An uncertainty interchange format for multi-agent systems based on imprecise probabilities

Pietro Baroni
Dipartimento di Elettronica per l'Automazione
University of Brescia
Via Branze, 38 - 25123 BRESCIA - Italy
e-mail: baroni@ing.unibs.it

Paolo Vicig
Dipartimento di Matematica Applicata "de Finetti"
University of Trieste
Piazzale Europa, 1 - 34127 TRIESTE - Italy
e-mail: paolo.vicig@econ.univ.trieste.it

Abstract

In multi-agent systems, there is the need to exchange uncertain information between distinct and independently developed software components. This requires that such components share a common uncertainty interchange format and poses, therefore, a serious and still poorly considered problem, in face of the variety of existing uncertainty theories. In fact, imposing that all components adopt the same uncertainty theory is often unrealistic. Defining a common uncertainty interchange format, able to guarantee compatibility with several different approaches, is an open research problem. In this paper we discuss the basic issues that need to be dealt with to face such a problem and formulate an initial proposal based on imprecise probabilities.

1 Introduction

One of the key issues in the development of multi-agent systems is the definition of an interchange format for communication and information exchange among different agents. The most influential proposal in this area is the Knowledge Interchange Format (KIF) [8], which is a language designed for use in the interchange of knowledge among disparate computer systems. Actually, KIF is a prefix version of the language of first order predicate calculus with various extensions to enhance its expressiveness.

As to our knowledge, KIF considers only the case where the information to be exchanged is represented by binary and certain sentences and simply does not address the case where such information is fuzzy and/or affected by uncertainty. This clearly represents a severe limitation, restricting the applicability of this and similar proposals to the application contexts where the importance of fuzziness and uncertainty is negligible.

In this paper we address the problem of defining a generic uncertainty interchange format for the exchange of uncertain information among heterogeneous agents, each one featuring a specific uncertain reasoning paradigm. The main objective of this effort is generality, namely the ability to guarantee compatibility with different existing theories and approaches.

Defining an interchange format for uncertainty is far more difficult than for 'certain' knowledge: in fact, it seems that there is no universally recognized common conceptual background underlying different uncertainty theories and even the same theory may have very different interpretations at the semantic level. The present work aims at analyzing the main issues involved in the definition of an uncertainty interchange format, pointing out the relevant open problems, and presenting a preliminary proposal based on imprecise probabilities.

The paper is organized as follows. In section 2 a generic architecture for interchange is described. In section 3 the problem of defining an uncertainty interchange format is analyzed and a proposal based on imprecise probabilities is formulated. Section 4 deals with inter-format translations and introduces transformation methods from imprecise probabilities into some other well-known representations, namely precise probabilities, belief functions, and possibilities. Finally section 5 summarizes and concludes the paper.

2 An architecture for multi-agent uncertainty interchange

The architectural scheme we consider is based on the assumption that an agent might completely ignore the features of the internal representation adopted by its partners and should not need to negotiate preliminarily any aspect of the interchange. While this assumption prevents the potential advantages of a tailored communication, it is coherent with the goal of maximum generality: for instance, it allows indirect communication through shared data bases, where the agent inserting new information does not know a priori who will eventually access it.

We will now examine the process of exchanging a piece of uncertain information between a sender, called *agent S*, and a receiver called *agent R*.

We assume that each *agent X* uses a specific internal representation language L_X, whereas the common interchange format language is denoted by L_{IF}.

Agent S produces the information to be exchanged and needs to translate it from L_S into L_{IF}. Translated information is then transmitted to *agent R*, which is then in charge of translating it from L_{IF} into L_R.

Since L_{IF} should be, in general, more expressive than any specific language L_X, any translation of the kind $L_X \rightarrow L_{IF}$ should involve a null (or minimal) distortion or loss of information. On the other hand, any translation of the kind $L_{IF} \rightarrow L_X$ may involve some possibly significant information loss or distortion, since it goes from a general format to a more specific one, whose expressiveness may be more limited.

