Toward multiple-agent extensions of possibilistic logic

Didier Dubois, Henri Prade

► To cite this version:

Didier Dubois, Henri Prade. Toward multiple-agent extensions of possibilistic logic. IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2007), Jul 2007, London, United Kingdom. pp.187–192, 10.1109/FUZZY.2007.4295363 . hal-03358861

HAL Id: hal-03358861

<https://hal.science/hal-03358861>

Submitted on 29 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Toward multiple-agent extensions of possibilistic logic

Didier Dubois, Henri Prade

Abstract— Possibilistic logic is essentially a formalism for handling qualitative uncertainty with an inference machinery that remains close to the one of classical logic. It is capable of handling graded modal information under the form of certainty levels attached to classical logic formulas. Such lower bounds of necessity measures are associated to the corresponding pieces of belief. This paper proposes extensions of the possibilistic logic calculus where such weighted formulas can be attached to a set of agents or which can be embedded inside another weighted formula, for the expression of mutual beliefs. It is possible to express that all the agents in a subset have some beliefs, or that there is at least one agent in a subset that has a particular belief. The case of all-or-nothing beliefs is first dealt with before presenting the inference rules for handling graded beliefs held by multiple agents. Illustrative examples are provided. The proposed framework offers a reasonable compromise between expressive power and a computational cost close to the one of classical logic.

1 INTRODUCTION

Possibilistic logic (Dubois and Prade, 2004; Dubois, Lang and Prade, 1994) offers a convenient tool for handling uncertain or prioritized pieces of information. Standard possibilistic logic expressions are propositional logic formulas associated with weights. Logical formulas with a weight strictly greater than the inconsistency level in a possibilistic logic base are immune to inconsistency and can be safely used in deductive reasoning. This property is at the basis of the representation of non-monotonic consequence relations in possibilistic logic (Benferhat et al., 1997).

A possibilistic logic formula is a pair made of i) any well-formed classical logic formula, and ii) a weight expressing its certainty or priority. Such classical logic formulas can be only true or false, and fuzzy statements with intermediary degrees of truth are not allowed in standard possibilistic logic. Then, a possibilistic logic base, viewed as a belief set, is equivalent to a conjunction of such pairs expressing graded pieces of belief (*graded beliefs*, for short). Possibilistic logic can be viewed as a particular labelled deductive system (Gabbay, 1996). In a possibilistic logic formula, the weight is formally interpreted in the framework of possibility theory (Zadeh, 1978) as a lower bound of the value of the necessity measure of the proposition appearing in the considered possibilistic logic formula. Thus, a possibilistic logic formula is similar to a modal logic formula with a graded modality (Farinas del Cerro and Herzig, 1991). Possibilistic logic expressive power is comparable to a fragment of modal logic that would use the box symbol $\Box$ only, with a computational cost close to the one of classical logic (Lang, 2001).

The paper investigates some extensions of possibilistic logic for reasoning about beliefs entertained by sets of agents, focusing as well on trust and mutual beliefs, issues that usually require the use of modal logics (e.g., Cohen and Levesque, 1990; Herzig and Longin, 2004; Lorini *et al.*, 2005). This attempt takes advantage of the fact that a possibilistic logic formula can be extended to a pair made of a statement that is either true or false associated with a kind of modality stating the circumstances under which the statement can be regarded as true. For example, the statement may represent a piece of information such as “if the agents in group A are rather certain that p is true, then they will have q as an important goal”, and this will be held as “quite certainly true by a group of agent B ”.

The paper is organized as follows. Section 2 restates the necessary background on standard possibilistic logic. Section 3 first describes an extension of propositional logic modeling the beliefs of sets of agents, before handling the general case of a graded collective possibilistic logic. In these extensions, one is able to express that all agents in some subsets, or that some agents in a subset have particular beliefs, and to reason both on beliefs and on subsets of agents. Section 4 explains how a possibilistic logic formula can be viewed as a Boolean proposition evaluated on information states. Negation or disjunction of graded beliefs can then be defined. Section 5 handle nested graded beliefs. Section 6 discusses how to interpret nested possibilistic multiagent formulas and provides an illustrative example.

2 BACKGROUND ON POSSIBILISTIC LOGIC

A standard possibilistic logic expression is a pair (p, α) , where p is a propositional (or first-order) formula and $\alpha \in (0,1]$ is interpreted as a lower bound of a necessity measure N . A graded belief (p, α) is semantically interpreted as $N(p) \geq \alpha$, where N is a necessity measure. Any discrete linearly ordered scale can be used in place of $[0,1]$. Formulas of the form $(p, 0)$, which do not contain any information ($\forall p, N(p) \geq 0$ always holds), are not part of the possibilistic language.

