
\end{aligned}$$

This remark is very useful for hypothetical reasoning, since by "transferring" a component of a disjunction to the weight part of the formula, assumptions can be explicitly distinguished. It is a way of considering a propositional variable p as a typed one. Indeed changing $(\neg p \vee q, \alpha)$ into $(q, \min(v(p), \alpha))$ leads to state the piece of knowledge under the form "q is certain at the degree α, *provided that* p is true". Then the weight is no more just a degree but in fact a label which expresses a context in which the piece of knowledge is more or less certain.

The semantic equivalence between $(\neg p \vee q, \alpha)$ and $(q, \min(v(p), \alpha))$ can be also checked at the syntactic level, observing that we can always add the formula $(p,v(p))$ to a possibilistic knowledge base. $(p,v(p))$ indeed expresses that p is certain if p is true. Then, using the inference rules:

i) $(\neg p \vee q, \alpha)$; $(p,v(p)) \vdash (q, \min(v(p), \alpha))$ by (6);
ii) conversely, $(q, \min(v(p), \alpha)) \vdash (\neg p \vee q, \min(v(p), \alpha))$ by (7), while $(\neg p,v(\neg p)) \vdash (\neg p \vee q, v(\neg p))$ also by (7). Then using (9) we get $(\neg p \vee q, \max(\min(v(p), \alpha), v(\neg p)))$, which entails $(\neg p \vee q, \alpha)$ by (8), since $\max(\min(v(p), \alpha), v(\neg p)) \geq \min(\alpha,1) = \alpha$ (letting $v(p) = 1$ and then $v(\neg p) = 1$).

3. MAKING THE ORDERING CONTEXT-DEPENDENT

Assessing the respective levels of certainty of formulas in practice is not always an easy matter. However, some ordering between the certainty levels should be satisfied if the knowledge base is to behave in a proper way.

Indeed, consider the simple example of a knowledge base $\mathcal{K} = \{(\neg p \vee q, \alpha), (\neg p \vee \neg r \vee \neg q, \beta),(p, 1),(r,1)\}$ with $\beta < \alpha$. Its degree of inconsistency is $\mathrm{Inc}(\mathcal{K}) = \beta$. Thus the conclusion (q, α) is valid since $\alpha > \mathrm{Inc}(\mathcal{K})$, and the rule 'if p and r then not q' (encoded by $(\neg p \vee \neg r \vee \neg q, \beta)$) plays no role, although it is a more specific rule than the rules 'if p then q' (encoded by $(\neg p \vee q, \alpha)$). Thus, if the more specific rule is to play any role and to apply when both p and r are known to be true, it requires that $\beta > \alpha$ be enforced. Intuitively, a more specific rule should have less exceptions than a general one (in the sense that it is more abnormal to encounter an exception to a specific rule than an exception to a more general one). For instance penguins which fly are more exceptional than birds which do not fly (for p = 'bird'; q = 'fly'; r = 'penguin'). A procedure is described in (Benferhat et al., 1992), that determines the constraints induced on the ordering of the certainty levels by the presence of rules which are more specific than others.

However, when specificity induces no ordering between two formulas, any ordering between them does encode the supplementary information that one of the formulas is more certain than the other. Such a statement may be context-dependent. Let us show on an example how context-dependent ordering can be handled in possibilistic reasoning, and more specifically how the choice of the context can affect the certainty ordering of conclusions.

Consider two formulas p and q, and a context expressed by a proposition c. Let α_1 and α_2 be two certainty levels such that $\alpha_1 > \alpha_2$; then the set of formulas

$$\{(p \vee \neg c, \alpha_1),$$
$$(p \vee c, \alpha_2),$$
$$(q \vee c, \alpha_1),$$
$$(q \vee \neg c, \alpha_2)\}$$

expresses that if c is true, the proposition p has a level of certainty α_1 greater than the level of certainty α_2 of proposition q. The converse holds if c is false. Indeed the above set of four formulas is equivalent to the following set of two formulas, using (9) and moving c to the weight slot:

$$(p,\ \max(\min(\alpha_1,v(c)),\ \min(\alpha_2,v(\neg c))));$$
$$(q,\ \max(\min(\alpha_1,v(\neg c)),\ \min(\alpha_2,v(c)))).$$

It can be checked on these formulas that if $v(c) = 1$ (and thus $v(\neg c) = 0$) the weight attached to p is α_1 and the one attached to q is $\alpha_2 < \alpha_1$. It is the converse if $v(c) = 0$. Thus the certainty ordering between conclusion p and q is depending on context c.