Therefore it may be useful for an *agent X* to include a distortion evaluation component DE_X, able to evaluate the distortion possibly introduced in the translation, by comparing the information originally received and the translation results. The output of DE_X may then be considered by an agent-specific filter F_X, that is in charge of deciding whether to discard the received information or forward it to the internal processing activity. In summary, an *agent X* able to both send and receive information has to be endowed with two mandatory modules, namely an internal-into-common and a common-into-internal translator, and may also include two optional modules, namely a distortion evaluator and an acceptance filter.

For each agent, the specification of all these modules partly depends on the features of the internal format adopted, however both translators and evaluator should be designed according to the same general principles in any agent, in order to guarantee the effectiveness and the correctness of the exchange. On the other hand, the acceptance filter is totally implementation dependent, as an agent should be completely free to decide what to take into account on the basis of any criterion.

In the rest of the paper we will focus our attention on the translation modules, as they represent the minimal mandatory endowment in order to enable the interchange.

3 Defining an uncertainty interchange format

As well-known, many different theories have been developed and are currently investigated in the area of uncertain reasoning. Their general validity, practical applicability and limitations have been and still are debated. Moreover, also the definition of a precise applicability scope, stating which theory should be adopted to face a specific uncertain reasoning problem, is lacking. The above remarks confirm how difficult is the problem of defining a general concept and format for uncertainty representation.

This problem calls for a series of questions to be faced, which are orderly discussed in the following.

3.1 Basic assumptions

Before proceeding with the analysis of the interchange problem, it is necessary to make some assumptions which allow us to better specify and delimit it:

1. the information exchanged concerns uncertainty judgements[1] about (binary) non-conditional events;
2. all agents share a common finite universe of discourse U, and the events about which they formulate their uncertainty judgements are taken from $\wp(U)$, which is the powerset of U. In other words, U is a finite partition made up of *atoms* (pairwise disjoint non-impossible events whose logical sum is the certain event Ω).

Clearly both assumptions are limiting and not completely realistic, but are useful in order to face gradually the problem; it is also important to note that, as will be stressed in (b) of sect. 3.5, they are not forced by the interchange format we adopt.

3.2 Uncertainty judgements

In presence of uncertainty, each possible truth value of an event is associated to a judgement, that qualifies the belief attitude of the agent with respect to the attribution of such a truth value to the event.

In general, it is possible to distinguish two main classes of uncertainty judgements:

- absolute judgements, which concern a single pair (Event, Truth Value)[2] and qualify agent's belief attitude about it through some quantification;
- relational judgements[3] (often also called comparative or qualitative), which define a relation between the belief attitudes concerning two or more ETV pairs (e.g. by stating that the truth of an event is more credible than that of another one, or is more

[1] This notion will be better specified in the sequel.

[2] In the sequel, such a pair will be termed ETV for short.

[3] It is worth noting that the so called 'symbolic approaches' to uncertainty, i.e. the various families of nonmonotonic logic, can be regarded as a special case of relational judgements.

credible than its falsity).

In general, a software agent may be able to deal with both absolute and relational judgements. Being based on two different and, in a sense, complementary basic notions, namely quantification and relation, absolute and relational judgements should be considered as distinct concepts, that can not be easily converted into each other. We believe that a complete uncertainty interchange format should encompass a distinct representation for absolute and relational judgements. However, in order to limit the scope and the extension of this paper, we focus here on absolute judgements.
Absolute judgements can be further classified as precise or imprecise:

- precise judgements associate a single quantification with an ETV pair;
- imprecise judgements associate a set of quantifications with an ETV pair.

Clearly precise judgements are a special case of imprecise judgements.
Imprecise judgements may in turn be classified as:

- crisp, if the sets of quantifications associated with ETV pairs are classical binary sets;
- fuzzy, if the sets of quantifications associated with ETV pairs are fuzzy.