The characteristic property of a necessity measure is its decomposability with respect to conjunction namely, the property $N(p \wedge q) = \min(N(p), N(q))$. Important valid syntactic inference rules are:

- $(\neg p \vee q, \alpha); (p, \beta) \vdash (q, \min(\alpha, \beta))$ (modus ponens)
- for $\beta \leq \alpha$ $(p, \alpha) \vdash (p, \beta)$ (weight weakening),

where $\vdash$ denotes the syntactic inference of possibilistic logic. The min-decomposability of necessity measures allows us to work with weighted clauses without lack of generality, since $N(\bigwedge_{i=1,n} p_i) \geq \alpha \Leftrightarrow \forall i N(p_i) \geq \alpha$. It means that in terms of possibilistic logic expressions we have the

logical equivalence $(\bigwedge_{i=1,n} p_i, \alpha) \equiv \bigwedge_{i=1,n} (p_i, \alpha)$. In other words, any weighted logical formula put in Conjunctive Normal Form is equivalent to a set of weighted clauses. This feature considerably simplifies the proof theory of possibilistic logic. The basic inference rule in possibilistic logic put in clausal form is the resolution rule:

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

Classical resolution is retrieved when all the weights are equal to 1. Other valid inference rules for propositional formulas are:

- if p entails q classically, $(p, \alpha) \vdash (q, \alpha)$
(formula weakening)
- $(p, \alpha); (p, \beta) \vdash (p, \max(\alpha, \beta))$ (weight fusion).

Refutation can be easily extended to possibilistic logic. Let K be a knowledge base made of possibilistic formulas, i.e., $K = \{(p_i, \alpha_i)\}_{i=1,n}$. Proving (p, α) from K amounts to adding $(\neg p, 1)$, put in clausal form, to K , and using the above rules repeatedly until getting $K \cup \{(\neg p, 1)\} \vdash (\perp, \alpha)$, where $\perp$ denotes the empty clause. Clearly, we are interested here in getting the empty clause with the greatest possible weight. It holds that $K \vdash (p, \alpha)$ if and only if $K_\alpha \vdash p$ (in the classical sense), where $K_\alpha = \{p \mid (p, \beta) \in K \text{ and } \beta \geq \alpha\}$. See Lang (2001) for algorithms and complexity issues.

An important feature of possibilistic logic is its ability to deal with inconsistency. The level of inconsistency of a possibilistic logic base is defined as

$$\text{inc}(K) = \max\{\alpha \mid K \vdash (\perp, \alpha)\}$$

(by convention $\max \emptyset = 0$). More generally, $\text{inc}(K) = 0$ if and only if $K^* = \{p_i \mid (p_i, \alpha_i) \in K\}$ is consistent in the usual sense.

Semantic aspects of possibilistic logic, including soundness and completeness results with respect to the above syntactic inference machinery based on resolution and refutation, are presented in (Dubois *et al.*, 1994). From a semantic point of view, a possibilistic knowledge base $K = \{(p_i, \alpha_i)\}_{i=1,n}$ is understood as the possibility distribution π_K representing the fuzzy set of models of K :

$$\pi_K(\omega) = \min_{i=1,n} \pi_{\{(p_i, \alpha_i)\}}(\omega) \quad (1)$$

(where $\pi_{\{(p_i, \alpha_i)\}}(\omega) = 1$ if $\omega \models p_i$; $\pi_{\{(p_i, \alpha_i)\}}(\omega) = 1 - \alpha_i$ otherwise). The degree of possibility of ω according to (1) is computed as the complement to 1 of the largest weight of a formula falsified by ω . Thus, ω is all the less possible as it falsifies formulas of higher weights. In particular, if ω is a counter-model of a formula with weight 1, then ω is impossible, i.e. $\pi_K(\omega) = 0$. It can be shown that π_K is the largest possibility distribution such that $N_K(p_i) \geq \alpha_i$, $\forall i=1,n$, i.e., the possibility distribution that allocates the greatest degree to each interpretation in agreement with the constraints induced by K . N_K is the necessity measure associated to distribution π_K , namely $N_K(p) = \min\{1 - \pi_K(\omega, \omega \models \neg p)\}$. It may be the case that $N_K(p_i) > \alpha_i$, for some i , due to logical constraints between formulas in K .

We now introduce a multiple-agent extension of possibilistic logic, by attaching sets of agents to classical and then to possibilistic logic formulas. Individual agents are denoted by letters a or b , or by indexed letter a_i for $i=1, m$. A set of agents is denoted by a capital letter A or B , which may be indexed, i.e. A_j . The set of all agents is denoted by ALL , and an un-instantiated set of agents by X . The case of binary beliefs is first considered, i.e., an agent believes or not a given proposition. Mind that if agent a does not believe p , it does not mean that it believes $\neg p$. Intermediary levels of certainty, as in standard possibilistic logic will be introduced in subsection 3.2.

3.1 Binary certainty

Let (p, a) and (p, A) denote the respective pieces of information “at least agent a believes p ” and “at least all agents in A believe p ”. Strictly speaking, (p, a) is a simplified notation for $(p, \{a\})$. Let $\text{Agent}(p)$ denotes the set of agents that believe p . Then, the semantic understanding of (p, A) is $\text{Agent}(p) \supseteq A$. Note that

$$\text{Agent}(p \wedge q) = \text{Agent}(p) \cap \text{Agent}(q),$$

$$\text{and } (p \wedge q, A) \equiv (p, A) \wedge (q, A),$$

which expresses that a set of agents believes p and q if and only if they both believe p and believe q . This agrees with $N(p \wedge q) = 1 \Leftrightarrow N(p) = 1 \text{ and } N(q) = 1$.

