Reasoning about the opinions of groups of agents

Asma Belhadi[1] and Didier Dubois[2] and Faiza Khellaf-Haned[1] and Henri Prade[2]

1. LRIA, Université des Sciences et de la Technologie Houari Boumediene, Algérie
BP 32 El Alia 16111 Bab Ezzouar, Alger, asma_b_07@yahoo.fr, hanedfaiza@yahoo.com
2. IRIT, CNRS & Université de Toulouse, France
118 route de Narbonne, 31062 Toulouse Cedex 09, France {dubois, prade}@irit.fr

Abstract. The paper presents a multiple agent counterpart and then a multiple agent extension of possibilistic logic. In the multiple agent counterpart, formulas are pairs (p, A) where p is a proposition and A a subset of agents, with the intended meaning that "all agents in A believe that p is true". Then the logic is extended with the introduction of certainty levels (e.g., "all agents in A believe that p is true at least at level α"). The semantics is in terms of interpretations associated with (fuzzy) sets of agents. Soundness and completeness results can be established. This setting enables us to distinguish between the global consistency of a set of agents, and their individual consistency (both can be a matter of degree). In particular, given a set of multiple agent possibilistic formulas, one can compute the subset of agents that are individually consistent to some degree.

1 Introduction

Agents have their own beliefs or opinions about the state of the world. Agents are not always fully consistent. Conflicting opinions are often encountered in groups of agents. Besides, in general, beliefs may be a matter of degree, when agents are not fully certain about them. In this paper, we consider *the problem of reasoning about the opinions of groups of agents*. Note that the problem is not so much to reason from the information provided by a set of agents viewed as distinct sources and to try to take the best of it in a fusion process (as in, e.g., [5]), in order to locate in the best possible way where the truth should be. Rather, the problem is to consider groups of agents, to understand what claims they support, what other groups they are in conflict with and about which issue, to distinguish the individual inconsistency of agents from the global inconsistency of a group of agents. In other words, we are interested in reasoning about group of agents as much as with their beliefs; an example of such a concern may be already found in [1] where possible answers in a fusion process are associated with the set of sources that supports the answer.

In the following, we provide a semantics for a multiple agent logic already briefly proposed in [7], whose formulas are made of pairs (p, A) where p is a proposition and A a subset of agents. Such a formula encodes the information that all the agents in A believe that p is true. In spite of its modal flavor, the formula (p, A) remains close to a classical logic formula, and has a semantics in terms of a labelled set of interpretations. Although a parallel can be made with possibilistic logic [6], where propositions are associated with certainty levels expressing the strength with which the propositions are believed to be true, (p, A) should not be just understood as another way of expressing the strength of the support in favor of p (the larger A, the stronger the support), but rather as a piece of information linking a proposition with a group of agents (which may be thought as a set defined by its extension, e.g., {Peter, Paul, Mary}, or intensionally by a set of attributes e.g., {over 50 years old, educated, female}). Still A cannot be equated here with a set of logical arguments in favor of p (as in [10]).

2 Reasoning on the beliefs of subsets of agents

The section presents an overview of multiple agent logic, and its semantics. Multiple agent logic parallels possibilistic logic. Then, the possibilistic logic extension of multiple agent logic, with the introduction of intermediary certainty levels, is outlined.

A detailed study of this possibilistic multiple agent logic can be found in [3].

2.1 Multiple agent logic

We consider a set of agents denoted by All. Capital letters $A, B, ...$ denote subsets of agents in All. At the syntactic level, a formula in multiple agent logic (MA-logic for short) is a pair (p, A) where p is a formula in a propositional language, and $A \subseteq All$. A MA-logic base K is a set of such pairs, namely $K = \{(p_i, A_i) \mid i = 1, m\}$, understood as the conjunction of these pairs. The formula (p, A) is supposed to express that (*at least*) all the agents in subset A believe that p is true. A may be any subset of All, including $\emptyset$ and All. However, if (p_i, A_i) belongs to a base K, it is normally assumed that $A_i \neq \emptyset$, since $(p_i, \emptyset)$ is a void information, meaning that the subset of agents believing p contains the empty set.