4. COMBINING NONMONOTONIC AND HYPOTHETICAL REASONING

Another example will exhibit the basic machinery where nonmonotonic reasoning and hypothetical reasoning are jointly used. Consider the following possibilistic knowledge base

$$\mathcal{K} = \{(\neg h \vee p, \alpha_1), (\neg q, \alpha_2),(\neg p \vee q,\ 1)\}.$$

It expresses that p entails q, that $\neg q$ is certain to the degree α_2 and that under hypothesis h, p is certain to the degree α_1.

Let us move the hypothetical proposition h to the weight slot, i.e.,

$$\mathcal{K}' = \{(p, \min(\alpha_1, v(h))), (\neg q, \alpha_2), (\neg p \vee q, 1)\}.$$

The degree of inconsistency of $\mathcal{K}'$ is

$$\mathrm{Inc}(\mathcal{K}') = \min(\alpha_1, \alpha_2, v(h)),$$

and now depends on v(h). Let us reason case by case about h.

– First assume $v(h) = 0$. Then the knowledge base is consistent, i.e., $\mathrm{Inc}(\mathcal{K}') = 0$ and we can infer $(\neg p, \alpha_2)$ and obviously $(\neg q, \alpha_2)$, assuming $\alpha_2 > 0$.
– If $v(h) = 1$, $\mathrm{Inc}(\mathcal{K}') = \min(\alpha_1, \alpha_2)$. Further assume $\alpha_2 < \alpha_1$, then $\mathrm{Inc}(\mathcal{K}') = \alpha_2$; then (q, α_1) can be non-trivially deduced and obviously (p, α_1) as well, since $\mathrm{Inc}(\mathcal{K}') < \alpha_1$. More formally, the following formulas have been derived:

$$(p \wedge q, \min(\alpha_1, v(h))) \text{ and } (\neg p \wedge \neg q, \min(\alpha_2, v(\neg h))).$$

Indeed if $N(r) \geq \alpha$ and $N(s) \geq \alpha$, then $N(r \wedge s) = \min(N(r),N(s)) \geq \alpha$.
On the whole, it follows that

$$\begin{aligned} N(p \Leftrightarrow q) = N((p \wedge q) \vee (\neg p \wedge \neg q)) &\geq \max(N(p \wedge q), N(\neg p \wedge \neg q)) \\ &\geq \max(\min(\alpha_1, v(h)), \min(\alpha_2, v(\neg h))) \\ &\geq \min(\alpha_1, \alpha_2) \end{aligned}$$

i.e., $(p \Leftrightarrow q, \min(\alpha_1, \alpha_2))$ can be derived.

Note that the conclusion concerning $p \wedge q$ alone $(p \wedge q, \min(\alpha_1, v(h)))$, which becomes non-trivial if $v(h) = 1$ is assumed, becomes *out of reach if we do not reason hypothetically* although $((p \wedge q) \vee \neg h, \alpha_1)$ can be derived. Indeed with $\alpha_2 < \alpha_1$

$$\mathcal{K} = \{(\neg h \vee p, \alpha_1), (\neg q, \alpha_2), (\neg p \vee q, 1)\} \vdash \{(\neg p \wedge \neg q, \alpha_2), (\neg h, \alpha_2)\}.$$

i.e., $\mathcal{K}$ is consistent but it is inferred that the hypothesis h is plausibly false. Observe that if the weights are omitted, we could conclude $\neg h$ for sure in classical logic. However if we force $v(h) = 1$ by adding $(h, 1)$ to $\mathcal{K}$, we shall have a nonmonotonic behaviour and we do obtain $(p \wedge q, \alpha_1)$ as a valid conclusion since $\alpha_1 > \alpha_2 = \mathrm{Inc}(\mathcal{K} \cup \{(h, 1)\})$. But the conclusion $(p \Leftrightarrow q, \min(\alpha_1, \alpha_2))$ would not be singled out and would remain out of reach since inhibited by the level of inconsistency $\min(\alpha_1, \alpha_2) = \alpha_2$.