Consider requirements for syntactic inference. For each agent a , classical logic is supposed to be valid:

- if p entails q classically, $(p, A) \vdash (q, A)$

Since classical logic holds for any agent, and (p, A) means that $\forall a \in A, (p, a)$, we can postulate beliefs are preserved by weakening of the set of agents :

- if $B \subseteq A$, $(p, A) \vdash (p, B)$

So syntactic entailment is defined by these two constraints:

$$(p, A) \vdash (q, B) \text{ if and only if } p \vdash q \text{ and } B \subseteq A.$$

Derived inference rules are:

- $(\neg p \vee q, a); (p \vee r, a) \vdash (q \vee r, a)$,

(which is nothing but the counterpart of the standard resolution rule for agent a);

- $(\neg p \vee q, A); (p \vee r, A) \vdash (q \vee r, A)$;

more generally

$$(\neg p \vee q, A); (p \vee r, B) \vdash (q \vee r, A \cap B)$$

(resolution rule)

since for each agent in $A \cap B$, both $\neg p \vee q$ and $p \vee r$ are true.

If $[p] \subseteq \Omega$ denotes the set of models of p , the set of models of (p, A) is a subset of the Cartesian product $\Omega \times ALL$. Namely

$$[(p, A)] = ([p] \times A) \cup (\Omega \times A^c),$$

since all agents in A believe (at least) p , while other agents are left uncommitted. On such a basis a semantic entailment can be defined such that $(p, A) \models (q, B)$ if and only if $[(p, A)] \subseteq [(q, B)]$. This is equivalent to $p \models q$ and $B \subseteq A$. Hence $(p, A) \vdash (q, B)$ if and only if $(p, A) \models (q, B)$. Indeed, this case if $A \subseteq \text{Agent}(p)$ then $A \subseteq \text{Agent}(q)$ since $p \models q$, hence $B \subseteq \text{Agent}(q)$.

The fusion of the set of agents is also a valid rule:

- $(p, A); (p, B) \vdash (p, A \cup B)$.

The beliefs of a subset of agents may not be consistent with the ones held by another disjoint subset of agents. Mind that this does not create inconsistency, as now illustrated by the following example.

Example. $K = \{(\neg p \vee q, A); (\neg p \vee r, A); (\neg q \vee r, ALL); (\neg p \vee \neg r, B); (p, ALL); (q, A); (r, C)\}$ with $A \cap B = \emptyset$, $C \subset A$ and $A \subset ALL$.

Note that K^* is inconsistent, while K is not. Indeed, one can infer from K , the formulas (q, A) , (r, A) , (r, C) and $(\neg r, B)$. Observe that (r, C) is subsumed by (r, A) since $A \cup C = A$. Moreover, there is no contradiction between (r, A) and $(\neg r, B)$ since $A \cap B = \emptyset$. This illustrates the fact that disjoint sets of agents may have opposite beliefs, without creating inconsistency. Indeed $(\perp, \emptyset)$ does not express any contradiction, while $(\perp, A)$ does if $A \neq \emptyset$.

As in standard possibilistic logic, in order to try to establish that it is true that a set of agents believe r , one can proceed by refutation, i.e., one adds to K the formula $(\neg r, ALL)$. Thus, in the above example, one can obtain the empty clause either under the form $(\perp, C)$ or under the form $(\perp, A)$, as it can be checked. Then, we get $(\perp, A \cup C) = (\perp, A)$, from which one can deduce that (r, A) .

3.2 Graded certainty

We now introduce the multiple agent extension of possibilistic logic, by supposing that a piece of belief encoded by (p, α) is held by an agent, say a , and should be distinguished from pieces of belief of the form (q, β) held by another agent b .

Reasoning will combine the basic patterns of reasoning of sections 2 and 3.1 together. Syntactically the fact that “agent a believes p at least at level α ” is denoted by the expression $(p, \alpha/a)$. Similarly, $(p, \alpha/A)$ will encode “all the agents in A believe p at least at level α ”, while $(p, \{\alpha/a, \beta/b\})$ means “agent a believes p at least at level α , and agent b at level at least β ”. We shall also write $(p, a/A \cup \beta/B)$ for “at least the agents in A believe p at least at level α , and the agents in B at level at least β ”. In particular, with these notations the following rewritings are allowed:

$$1/A \cup 1/B = 1/(A \cup B) = A \cup B$$

$$\alpha/\{a, a'\} = \{\alpha/a, \alpha/a'\} = \alpha/\{a\} \cup \alpha/\{a'\} = \alpha/a \cup \alpha/a',$$

and more generally $\{\alpha/A, \beta/B\} = \alpha/A \cup \beta/B$.