Again, crisp sets are a special case of fuzzy sets.
Clearly, the notion of fuzzy imprecise quantifications is the most general and therefore appears to be a good basis for the definition of the interchange format. However, two difficulties have to be acknowledged:

- since a fuzzy set can in general be defined by any membership function, the formalism may turn out to be very complex, unless the class of admissible membership functions is constrained in some way;
- some of the most known uncertainty theories (imprecise probabilities, belief functions, possibility theory) can be characterized in terms of crisp imprecise quantifications, namely intervals.

For these reasons, while recognizing the potential importance of fuzzy sets for future developments, we limit the scope of the present paper to quantifications expressed through crisp intervals included in the real interval [0, 1], namely through a couple of numbers (a, b): $a \le b; a, b \in [0, 1]$.

3.3 Message structure

After deciding the format for a single uncertainty judgement, it should be defined which is the standard format for an exchange of uncertain information.
We assume that the agents adopt one of the existing standards for inter-agent communication, such as KQML [6] or FIPA ACL [7]: the choice of the communication language is rather indifferent, since our work focuses on contents language, namely on the representation of the information carried by the message rather than on the structure of the message itself. Clearly these two aspects are independent.
As for the contents language, we suggest that the minimal amount of information to be exchanged consists of the complete belief state concerning an event, namely of the uncertainty judgements concerning all its possible truth values. In the case of binary events, the belief state is clearly represented by a pair of judgements. There are two main reasons for exchanging complete belief states:

- in general, judgements concerning different truth values are not tightly constrained (in precise probability theory they are, but in other theories they are not), therefore a complete belief state is necessary to express the overall information one agent has about an event;
- providing partial information about an event (namely an incomplete belief state) may cause an undesirable distortion on the receiver side since the receiver might tend to integrate the incomplete information according to its own internal theory, possibly yielding significantly distorted results with respect to the complete belief state of the sender.

A message contents carrying uncertain information is therefore constituted by the specification of an event, which can be provided using KIF, and by the specification of the relevant belief state, which, in the binary case, is a couple of couples of real numbers within the [0, 1] interval, as specified above.

3.4 Semantics

Considering the semantic aspects of the interchange format, it has to be noted that defining 'the' semantics of a representation is a serious problem even if one limits himself to a single uncertainty theory. Just to mention some of the most known ones: in probability theory, a subjective and a frequentist interpretation do exist; belief functions have been regarded either as a special case of imprecise probabilities or as an autonomous concept related to the representation of evidence [12]; possibilistic reasoning has been given both a preference-based and a similarity-based semantics [5].
Given that even individual theories have not a universally accepted semantics, it is hard to pretend that a generic interchange format has one. What seems to be lacking is a basic ontology of uncertainty and uncertain reasoning, whose definition would provide a reference

framework for comparing different approaches.

Information interchange, however, necessarily requires some common background, so that a piece of information still preserves at least part of its initial meaning once translated first into the interchange format by the sender and then into a new specific format by the receiver. In our opinion, such a background can be provided by imprecise probability theory.

3.5 Imprecise probabilities and the interchange format

In the theory of coherent imprecise probabilities, as developed in [13], it is assumed that the following conjugacy relation holds between the lower ($\underline{P}$) and the upper ($\overline{P}$) probability of an event E:

(1) $\underline{P}(E) = 1 - \overline{P}(\neg E)$.

This enables us to consider lower (alternatively, upper) probabilities only, whenever they are defined on a set of events closed under complementation; in sect. 4 we shall sometimes exploit upper probabilities, here we mainly refer to lower probabilities.