The inference rules are

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

whose meaning fits the intuition (any agent in $A \cap B$ believes both $\neg p \vee q$ and $p \vee r$, and if any agent in A and any agent in B believes p, then any agent in the union believes it as well. Note that $(p, A) \vdash (q, B)$ as soon as $p \vdash q$ and $A \supseteq B$.

Proving (p, A) from K, denoted $K \vdash (p, A)$, amounts to obtaining $(\bot, A)$ by refutation, namely to get it from $K \cup \{(\neg p, All)\}$, by repeated application of the inference rules.

The inconsistency associated to a MA-logic base K is the set of agents

$$inc(K) = \bigcup \{A | K \vdash (\bot, A)\}$$

The set $inc(K) \subseteq All$ is the set of agents that are *individually* inconsistent. Note that one may have $inc(K) = \emptyset$ even if $K^* = \{p_i | (p_i, A_i) \in K\}$ is classically inconsistent, as shown by the following example $K = \{(p, A), (\neg p, \overline{A})\}$, where the set of agents is partitioned into the two subsets of agents A and $\overline{A}$ having opposite consistent beliefs.

2.2 Comparison with possibilistic logic

One can notice a striking parallel with standard possibilistic logic, where formulas are pairs (p_i, α_i), p_i is a proposition, and α_i is a certainty level usually belonging to a totally ordered scale $\mathcal{S}$, and a base $K = \{(p_i, \alpha_i) \mid i = 1, m\}$ is a conjunction of such pairs. The inference rule is $(\neg p \vee q, \alpha), (p \vee r, \beta) \vdash (q \vee r, \min(\alpha, \beta))$. There is no need here of an explicit use of the rule $(p, \alpha), (p, \beta) \vdash (p, \max(\alpha, \beta)))$, since scale $\mathcal{S}$ is totally ordered, which is not the case for All. Then, the inconsistency level of K is defined by $inc(K) = \max\{\alpha | K \vdash (\bot, \alpha)\}$, where, $K \vdash (p, \alpha)$ iff $K^*_\alpha \vdash p$ and $\alpha > inc(K)$ where $K^*_\alpha = \{p_i \mid (p_i, \alpha_i) \in K, \alpha_i \geq \alpha\}$. Then $inc(K) = 0$ iff K^* is consistent, with $K^* = \{p_i | (p_i, \alpha_i) \in K\}$. As we have just seen, the exact counterpart of this latter result does not hold any longer in MA-logic, and we cannot either reason from the α-level cuts of K, since All is only partially ordered.

A standard possibilistic logic base K has a semantics in terms of a possibility distribution $\forall \omega \in \Omega, \pi_K(\omega) = \min_{i=1,m} \max([p_i](\omega), 1 - \alpha_i)$ where $[p_i](\omega) = 1$ if $\omega \vDash p_i$ and $[p_i](\omega) = 0$ otherwise. Thus, the possibility distribution π_K is a mapping from the set of interpretations Ω to the scale $\mathcal{S}$ which computes the level of possibility of each interpretation. Thus, π_K is the fuzzy set conjunction of the possibility distributions associated with each formula $\forall \omega, \pi_{\{(p_i,\alpha_i)}(\omega) = [p_i](\omega)$ if $\omega \vDash p_i$, while the interpretations that violate p_i receive a possibility level $\pi_{\{(p_i,\alpha_i)}(\omega)$ equal to $1 - \alpha_i$, which is all the smaller as the violated formula has a high certainty level ($1 - (\cdot)$ is the reversing map of $\mathcal{S}$). Then, it can be checked that the semantics of (p_i, α_i) is the constraint $N(p_i) \geq \alpha_i$ where N is a necessity measure such that $N(p \wedge q) = \min(N(p), N(q))$, dual of the possibility measure $\Pi(p) = 1 - N(\neg p)$, where $\Pi(p) = \max_{\omega:\ \omega \vDash p} \pi_K(\omega)$ is the possibility measure associated with the possibility distribution π_K. Moreover, the semantic entailment is defined by $K \vDash (p, \alpha)$ iff $\forall \omega, \pi_K(\omega) \leq \pi_{\{(p,\alpha)\}}(\omega)$. It can be shown that $inc(K) = 1 - \max_\omega\ \pi_K(\omega)$ and possibilistic logic is sound and complete wrt this semantics.