More generally, we may have several hypothetical propositional variables. They may correspond to variables h_i for which it is not possible, from the knowledge base $\mathcal{K}$, to conclude either (h_i, α_i) with $\alpha_i > 0$ or $(\neg h_i, \alpha'_i)$ with $\alpha'_i > 0$. But other propositional variables can be made hypothetical as well, since, in the above example, h is used as an hypothetical variable although $\mathcal{K} \vdash (\neg h, \min(\alpha_1, \alpha_2))$ non-trivially. However, it would not be possible to reason hypothetically on h if it were the case that $\mathcal{K} \vdash (\neg h, 1)$ or $\mathcal{K} \vdash (h, 1)$.

A general procedure for combined hypothetical and plausible inference first computes the level of inconsistency $\mathrm{Inc}(\mathcal{K})$ of $\mathcal{K}$ as a symbolic expression, where the hypothetical variables appear. Then reasoning proceeds case by case by fixing the context in each case, assigning the truth values 'true' or 'false' to each hypothetical variable. The next section shows how this procedure may be applied to updating problems.

In classical logic, reasoning case-by-case in a consistent way with respect to the knowledge base, and then performing the disjunction of the conclusions obtained in each case, provides the same results as the ones which would be inferred from the knowledge base after adding to it the disjunction of the considered cases. This behaviour is no longer observed in the possibilistic framework. This is due to the following fact.

In each particular hypothetical situation, the *plausible* (that is, the non-trivial) conclusions which can be derived depend on the level of inconsistency. This level varies from one case to another, inhibiting the formulas which are below this level; these formulas may not be inhibited for another value of this level of inconsistency in another case. In order words, the useful part of the knowledge base (above the level of inconsistency) used for inferring plausible conclusions, is not exactly the same from one case to another. As recalled in the preceding sections, possibilistic logic rank-orders the possible worlds (or, if we prefer, the interpretations), on the basis of the weights attached to the formulas in the knowledge base. By moving hypothetical literals into the weights and then by reasoning case by case, different orderings between interpretations are used in each case, representing a context. Thus, the possibility distribution and the associated necessity measure which represent the semantics of the knowledge base are made context-dependent.

It should be clear that by reasoning hypothetically as described above, we are adding a particular piece of information to the knowledge, base, which is not expressed as a logical formula. Namely, we are enforcing the fact we *do not know* if *h* is either true or false in the above example, and more generally we are enforcing a situation where the different considered cases are a priori equally plausible (even if the knowledge base would conclude that some cases are more plausible than others without reasoning hypothetically). This use of hypothetical reasoning has not here the same intended purpose as in Assumption-based Truth Maintenance Systems whose aim is to determine under what sets of hypotheses a particular conclusion can be inferred.

5. APPLICATION TO UPDATING

Updating refers to information maintenance when the world changes. See Léa Sombé (1994) for a general introduction. The way a system (called "the world" here) evolves or may evolve with time can be described by logical formulas parametrized by time. This is usual in situation calculus, for instance. However, some of these pieces of knowledge may be default rules yielding plausible but defeasible conclusions. Moreover at a given time instant, the knowledge base only provides an incomplete description of the world of interest usually. This incomplete information leads us to reason case by case since the resulting state of the world depends on the possible transitions from an initial state which may be incompletely known.

Consider an example taken from Sandewall (1993) and also discussed by Cordier and Siegel (1993).

The deaf person example:

The following pieces of information are given (the initial time point is denoted t and the next (final) time point is denoted t'):

. if somebody is not deaf at time t($\neg d_t$), (s)he is still not deaf at time t';
. if somebody is asleep at time t(a_t), (s)he is still asleep at time t';
. somebody is asleep at time t;
. if a clock rings at time t'($c_{t'}$) and if the person is not deaf then (s)he is instantaneously awake ($\neg a_{t'}$);
. a clock rings at time t'.