This kind of information can be compactly written as (p, F) where F is the fuzzy set of agents that believe p , is understood as $\forall a \in ALL, N_a(p) \geq \mu_F(a)$. Then the fuzzy set of models of (p, F) on $\Omega \times ALL$, is defined as

$$\mu_{[p, F]}(\omega, a) = \max(\mu_{[p]}(\omega), 1 - \mu_F(a))$$

For instance, if $F = \alpha/A$, it yields $\mu_{[p, F]}(\omega, a) = 1$ if $\omega \models p$ and $a \in A$, or if $a \notin A$ whatever ω ; $\mu_{[p, F]}(\omega, a) = 1 - \alpha$ if $\omega \models \neg p$ and $a \in A$.

Then the following inference rule generalizes the resolution rules of sections 2 and 3.1.

$$(\neg p \vee q, \{\alpha/a, \beta/b\}); (p \vee r, \{\gamma/a, \delta/b\})$$

$$\vdash (q \vee r, \min(\alpha, \gamma)/a, \min(\beta, \delta)/b)$$

where possibly $\alpha = 0$, $\beta = 0$, $\gamma = 0$, or $\delta = 0$, with $\{\alpha/a, 0/b\} = \{\alpha/a\}$. It can be written for any subset of agents rather than only two. For homogeneous subsets of agents the resolution rule can be written more compactly as

$$(\neg p \vee q, \alpha/A); (p \vee r, \beta/B) \vdash (q \vee r, \min(\alpha, \beta)/(A \cap B))$$

The generalized weakening and fusion rules write:

- if p entails q classically, $(p, \alpha/A) \vdash (q, \alpha/A)$

(formula weakening)

- if $B \subset A$, $\alpha \geq \beta$, $(p, \alpha/A) \vdash (p, \beta/B)$

(weakening of a fuzzy set of agents)

- $(p, \alpha/A); (p, \beta/B) \vdash (p, \alpha/A \cup \beta/B)$

(fusion of fuzzy sets of agents)

where $\cup$ is Zadeh’s fuzzy set union, i.e. if $a \in A \cap B$, $\max(\alpha, \beta)/a$ appears in $\alpha/A \cup \beta/B$.

4 A BOOLEAN VIEW OF POSSIBILISTIC FORMULAS

Observe that a possibilistic formula such as (p, α) is crisply entailed or not from a possibilistic knowledge base K , since either the inequality $N_K(p) \geq \alpha$ holds and then (p, α) can be regarded as certainly true, or $N_K(p) < \alpha$ holds and (p, α) is certainly false. In this sense, it is tempting to consider (p, α) as a Boolean formula. On the other hand, the semantics of possibilistic logic leads to interpreting (p, α) as a multiple-valued proposition having a fuzzy set of models.

In order to make sense of this apparent contradiction, consider the formula $(p, 1)$. It expresses that p is certainly true (and stands for $\models p$ in classical notation). While the set of models $[p]$ of p is a subset of interpretations (objective states of the world), the set of (meta-)models $[(p, 1)]$ of $(p, 1)$ is a subset of information states, that is a subset of the power set of Ω , namely all $A \subseteq \Omega$ such that $A \subseteq [p]$. Since from “ p is certainly true”, it follows that (for the agent asserting it) “ p is true”, the set of information states compatible with $(p, 1)$ only leaves the set of models of p as possible states of the world ($[p]$ is the weakest information state coherent with $(p, 1)$).

Similarly the set of meta-models $[(p, \alpha)]$ of (p, α) is a subset of (fuzzy) information states, namely the set of possibility distributions π on Ω such that $N(p) \geq \alpha$. The possibility distribution $\pi = \max(\mu_{[p]}(\omega), 1 - \alpha_i)$ is the weakest information state compatible with (p, α) .

Since the act of asserting (p, α) is crisp, it can be encoded as a crisp formula to which classical connectives can be applied, considering $[0, 1]^\Omega$ as the corresponding set of (meta-)interpretations. Then $\neg(p, \alpha)$ means $N(p) < \alpha$, not to be confused with $(\neg p, \alpha)$, that encodes $N(\neg p) \geq \alpha$. Similarly, a disjunction $(p, \alpha) \vee (q, \beta)$ refers to the set of information states such that one or both of $N(p) \geq \alpha$, $N(q) \geq \beta$, hold. Clearly $(p, \alpha) \vee (q, \beta)$ is not equivalent to $(p \vee q, \max(\alpha, \beta))$, it only entails the latter, since $N(p \vee q) \geq \max(N(p), N(q))$,

while $(p, \alpha) \wedge (q, \beta)$ is equivalent to $(p \wedge q, \min(\alpha, \beta))$, an equivalence agreeing with usual possibilistic logic bases.

For instance, it is possible to handle possibilistic inference as a classical inference rule. Namely, modus ponens can be rewritten as

$$\neg(\neg p \vee q, \alpha) \vee \neg(p, \beta) \vee (q, \min(\alpha, \beta)); (\neg p \vee q, \alpha); (p, \beta) \vdash (q, \min(\alpha, \beta)).$$

Indeed the first premise reads “if $(\neg p \vee q, \alpha)$ and (p, β) hold then $(q, \min(\alpha, \beta))$ holds also”.