3.5.1 Definition [13] Given an *arbitrary* (finite or not) set of events S, $\underline{P}(\cdot)$ is a *coherent lower probability* on S iff, $\forall$ m, $\forall$ $E_0,\ldots,E_m \in S$, $\forall$ $s_i \geq 0$, i = 0,…,m, defining $I(E)$ as the indicator of E ($I(E) = 1$ if E is true, $I(E) = 0$ if E is false) and putting

$$\underline{G} = \sum_{i=1}^{m} s_i \, [I(E_i) - \underline{P}(E_i)] - s_0[I(E_0) - \underline{P}(E_0)] \, ,$$

it is true that max $\underline{G} \geq 0$.

Def. 3.5.1 weakens de Finetti’s coherence principle [2], and a precise probability (coherent by de Finetti’s definition) is indeed a special case of (coherent) imprecise probability (where $\underline{P}(\cdot) = \overline{P}(\cdot) = P(\cdot)$).

Coherent[4] imprecise probabilities are a very general tool, which generalizes various uncertainty measures. In particular, a *belief function* as defined in [11] is a special case of lower probability [14], and so is a *necessity measure* which can be seen as a special case of belief function (actually, as a consonant belief function) [3] [12], while a *possibility measure* is a special case of upper probability [14].

The main advantages assured by the use of imprecise probabilities for an interchange format are:

(a) no translation is needed from the agent internal representation to the common interchange format whenever the internal representation is based on an uncertainty measure which is a special case of imprecise probability, like the (quite common) ones mentioned above (and others, for instance 2-monotone probabilities [14]). In fact, in such instances the information produced by the agent may be simply read as an imprecise probability in the interchange format, without modifying any of its numerical values.

(b) The constraints in sect. 3.1 might be widely relaxed while keeping on using an interchange format based on imprecise probabilities. In fact, as appears from def. 3.5.1, coherent lower (and upper) probabilities are defined on *arbitrary* sets of events, so there is no need either to put $S = \wp(U)$ or to consider finite sets of events. Also, by the extension theorem [13], an imprecise probability on S can *always* be coherently extended to *any* superset of S, and this allows exchanging information in a *dynamic* setting where the universe of discourse is not fixed.

Further, these features are shared also by generalizations of imprecise probabilities to conditional imprecise probabilities and to (coherent) conditional previsions (the latter are suited for handling information on conditional random numbers) [13].

Of course, there are important ways of expressing uncertainty which are not special cases of imprecise probabilities and therefore would need some translation before using the proposed interchange format: we mention fuzzy judgements, which should be reduced to crisp intervals, and comparative probabilities, for which a realization problem (by means of an imprecise probability) arises, and it is not guaranteed a priori that it always has a solution.

Imprecise probabilities (and in particular belief functions and possibility measures) weaken the tight additivity constraint of precise probabilities $P(E) + P(\neg E) = 1$ replacing it by (1) (substitute $\underline{P}$, $\overline{P}$ with, respectively, *Bel*, *Pl* for belief functions, and N, Π for possibilities, N being a necessity measure, Π a possibility). At a semantic level, the question arises whether constraints among the uncertainty judgements concerning the different truth values of an event should be enforced in the interchange format. Our answer is intermediate: on the one hand, a constraint-free representation is more general than a constrained one, on the other hand, reasonably complete and solid theoretical treatments mainly exist for coherence-based theories.

Therefore, we consider in this work only quantifications that respect the constraints imposed by coherent imprecise probability theory.

[4] We shall usually omit the term ‘coherent’ in the sequel, when referring to coherent imprecise probabilities.