This can be encoded in the following way in possibilistic logic:

$$\mathcal{K} = \{(d_t \vee \neg d_{t'}, \alpha_1), (\neg a_t \vee a_{t'}, \alpha_2), (a_t, 1), (\neg c_{t'} \vee d_{t'} \vee \neg a_{t'}, 1), (c_{t'}, 1)\}.$$

We further assume $1 > \alpha_1 > \alpha_2 > 0$ in order to acknowledge the fact that the persistence of 'not deaf' in the first default rule is greater than the one of 'asleep' in the second default rule.

Note that the set of formulas is consistent and that, in classical logic (dropping the weights) the only possible conclusion is d_t, which is hard to believe since it comes down to suggesting that learning that the clock rings at t' implies that the person was deaf at time t. In plain possibilistic logic, d_t is obtained with weight α_2.

Applying the procedure described in the preceding section and selecting d_t as an hypothetical variable, we rewrite the first clause as $(\neg d_{t'}, \min(\alpha_1, v(\neg d_t)))$ thus modifying $\mathcal{K}$ into

$$\mathcal{K}' = \{(\neg d_{t'}, \min(\alpha_1, v(\neg d_t))), (\neg a_t \vee a_{t'}, \alpha_2), (a_t, 1), (\neg c_{t'} \vee d_{t'} \vee \neg a_{t'}, 1), (c_{t'}, 1)\}.$$

Then

$$\mathrm{Inc}(\mathcal{K}') = \min(\alpha_1, \alpha_2, v(\neg d_t)) = \min(\alpha_2, v(\neg d_t)).$$

Thus, if $v(\neg d_t) = 1$, we get $\mathcal{K}' \vdash (\neg d_{t'}, \alpha_1)$ and $\mathcal{K}' \vdash (\neg a_{t'}, \alpha_1)$ since $\alpha_1 > \mathrm{Inc}(\mathcal{K}') = \alpha_2$.

If $v(\neg d_t) = 0$, $\mathrm{Inc}(\mathcal{K}') = 0$ and we get $\mathcal{K}' \vdash (d_{t'}, \alpha_2)$ and $\mathcal{K}' \vdash (a_{t'}, \alpha_2)$.

Notice that this example is formally equivalent to the one of Section 4. Please also note that the undesirable conclusion (d_t, α_2) is no longer systematically obtained. That is, we obtain that either the person is deaf and asleep at time t', or not deaf and not asleep, as intuitively expected.

N.B. Clearly the rule "if a clock rings at time t' and if the person is not deaf then (s)he is immediately no longer asleep" is true at any time point, and the knowledge base should also contain the clause

$$(\neg c_t \vee d_t \vee \neg a_t, 1).$$

If $\mathcal{K}$ stipulates that the clock does not ring at time t (in order to remain in the spirit of the scenario of the above example), i.e., it contains $(\neg c_t, 1)$, the previous results are still obtained. In the above example, nothing is said about the clock at time t, which means that we do not know if c_t is true or false. We may also reason hypothetically about c_t. In case $v(c_t) = 1$ it is immediate that we conclude $(d_t, 1)$; in case $v(\neg c_t) = 1$, nothing more is inferred than previously. It can be checked that reasoning case by case on the basis of the four contexts induced by c_t and d_t yields the same conclusion as above. However, note that the context $v(c_t)=1$ and $v(\neg d_t)=1$ is forbidden, since $\{(\neg c_t \vee d_t \vee \neg a_t, 1), (a_t, 1)\} \vdash (\neg c_t \vee d_t, 1)$ and that we cannot reason hypothetically on a situation for which the knowledge base $\mathcal{K}'$ entails with total certainty that it cannot occur (provided that $\mathrm{Inc}(\mathcal{K}') < 1$).

Besides, for the sake of agreement with reality, we may like to add that being deaf is a strongly persistent state, i.e.,$(\neg d_t \vee d_{t'}, \alpha_0)$ with $\alpha_0 > \alpha_1$, this clause being added to $\mathcal{K}$. Again it changes nothing to the results, which remain as expected.