Other valid possibilistic inference rules can be handled in the same way. For instance, $\neg(p, \alpha) \vee (p, \beta)$ holds provided that $\beta \leq \alpha$. Assuming (p, α) , we get (p, β) by modus ponens. This result perfectly agrees with the ‘weight weakening’ inference rule recalled in section 2.

Note as a particular case the tautology $\neg(p, \alpha) \vee (p, \alpha)$, which itself generalizes $\neg p \vee p$, retrieved for $\alpha = 1$, i.e., $\neg(p, 1) \vee (p, 1)$. The latter expression should not be confused with $(\neg p, 1) \vee (p, 1)$ which is not a tautology, since this is the syntactic counterpart of the false claim “either $N(\neg p) = 1$ or $N(p) = 1$ ”. The latter does not question the excluded middle law, but sheds light on the difference between truth and certainty.

More generally, any entailment about necessity measures can be written at the syntactic level. Thus, for instance, $N(p \wedge q) = \min(N(p), N(q))$ entails $N(p) = 0$ or $N(\neg p) = 0$, reads : $\neg(p, \alpha) \vee \neg(q, \beta) \vee (p \wedge q, \min(\alpha, \beta))$ for any $\alpha > 0$ and $\beta > 0$, which once particularized with $q = \neg p$, yields $\neg(p, \alpha) \vee \neg(\neg p, \beta) \vee (p \wedge \neg p, \min(\alpha, \beta))$, i.e. $\neg(p, \alpha) \vee \neg(\neg p, \beta) \vee (\perp, \min(\alpha, \beta))$, still equivalent to $\neg(p, \alpha) \vee \neg(\neg p, \beta)$. It indeed means $N(p) < \alpha$ or $N(\neg p) < \beta$ for any $\alpha > 0$ and $\beta > 0$.

Remark: Other modalities

Another interesting type of piece of information worth handling in reasoning is of the form “at least one agent in A believes p ”. This will be denoted by $(p, \exists A)$, with the intended meaning that $\exists a \in A (p, a)$, or if we prefer $\text{Agent}(p) \cap A \neq \emptyset$, where $\text{Agent}(p)$ is the set of agents that believe p . Observe that $(p, \exists ALL)$ encodes the piece of information “there exists at least one agent that believes p ”. Then the following hybrid resolution rule should be endorsed : $(\neg p \vee q, A); (p \vee r, \exists A) \vdash (q \vee r, \exists A)$

This kind of information could be expressed as a disjunction $(p, a) \vee (p, a') \equiv (p, \exists \{a, a'\})$. More generally,

$$(p, \exists A) \vee (p, \exists B) \equiv (p, \exists (A \cup B)).$$

Note also that the following inference rule should obviously hold

$$(p, A) \vdash (p, \exists A) \quad (\text{quantification weakening})$$

Lastly, the information that the set of agents that believe p is included in A , i. e., $\text{Agent}(p) \subseteq A$ means that no agent outside A believes p . It could be encoded as $\neg(p, \exists A^c)$ with $A^c = ALL - A$. Indeed, $\neg(p, \exists A^c)$ means that $\exists a \in A^c (p, a)$, i.e. no agent in A^c can believe p .

The hybrid resolution rule generalizes into

$$\begin{aligned} & \bullet (\neg p \vee q, \alpha/A); (p \vee r, \beta/\exists B) \vdash \\ & (q \vee r, \min(\alpha, \beta)/\exists B) \text{ if } B \subseteq A. \end{aligned}$$

Another rule applies when $B \subseteq A$ does not hold:

$$\begin{aligned} & \bullet (p, \alpha/A); (\neg p, \beta/\exists B) \vdash (\neg p, \beta/\exists (B - A)) \\ & \text{if } B \cap A^c \neq \emptyset. \end{aligned}$$

while

$$\begin{aligned} & (p, \alpha/A); (\neg p, \beta/\exists B) \vdash (\perp, \min(\alpha, \beta)/\exists B) \\ & \text{if } B \cap A^c = \emptyset, \end{aligned}$$

which is a particular case of the previous rule when $B \subseteq A$ holds. Lastly, the following weakening rule obviously holds.

$$\begin{aligned} & \bullet \text{ If } \beta \leq \alpha \text{ then } (p, \alpha/A) \vdash (p, \beta/\exists A) \\ & \quad (\text{quantification weakening}) \end{aligned}$$

Note that all the above inference rules follow easily from the ones of possibilistic logic writing them for each agent that is involved, according to the semantics of formulas of the form (p, A) and $(p, \exists A)$.