4 Inter-format translations

In the following, we shall consider translations with reference to precise and imprecise probability, belief functions, and possibility theory.
Recalling (a) of sect. 3.5, translations from L_S into L_{IF} are trivial.
Translations from L_{IF} into L_R involve the problem of transforming an imprecise probability into: a precise probability, a belief function, or a possibility. Problems of this kind have no throughout accepted solution: every translation mechanism involves some arbitrariness and is questionable in some respect.
Let μ_L, ν_L be two uncertainty measures on the same set of events J. The subscript L denotes lower uncertainty measures (i.e. lower probabilities, beliefs, necessities). The following points clarify the criteria underlying the translation methods we propose:

(a) when comparing μ_L, ν_L it may appear that μ_L is intrinsically more precise than (or at least as precise as) ν_L. In this case, translations from a format using μ_L into a format based on ν_L should give a uniformly more imprecise evaluation (vice versa when passing from ν_L to μ_L). This *consistency principle* leads operationally to the dominance condition

(2) $\mu_L(E) \geq \nu_L(E), \forall\, E \in J$

but does not determine, in general, a unique ν_L.

(b) The transformed measure ν_L should be as close as possible, in some sense, to the original measure μ_L. If (a) is applied, a simple way to interpret (b) is to choose ν_L in order to minimize

(3) $S = \sum_{E \in J} (\mu_L(E) - \nu_L(E)).$

In the sequel we shall apply (2) and (3) (with the exceptions noted below), as well as analogous conditions when passing from a more precise to a less precise measure or for translations concerning upper uncertainty measures.
The consistency principle will be applied without exceptions. To use it we need to compare the precision of imprecise probabilities with that of each of the other three measures.
Imprecise probabilities are less precise than probabilities, but are more precise than possibilities/ necessities: a couple $(N(E), \Pi(E))$ expressing the necessity and possibility of E is constrained to have the form $[0, \Pi]$ or $[N,1]$ by elementary properties of these measures, but this means that $[N(E), \Pi(E)]$, viewed as an imprecise probability, is always either lower (if $N = 0$) or upper (if $\Pi = 1$) maximally imprecise.

Belief functions appear to be less precise than imprecise probabilities, by the following inferential argument, showing that they may produce less precise inferences: if $\underline{P}(\cdot)$ is an unconditional lower probability defined on the relevant events and $\underline{P}(B) > 0$, it is known [14] that its vaguest (or least-committal) coherent extension on $A|B$ is such that

(4) $\underline{P}(A|B) \geq \underline{P}(A \wedge B)/(\underline{P}(A \wedge B) + \overline{P}(A \wedge \neg B))$

and that equality holds in (4) if $\underline{P}(\cdot)$ is a belief function (actually, also if $\underline{P}(\cdot)$ is 2-monotone).

Principle (b) and (3) cannot be applied when translating imprecise into precise probabilities, and this is easily seen to depend on the equality $P(E) = 1 - P(\varnothing E)$. In this case we shall adopt a 'centrality' criterion, so that $P(E)$ tends to the midpoint between $\underline{P}(E)$ and $\overline{P}(E)$ (see 4.2.1).
Imposing (b) is also conflicting with *preference preservation* requirements like $\mu(E) \geq \mu(F) \Rightarrow \nu(E) \geq \nu(F), \forall\, E, F \in \wp(U)$. In fact, the additional constraints due to preference preservation are often very strong and necessarily widen the imprecision gap between μ and ν. We shall impose preference preservation only occasionally (see 4.2.3), partly giving up (b) (see also [4] for comparisons among conflicting principles in the probability/possibility transformation case).

4.1 Translating belief states about a single event

As noted in sect. 2, the belief state of an event is represented by two couples, concerning respectively the truth and falsity of the event, but since (1) holds in the theories we consider, one couple can be deduced from the other in our framework. Although it is convenient to use both couples in the definition of L_{IF}, especially to allow compatibility with less constrained theories, we shall consider in the sequel, for conciseness, only the couple (l_T, u_T) concerning the truth of E, where l_T (u_T) is interpreted in L_{IF} as the lower (upper) probability of E.
Translations from L_{IF} are straightforward if L_R is based on imprecise probabilities or belief functions (in the latter case, because every lower probability on $\{E, \neg E\}$ is a belief function).
In the case of precise probability, the translation consists of selecting a single value from an interval: the midpoint of the interval is a natural choice.
We thus obtain: $P(E) = (l_T + u_T)/2$.
Translations into possibilities are more articulated: by elementary necessity/possibility properties, and unless $u_T = 1$ or $l_T = 0$, a sort of stretching of the probability interval in one direction is required.