Default rules may not only express persistence over time, but also other generic rules with exceptions which do not refer to time. Let us suppose that in our example we have the generic knowledge that people are usually not deaf, encoded by

$$(\neg d_t, \beta) \text{ for any t or t' and added to } \mathcal{K}.$$

Then $\mathcal{K}$ is partially inconsistent and $\mathrm{Inc}(\mathcal{K}) = \min(\alpha_1, \alpha_2, \beta)$. Without performing an hypothetical reasoning, and still assuming $\alpha_1 > \alpha_2$

$\mathcal{K} \vdash (\neg d_{t'}, \beta)$ provided that $\beta > \alpha_2$;
$\mathcal{K} \vdash (\neg a_{t'}, \min(\alpha_1, \beta))$ provided that $\beta > \alpha_2$;

if $\beta < \alpha_2$, $\mathrm{Inc}(\mathcal{K}) = \beta$ and $\mathcal{K} \vdash (d_{t'}, \alpha_2)$ and $\mathcal{K} \vdash (a_{t'}, \alpha_2)$.

However, hypothetical reasoning can be again performed by using $(\neg d_t, \max(\beta, v(\neg d_t)))$ instead of $(\neg d_t, \beta)$ and $(\neg d_{t'}, \min(\alpha_1, v(\neg d_t)))$ instead of $(d_t \vee \neg d_{t'}, \alpha_1)$. Then if the generic rule is not certain enough, i.e., $\beta < \alpha_2$, the same hypothetical reasoning as before still takes place with

$$\mathrm{Inc}(\mathcal{K}') = \min(\alpha_2, \max(\beta, \min(\alpha_1, v(\neg d_t)))).$$

If β is larger than α_2, we shall have $\mathrm{Inc}(\mathcal{K}') = \alpha_2$, and $(d_{t'}, \alpha_2)$ and $(a_{t'}, \alpha_2)$ will no longer be valid conclusions in the case where $v(\neg d_t) = 0$.

This example raises the problem of the commensurability of the two scales used for ranking persistence levels and for ranking generic defaults having different levels of specificity.

6. CONCLUDING REMARKS

This preliminary study shows the potential interest of combining nonmonotonic and hypothetical reasoning in the possibilistic logic framework. The problem of the choice of hypothetical variable(s) leading to conclusions of interest has not been discussed here. However in the update example it is clear that hypothetical variables should be the ones whose truth values are not known with full certainty in the initial state. The basic idea is that hypothetical reasoning can introduce inconsistency which enables us to block undesirable conclusions, and is thus more cautious.

The inference machinery discussed here presents some similarities with the McCarthy (1986) prioritized circumscription-based approach to updating proposed by Cordier and Siegel (1993) where predicates referring to the initial instant t are not allowed to "vary". It means in practice that the minimization of the abnormality is made without making any further assumption about the state of facts at time t. It is what is also assumed in our approach by considering all the different states of facts at time t compatible with the available knowledge. Our hypothetical reasoning approach agrees with the idea that the predicates corresponding to the hypothetical variables, are not allowed to vary in the circumscription approach. Indeed by fixing the truth values of these variables we make sure in an obvious manner that in the underlying preferential ordering we are only ranking interpretations which behave in the same way with respect to the extensions of these predicates (here the truth values of the hypothetical variables, since we are in a propositional case).

Further research should concentrate on the relation between this inference machinery and the possibilistic approach to update problems (Dupin de Saint-Cyr et al. 1994), and on the links with circumscription (see Benferhat et al. (1993) for a preliminary study).

ACKNOWLEDGEMENTS

The authors wish to thank Marie-Odile Cordier for many stimulating discussions about updating.

This work is partially supported by the French Inter-PRC project "Gestion de l'Evolutif et de l'Incertain" (Handling of Uncertainty in a Dynamical World").