5 NESTED POSSIBILISTIC FORMULAS

Booleanization of possibilistic formulas gives us the capability of embedding them inside other possibilistic formulas. Remember that (p, α) already means that “it is certain at least at level α that p is true”. Thus a nested formula such as $((p, \alpha), 1)$ should mean “it is completely certain that (p, α) is true”, i. e., “it is completely certain that, the statement “it is certain at least at level α that p is true”, is true”. Let us more generally consider the nested formula $((p, \alpha), \beta)$, which reads “it is at least certain at level β that p is certain at level α ”. Since the set of models of $(p, 1)$ is a family of subsets of Ω , the set of models of $((p, 1), 1)$ should be a family of subsets of the power set of Ω (actually, the power set of $[p]$, here). So, interpreting $((p, \alpha), \beta)$ in the crisp way comes down to considering the set of meta-possibility distributions π^m on 2^Ω such that

$$N^m(\pi \models (p, \alpha)) \geq \beta,$$

where π is a standard information state. This is in line with Zadeh’s approach to possibility qualification [13]. Using the many-valued semantics of possibilistic logic on $((p, \alpha), \beta)$ and the Boolean metasemantics on (p, α) leads to reducing $((p, \alpha), \beta)$ to a standard possibilistic formula. Namely:

- it is fully possible that (p, α) represents the information i.e., $\pi^m(\pi) = 1$ if $\pi \leq \pi_{\{(p, \alpha)\}}$,
- and it is possible at level $1 - \beta$ that the information (p, α) is not believed, i. e. consider $\pi^m(\pi) = 1 - \beta$ if $\pi \leq \pi_{\{(p, \alpha)\}}$ does not hold.

Focusing on the least specific possibility distributions on Ω restricted by π^m yields

$$\begin{aligned} \pi^{m*}(\pi) &= 1 \text{ if } \pi = \pi_{\{(p, \alpha)\}}, \\ &= 1 - \beta \text{ if } \pi(\omega) = 1 \ \forall \omega \in \Omega. \end{aligned}$$

Then, in agreement with necessity-qualification, the meta-information $((p, \alpha), \beta)$ is reduced via a disjunctive weighted aggregation to a new possibility distribution π^* . Namely, it yields

$$\begin{aligned} \pi^* &= \max(\min(\pi_{\{(p, \alpha)\}}, 1), \min(1, 1 - \beta)) \\ &= \max(\mu_{[p]}, 1 - \min(\alpha, \beta)). \end{aligned}$$

Thus, the meta-information of the form $((p, \alpha), \beta)$ is reduced to a standard possibilistic formula $(p, \min(\alpha, \beta))$. In

practice, such a reduction $((p, \alpha), \beta)$ expresses a discounting of the information (p, α) . It can be also viewed as a counterpart of the fact that in modal logic systems such as KD45, $\Box\Box p \equiv \Box p$ holds. In particular, $((p, \alpha), 1)$ reduces to (p, α) , and is also equivalent to $((p, 1), \alpha)$. Similarly $(p, 1)$ can be reduced to p : in this view, the certainty that “ p is true” entails that “ p is true”.

Under the discounting view, $(p, \alpha) \vdash ((p, \alpha), \alpha)$ and $((p, \alpha), \alpha) \vdash (p, \alpha)$ are valid inference rules. Then, nesting formulas brings no additional expressivity to standard possibilistic logic, but this will be no longer the case in multiple agent logic, as we shall see in the next section.

More interestingly, it is also allowed to deal with embedded formulas such as $(\neg(p, \alpha) \vee q, \beta)$ for instance (since (p, α) is either true or false), where formulas of different levels are mixed. Thus, the previous formula expresses that it is certain at least at level β that if p is at least certain at level α then q is true”, where q may be also a possibilistic logic formula.

6 TRUSTING OTHER AGENTS VS EXPLOITING OTHER AGENT BELIEFS

Reasoning with mutual beliefs requires the handling of nested expressions such as $((p, \alpha/a), \beta/b)$, $((p, \alpha/A), \beta/b)$, or $((p, \alpha/A), \beta/\exists B)$. However, the meaning of such formulas can be ambiguous. A nested formula like $((p, \alpha/a), \beta/b)$ may mean

- either “agent b believes at level β that agent a believes p at least at level α ”
- or “agent b is believes agent a at level β when a believes p at least at level α ”

The first case is a matter of reasoning with our belief in other agents beliefs. The information state of agent b depicted by this formula is a possibility distribution over possible information states of a (each of which is a possibility distribution on Ω). The weight β bears on the fact that agent b considers (p, α) appears or not among agent a ’s graded beliefs. If $BEL(a)$ is the set of beliefs of a , $\beta \leq N_b((p, \alpha) \in BEL(a))$. In particular agent b may be certain that agent a believes p with some certainty $((p, \alpha) \in BEL(a))$, but agent b may believe $\neg p$ at the same time, without any contradiction. Common knowledge is represented by formulas of the form $((p, ALL), ALL)$, i.e. “each agent believes that each agent believes that p is true”. It entails that $((p, a), b)$ for all pairs of agents a and b .

The second case is a matter of acquiring new beliefs from other agents. Agent b considers agent a as a more or less reliable source of information. The weight β is now the degree to which b trusts a about p . Here the reduction process outlined in the previous section can apply, namely b endorses (p, α) with a discounting effect modelled by weight β , and b eventually ends up with the graded belief $(p, \min(\alpha, \beta))$ which is contradictory with b believing $\neg p$.