This can be obtained by translating u_T into a possibility value 1, if $1 - u_T < l_T$, i.e. if u_T is closer to 1 than l_T is to 0, or translating l_T into a necessity value 0, if $1 - u_T > l_T$. The case

(5) $1 - u_T = l_T$

shows however a singularity of this method, that is present, though unnoticed, when $P(E) = 0.5$ (which is a special case of (5) with $u_T = l_T$) also in the version proposed for precise probabilities in [4]. In fact, the case of uniform probabilities is equated, in possibility theory, to 'total ignorance': in absence of any preference, we get the extreme assignment $\Pi(E) = 1$, $N(E) = 0$ (and hence $\Pi(\neg E) = 1$, $N(\neg E) = 0$), which should also be the translation of an imprecise probability assignment obeying (5), since again we have no ground for modifying one rather than the other imprecise probability measure.

However in these cases the probability $\rightarrow$ possibility translation operator behaves discontinuously. To exemplify, put $l_T = 0.5$, $u_T = 0.5 + \varepsilon$, with a quite 'small' ε (being therefore close to the total ignorance case considered above): this gives $N(E) = 0.5 + \varepsilon$, $\Pi(E) = 1$ (and hence $N(\neg E) = 0$, $\Pi(\neg E) = 0.5 - \varepsilon$) with a large discontinuity of $N(E)$ (and of $\Pi(\neg E)$).

This shows that even in the simple case of a single event, translations between different theories may involve some inherently problematic aspects and unavoidable distortions.

4.2 Translating belief states about $\tilde{A}(U)$

We propose in this section some methods of translating from L_{IF} into the specific formats considered when the exchanged information concerns all the non-trivial events in $\wp(U)$ and $U = \{e_1, \dots, e_n\}$.

4.2.1 From imprecise to precise probabilities

Given a lower probability $\underline{P}$ on $\wp(U)$, the consistency principle requires that the precise probability P resulting from translation is such that $P(E) \geq \underline{P}(E)$, $\forall E \in \wp(U)$ or equivalently, in terms of upper and lower probability, that $\underline{P}(E) \leq P(E) \leq \overline{P}(E)$, $\forall E \in \wp(U)$.

A straightforward extension of the single event case would lead to translating $\underline{P}$ into $P_m(E)=(\underline{P}(E) + \overline{P}(E))/2$ for all E, but in general P_m is not a coherent precise probability. One way out is that of finding a precise probability P^* which is as close as possible to P_m in some sense[5]. If this is meant as quadratic approximation, we are lead to solve the following minimization problem:

(6) $\min \varphi = \sum_{E \in \tilde{A}(U)} (P^*(E) - P_m(E))^2$

subject to:

$\underline{P}(E) \leq P^*(E)\ \forall E \in \wp(U)$; $\sum_{i=1}^{n} P^*(e_i) = 1$

Problem (6) minimizes a strictly convex function on a convex set[6] and therefore its solution (exists and) is unique. Further, P^* coincides with P_m whenever P_m is a coherent precise probability. Operationally, problem (6) may be solved with standard quadratic programming techniques (see for instance [1]).

4.2.2 From imprecise probabilities to belief functions

To translate a lower probability $\underline{P}(\cdot)$ into a belief function $Bel(\cdot)$ we apply (2) and (3) and determine Bel through its Möbius inverse m_B (also called mass function) which, as well-known [11], is non negative, normalized to 1 over all events in $\wp(U)$ and such that $Bel(E) = \sum_{A \subseteq E} m_B(A)$.