References

Adams E.W., Levine H.P. (1975) On the uncertainties transmitted from premises to conclusions in deductive inferences. Synthese 30, 429-460

Benferhat S., Dubois D., Lang J., Prade H. (1994) Hypothetical reasoning in possibilistic logic: Basic notions, applications and implementation issues. In: Between Mind and Computer, Fuzzy Science and Engineering, (Wang P.Z., Loe K.F., eds), World Scientific, 1-29

Benferhat S., Dubois D., Prade H. (1992) Representing default rules in possibilistic logic. Proc. of the 3rd Inter. Conf. on Principles of Knowledge Representation and Reasoning, Cambridge, Ma., 673-684

Benferhat S., Dubois D., Prade H. (1993) Possibilistic logic: From nonmonotonicity to logic programming. Proc. of the Europ. Conf. ECSQARU'93 (Clarke M., Kruse R. and Moral S., eds), Lecture Notes in Computer Science, Vol. 747, Springer-Verlag, 17-24

Cordier M.O., Siegel P. (1993) Prioritized transitions for udpates. Proc. of the AI'93 Workshop on Belief Revision, Melbourne.

De Kleer J. (1986a) An asssumption-based TMS. Artificial Intelligence, 28 127-162

De Kleer J. (1986b) Extending the ATMS. Artificial Intelligence, 28, 163-196
Dubois D., Dupin de Saint Cyr F., Prade H. (1994) Updating, transition constraints and possibilistic markov chains. Proc. of the 4th Conf. on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Paris, July 4-8, 826-831
Dubois D., Lang J., Prade H. (1987) Theorem-proving under uncertainty — A possibilistic theory-based approach. Proc. of the 10th Inter. Joint Conf. on Artificial Intelligence, Milano, Italy, 984-986
Dubois D., Lang J., Prade H. (1991) A possibilistic assumption-based truth maintenance system with uncertain justifications, and its application to belief revision. Proc. of the ECAI Workshop on Truth-Maintenance Systems (J.P. Martins, M. Reinfrank, eds.), Stockholm, Aug. 6, 1990, Lecture Notes in Computer Sciences, Vol. 515, Springer Verlag, 87-106
Dubois D., Lang J., Prade H. (1994) Automated reasoning using possibilistic logic: semantics, belief revision, variable certainty weights. IEEE Trans. on Data and Knowledge Engineering, 6(1), 64-71
Dubois D., Prade H. (1987) Necessity measures and the resolution principle. IEEE Trans. on Systems, Man and Cybernetics, 17, 474-478
Dubois D., Prade H. (1988) Possibility Theory. Plenum Press, New-York.
Dubois D., Prade H. (1991) Possibilistic logic, preference models, non-monotonicity and related issues. Proc. of the 12th Inter. Joint Conf. on Artificial Intelligence, Sydney, Australia, Aug. 24-30, 419-424
Hisdal E. (1978) Conditional possibilities: Independence and non-interactivity. Fuzzy Sets and Systems, 1, 283-297.
Lang J., Dubois D., Prade H. (1990) A logic of graded possibility and certainty coping with partial inconsistency. Proc. of the 7th Conf. on Uncertainty in A.I., UCLA, Los Angeles, July 13-15, Morgan & Kaufmann, 188-196
Léa Sombé (1994) A glance at revision and updating in knowledge bases. Int. J. of Intelligent Systems, 9(1), 1-27
Lehmann D., Magidor M. (1992) What does a conditional knowledge base entail?. Artificial Intelligence, 55(1), 1-60.
McCarthy J. (1986) Applications of circumscription to formalizing commonsense knowledge. Artificial Intelligence, 28, 89-116.
Pavelka, J. (1979) On fuzzy logic. Zeitschr. f. Math. Logik und grundladen d. Math. 25, Part I: 45-72
Rescher N. (1969) Many-valued logics. MacGraw Hill, New York
Rescher N. (1976) Plausible Reasoning. Van Gorcum, Assen/Amsterdam.
Sandewall E. (1993) The range of applicability of nonmonotonic logics for the inertia problem. Proc. of the Inter. Joint Conf. on Artificial Intelligence, Chambéry, France, 738-743.
Shoham Y. (1988) Reasoning About Change — Time and Causation from the Standpoint of Artificial Intelligence. The MIT Press, Cambridge, MA
Yager R. (1992) On the specificity of a possibility distribution. Fuzzy Sets and Systems, 50, 279-292

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

Zadeh L.A. (1979) A theory of approximate reasoning. In: Machine Intelligence, Vol. 9 (J.E. Hayes, D. Michie, L.I. Mikulich, eds.), Elsevier, New York, 149-194