2.3 Semantics of MA-logic

Similarly, the semantics of a MA-logic base is a set-valued distribution from Ω to 2^{All}. The distribution $\pi_{\{(p,A)\}}$ associated with formula (p, A) is defined by

$$\pi_{\{(p,A)\}}(\omega) = All \text{ if } \omega \models p;$$
$$\pi_{\{(p,A)\}}(\omega) = \overline{A} \text{ if } \omega \models \neg p.$$

It acknowledges the fact that given the piece of information (p, A) the set of agents that may find $\neg p$ true possible are *at most* the agents that are outside

A, i.e. in $\overline{A}$. More generally, the MA-possibility distribution associated with an MA-base $K\{(p_1, A_1), ..., (p_m, A_m)\}$ is

$$\pi_K(\omega) = All \text{ if } \forall (p_i, A_i) \in K, \omega \models p_i;$$
$$\pi_K(\omega) = \cap\{\overline{A_i} : (p_i, A_i) \in K, \omega \models \neg p_i\} \text{ otherwise.}$$

The set-valued distribution π_K defines set-valued possibility-like and necessity-like measures

$$\Pi(p) = \bigcup_{\omega:\ \omega \vDash p} \pi_K(\omega) \quad \text{and} \quad N(p) = \overline{\Pi(\neg p)}$$

They are respectively the set of agents that find p possible and the set of agents that believe p is true including the inconsistent agents. Then it can be shown that $N(p_i) \supseteq A_i$ if $(p_i, A_i) \in K$ and more generally that $N(p \wedge q) = N(p) \cap N(q)$. The semantic entailment is then defined by

$$K \vDash (p, A) \text{ iff } \forall \omega, \pi_K(\omega) \subseteq \pi_{\{(p,A)\}}(\omega).$$

The set of agents that are individually inconsistent $inc(K)$ is semantically expressed as $inc(K) = \bigcap_{\omega \in \Omega} \overline{\pi_K(\omega)}$. It can be shown that MA-logic is sound and complete wrt this semantics.

It is important to notice that $inc(K) = \emptyset$ means that any agent is *individually* consistent, which is weaker than the condition $\exists \omega, \pi_K(\omega) = ALL$, which in turn means that the agents are *collectively* consistent. This latter condition is equivalent to the consistency of the classical logic base $K^* = \{p_i | (p_i, A_i) \in K\}$.

2.4 Multiple agent possibilistic logic

MA-logic handles all-or-nothing beliefs and can be extended to the handling of multiple agent graded beliefs of the form "all the agents in A believes p at least at level α" denoted $(p, \alpha/A)$. α/A denotes a fuzzy set of agents defined by $\forall a \in All, \alpha/A(a) = \alpha$ if $a \in A$ and $\alpha/A(a) = 0$ if $a \notin A$. The inference rules becomes (where $\cup$ is the max-based fuzzy set union)

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

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

When $\alpha = 1 = \beta$, we retrieve the MA-logic resolution rule, identifying (p, A) with $(p, 1/A)$. When $A = All = B$, we retrieve the possibilistic resolution rule. $inc(K) = \bigcup\{\alpha/A \mid K \vdash (\bot, \alpha/A)\}$ then yields a fuzzy set of individually inconsistent agents, which describes to what extent different subgroups of agents are inconsistent.The semantics of a base $K = \{(p_i, \alpha_i/A_i)\}|i = 1, m$ is now in terms of a *fuzzy* set-valued distribution (i.e. each interpretation is associated with a maximal fuzzy set of agents for whom the interpretation is more or less possible), and such that $N(p_i) \supseteq \alpha_i/A_i$, where $N(p) = \overline{\Pi(\neg p)}$ and $\Pi(p) = \bigcup_{\omega:\ \omega \vDash p} \pi_K(\omega)$. Soundness and completeness results of MA-logic and possibilistic logic can be extended to this case.