Nested well-formed formulas can be recursively defined in the following way.

$\varphi = (p, \alpha_1/A_1 \cup \dots \cup \alpha_k/A_k)$ where p is a proposition (built

from a given language) and the A_i ’s are subsets of agents, and the α_i ’s are certainty weights.

$(\varphi, \alpha_1/A_1 \cup \dots \cup \alpha_k/A_k)$ where φ is a logical combination of (Boolean-viewed) possibilistic formulas, with connectives $\neg, \wedge, \vee$.

Nested formulas can be handled by means of the following recursive inference rules, valid whatever their interpretation, in terms of trust or mutual beliefs. Namely,

- if $\varphi \vdash \psi$ then $(\varphi, \alpha/A) \vdash (\psi, \alpha/A)$
- $(\varphi, \alpha/A); (\psi, \beta/B) \vdash (Res(\varphi, \psi), \min(\alpha, \beta)/(A \cap B))$
- $(\varphi, \{\alpha/a, \beta/b\}); (\psi, \{\gamma/a, \delta/b\}) \vdash (Res(\varphi, \psi), \{\min(\alpha, \gamma)/a, \min(\beta, \delta)/b\})$

where $\varphi; \psi \vdash Res(\varphi, \psi)$, and $\vdash$ denotes the recursive application of any valid inference rule, such as the ones described in the previous sections. For instance, consider the following rule involving hybrid labels:

$$((p, \beta/B), \alpha/A); ((\neg p \vee q, \beta'/\exists B), \gamma/C) \vdash ((q, \min(\beta, \beta')/\exists B), \min(\alpha, \gamma)/(A \cap C)).$$

Moreover, the following particularization rule is obviously valid:

- $((p, \alpha/A), \beta/B) \vdash ((p, \alpha/a), \beta/b)$
where $a \in A$ and $b \in B$.

Also, from $((p, \alpha/a), \beta/b)$ and $((\neg p \vee q, ALL), ALL)$ one can deduce $((q, \alpha/a), \beta/b)$.

When nested formulas represent one agent’s confidence into a piece of information believed by another agent (not the confidence in the fact that the other agent possesses this belief), a special inference rule which can be called “endorsement” makes sense:

$$((p, \alpha/a), \beta/b) \vdash ((p, \alpha/b), \beta/b)$$

by which agent b adds agent a ’s belief to his own belief base, with some caution (expressed by the discounting factor β). Then, one can apply the discounting rule

$$((p, \alpha/b), \beta/b) \vdash (p, \min(\alpha, \beta)),$$

which reduces meta-beliefs to simple beliefs. Lastly, let us suggest the practical interest of this inference machinery on an illustrative example involving mutual trust between agents.

Example: Let Peter (pt) and Mary (m) be two agents. They have some common knowledge about the fact that for being able to go (ag) together to a particular place on a Sunday, what they plan to do, they can use either the train (t) or Mary’s car (c), but nothing else. Thus we have

$$((\neg ag \vee c \vee t, ALL), ALL); ((\neg t \vee ag, ALL), ALL); ((\neg c \vee ag, ALL), ALL).$$

Mary believes that there is a train running on Sunday, and that her car has a breakdown, while Peter is sure that there is no train going to the place they want to go on Sunday, while he believes Mary’s car is working. Thus we have

$$\begin{aligned} & - (t, \alpha/m); (\neg c, \delta/m) \\ & - (c, \beta/pt); (\neg t, 1/pt) \end{aligned}$$

Thus, by particularization from $((\neg t \vee ag, ALL), ALL)$ and $((\neg c \vee ag, ALL), ALL)$, one gets $((\neg t \vee ag, m), m)$ and $((\neg c \vee ag, pt), pt)$, and then by reduction $(\neg t \vee ag, m)$ and $(\neg c \vee ag,$

pt). Then, it yields $(ag, \alpha/m)$ and $(ag, \beta/pt)$, thus $(ag, \{\alpha/m, \beta/pt\})$ and then $(c \vee t, \{\alpha/m, \beta/pt\})$ follows using $((\neg ag \vee c \vee t, ALL), ALL)$.

Suppose now that Mary tells Peter about her car problem, i. e., Peter adds $((\neg c, \delta/m), \gamma/pt)$ to his belief set, where γ reflects his trust in Mary's information relevance. Note that if Peter fully trusts Mary then $\gamma = 1$, and by endorsement $((\neg c, \delta/pt), 1/pt)$, which reduces to $(\neg c, \delta/pt)$. It expresses that Peter endorses Mary's judgement. Since we have $(\neg ag \vee c \vee t, pt)$, it yields $(\neg ag \vee t, \delta/pt)$. This reveals to Peter that $(\neg ag, \delta/pt)$, since $(\neg t, 1/pt)$. Hence, Peter is convinced by Mary that they cannot go. The conclusion about 'ag' has been reversed by addition of the new piece of information $((\neg c, \delta/m), 1/pt)$, provided that Mary believes here car is down sufficiently strongly, i.e. this piece of information does not get drowned by the level of inconsistency (here equal to β) of Peter's belief base, i. e., $\beta < \delta$. However, observe that:

- in case Peter *does not trust* Mary on her judgment, he adds to his epistemic state $((\neg c, \delta/m), \gamma/pt)$. After endorsement and discounting it yields $(\neg ag \vee t, \min(\gamma, \delta)/pt)$. If $\gamma < \beta$ Mary's information is drowned in Peter's beliefs regardless of how strong she believes her car to be down, and he will believe they can go (ag), although he knows there is no train because he does not believe Mary is competent about cars (not that he thinks she is lying).
- in case Peter only believes that Mary thinks her car is down, even if she did not tell him, he may only add $((\neg c, \delta/m), 1/pt)$ to his belief base, while not endorsing $(\neg c, \delta)$. Then, this meta-information cannot be used without additional meta-knowledge about what beliefs Peter would entertain if he knew Mary's beliefs. Here, Peter sticks to his belief they can go (ag) with Mary's car, since he knows there is no train.

7 CONCLUDING REMARKS

In this paper, we have outlined extensions of possibilistic logic that handle reasoning about the beliefs of agents and their mutual beliefs. An important distinction is laid bare between agents that consider other agents as more or less reliable information sources, and agents that think they know what other agents' beliefs are and take action on this basis. Potentially, it considerably extends the representation and the inferential power of standard possibilistic logic. This framework seems to offer a reasonable compromise between expressive power and a computational cost. The latter should remain close to the one of classical logic, since we continue to benefit from the computational simplicity of possibilistic logic.

Clearly, many issues remain to discuss. In a longer paper,

it would be interesting to develop the formal semantics of this extended possibilistic logic in terms of fuzzy sets of agents, following ideas already suggested in (Dubois *et al.*, 1992). It remains also to separately develop the multiagent possibilistic logic of trust and the one for reasoning about other agent's belief. The effective potential of such representations in multiagent deliberation systems should be explored, as opposed to modal multiagent logics systems.

We have implicitly assumed the commensurability of the scales used by the different agents for grading their beliefs. It should be possible to handle different scales with partial information about the relative positions of their levels using a recent extension of possibilistic logic with symbolic weights (Benferhat and Prade, 2005). One may also handle cardinality information, such as "at least k agents in A believe p at least at level α ", denoted $(p, \alpha/k:A)$. For instance, a rule such as $(p, \alpha/k:ALL); (\neg p \vee q, \beta/ALL) \vdash (q, \min(\alpha, \beta)/k:ALL)$ clearly holds.

It would be also of interest to consider agents having more or less strong beliefs and preference, and beliefs on other agent beliefs and preferences in order to be able to reason about actions agents should perform. The framework of possibilistic can handle both preference and beliefs, and introducing multiagent decision processes in the possibilistic logic format should be feasible.

REFERENCES

- [1] Benferhat S., Dubois D., Prade H. (1997) Non-monotonic reasoning, conditional objects and possibility theory. *Art. Intel.*, 92, 259-276.
- [2] Benferhat S., Prade H. (2005) Encoding formulas with partially constrained weights in a possibilistic-like many-sorted propositional logic. *Proc. IJCAI'05, Edinburgh*, 1281-1286.
- [3] Cohen P. R., Levesque H. J. (1990) Intention is choice with commitment. *Artificial Intelligence*, 42, 213-261.
- [4] Dubois D., Lang J., and Prade H. (1992) Dealing with multi-source information in possibilistic logic. *Proc. ECAI'92, Vienna, Austria*, Aug. 3-7, 1992, Wiley, New York, 1992, 38-42.
- [5] Dubois D., Lang J., and Prade H. (1994) Possibilistic logic. In: *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3 (D.M. Gabbay *et al.*, eds.), Oxford University Press, 439-513.
- [6] Dubois D., Prade H. (2004) Possibilistic logic: a retrospective and prospective view. *Fuzzy Sets and Systems*, 144, 3-23.
- [7] Fariñas del Cerro L., Herzig A., A modal analysis of possibility theory, *Fundamentals of Artificial Intelligence Research (FAIR'91) LNCS*, Vol. 535, Springer Verlag, Berlin, 1991, pp. 11-18.
- [8] Gabbay D. (1996) *Labelled Deductive Systems* vol. 1, Oxford University Press, Oxford, UK.
- [9] Herzig A., Longin D. (2004) C&L intention revisited. *Proc. KR2004*, 527 - 535.
- [10] Lang J. (2001) Possibilistic logic: complexity and algorithms. In: Vol. 5 of the *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Kluwer Acad. Publ., Dordrecht, 179-220.
- [11] Lorini E., Tummlini L., Herzig A. (2005) Establishing mutual beliefs by joint attention: towards a formal model of public events. *Proc. of CogSci05, Stresa (Italy)*, July 21 -23.
- [12] Zadeh L. A. (1978) Fuzzy sets a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1, 3-28.
- [13] Zadeh L. A. (1978) PRUF-a meaning representation language for natural languages, *Int. J. Man-Machine Studies* 10, 395-460.