3 Concluding remarks

The paper has presented MA-logic and its possibilistic extension that enables a rich handling of inconsistency both in terms of subsets of agents and in terms of levels of certainty. In particular, two formulas such as $(\neg p, A)$ and (p, B) are contradictory only if $A \cap B \neq \emptyset$, i.e. if there exists an agent that believes both p and believe $\neg p$.

One may think of several extensions. On the one hand, one may develop an extension of generalized possibilistic logic [8], a logic that allows not only for conjunctions but also for disjunctions and negations of possibilistic logic formulas such as (p, α), for handling disjunctions and negations of formulas of the form $(p, \alpha/A)$. One may start with a multiple agent extension of MEL (short for "Meta Epistemic Logic") [2] (which can be viewed as the Boolean restriction of generalized possibilistic logic), which would allow us to consider the disjunction and the negation of formulas like (p, A), and to express for instance that *at most* the agents in some subset believe p.

On the other hand, one might also take into account trust information about information transmitted between agents [4, 9]. For instance, assume agent a trusts agent b at level θ, which might be written $(b, \theta/\{a\})$, assimilating b to a proposition standing for "what b says". Then together with $(p, \alpha/\{b\})$ (agent b is certain at level α that p is true), it would enable us to infer $(p, \min(\alpha, \theta)/\{a\})$ in a possibilistic logic manner. Indeed, substituting (p, α) to b in $(b, \theta/\{a\})$ yields $((p, \alpha), \theta/\{a\})$, which indeed reduces to $(p, \min(\alpha, \theta)/\{a\})$ (see [7] for a justification).

References

1. Assaghir, Z., Napoli, A. , Kaytoue, M., Dubois, D., Prade, H. Numerical information fusion: Lattice of answers with supporting arguments. Proc. 23rd IEEE Int. Conf. on Tools with Artificial Intelligence (ICTAI 2011), Boca Raton, Fl. Nov. 7-9, 621-628, 2011.
2. Banerjee, M., Dubois, D. A simple modal logic for reasoning about revealed beliefs. Proc. 10th European Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09), (C. Sossai, G. Chemello, eds.), Verona, July 1-3, Springer, LNCS 5590, 805-816, 2009.
3. Belhadi, A., Dubois, D., Khellaf-Haned, F., Prade, H. Multiple-agent possibilistic logic. J. of Applied Non-Classical Logics, doi = 10.1080/11663081.2013.864470, 2013.
4. Cholvy, L. How strong can an agent believe reported information? Proc. 11th Europ. Conf. on Symb. and Quantit. Appr. to Reas. with Uncert. (ECSQARUÕ11), (Weiru Liu, ed.), Belfast, June 29-July 1, LNCS 6717, Springer, 386–397, 2011.
5. Dubois D., Lang J., Prade H. Dealing with multi-source information in possibilistic logic. Proc. of the 10th Europ. Conf. on Artificial Intelligence (ECAI'92), Vienna, Aug. 3-7, 38–42, 1992.
6. Dubois D., Prade H. Possibilistic logic: A retrospective and prospective view. Fuzzy Sets and Systems, 144, 3–23, 2004.

7. Dubois D., Prade H. Toward multiple-agent extensions of possibilistic logic. Proc. IEEE Int. Conf. on Fuzzy Systems, London, July 23-26,187–192, 2007.
8. Dubois, D., Prade, H., Schockaert, S. Stable models in generalized possibilistic logic. Proc. 13th Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR'12), (Gerhard Brewka, Thomas Eiter, Sheila A. McIlraith, eds.), Roma, June 10-14, AAAI Press, 519–529, 2012.
9. Gutscher, A. Reasoning with uncertain and conflicting opinions in open reputation systems. Electron. Notes Theor. Comput. Sci., 244, 67–79, 2009.
10. Lafage, C., Lang, J., Sabbadin, R. A logic of supporters. In : Information, Uncertainty and Fusion, (B. Bouchon-Meunier, R. R. Yager, L. A. Zadeh, eds.), Kluwer Acad. Publ., 381–392, 1999.