This leads us to the following problem:

(7) $\min \varphi = \sum_{E \in \tilde{A}(U)} (\underline{P}(E) - \sum_{A \subseteq E} m_B(A))$

subject to:

$\sum_{A \subseteq E} m_B(A) \leq \underline{P}(E)$, $m_B(E) \geq 0, \forall E \in \wp(U)$;

$\sum_{E \in \tilde{A}(U)} m_B(E) = 1$.

The feasible region of this problem is always non-empty ($m_B(E) = 0\ \forall E \neq \Omega$, $m(\Omega) = 1$ is a feasible point).

4.2.3 From imprecise probabilities to possibilities

Given an upper probability $\overline{P}(\cdot)$ on $\wp(U)$, a possibility measure $\Pi(\cdot)$ which translates it should respect the consistency constraint:

(8) $\Pi(E) \geq \overline{P}(E)\ \forall E \in \wp(U)$.

Applying also (3) we are lead to minimize

(9) $S = \sum_{E \in \tilde{A}(U)} (\Pi(E) - \overline{P}(E))$.

However this request cannot be always fulfilled: for instance, if $\overline{P}$ is a precise probability P, $U = \{e_1, e_2\}$, and $P(e_1) = P(e_2)$, it is natural to translate P into

---

[5] Other solutions could be considered. The emphasis on approaching P_m may be also motivated by the uses of P_m in a decision theoretic framework [16].

[6] By the coherence of the imprecise probability, the convex set is not empty.

$\Pi(e_1) = \Pi(e_2) = 1$, which maximizes S. Here the preference preservation principle is rather pursued: the resulting possibility must not introduce differences among events which are not given different upper (and lower) probabilities.

We shall now propose a procedure which obeys (8) and generalizes to imprecise probabilities what suggested for precise probability $\to$ possibility transformations in [4].

The procedure works building up the possibility distribution function $\pi(\cdot)$ of $\Pi(\cdot)$, which is equivalent to assigning a possibility value $\Pi(e_i)$ to each atom $e_i \in U$, under the normality condition: $\exists\ e_i : \Pi(e_i) = 1$. Then, as well-known, $\Pi(E)$ is given by:

(10) $\Pi(E) = max_{e_i \Rightarrow E} \{\Pi(e_i)\}$.

The procedure consists of the following steps:

1.a: $\forall\ e_i \in M = \{e_i \in U: 1 - \overline{P}(e_i) < \underline{P}(e_i)\}$,
put $\Pi(e_i) = 1$;

1.b: If $M = \varnothing$ then:
let $M^* = \{e_i \in U: \overline{P}(e_i) = max_{e_j \in U} \{\overline{P}(e_j)\}\}$,
and assign $\Pi(e_i) = 1$ to any $e_i \in M^*$ which minimizes: $\overline{P}(e_i) - \underline{P}(e_i)$.

2: order the atoms of U which have not been assigned a possibility value in step 1 by decreasing upper probability; tied atoms (i.e. atoms having equal upper probability) are ordered by decreasing lower probability. Suppose for simplicity that the ordered sequence is $e_1, \ldots, e_m$, $m < n$, so that
$\overline{P}(e_1) \geq \overline{P}(e_2) \geq \ldots \geq \overline{P}(e_m)$.
For $i = 1, \ldots, m$, define $A_i = e_i \vee e_{i+1} \vee \ldots \vee e_m$;
put $\Pi(e_i) = \overline{P}(A_i)$, unless it is $\overline{P}(e_i) = \overline{P}(e_{i+1})$ and $\underline{P}(e_i) = \underline{P}(e_{i+1})$: in this latter case,
put $\Pi(e_i) = \Pi(e_{i+1}) = \overline{P}(A_i)$.

The motivation for step 1.a is analogous to the single event case of 4.1: having to transform $[\underline{P}(e_i), \overline{P}(e_i)]$ into either $[0, \Pi(e_i)]$ or $[N(e_i), 1]$, the solution introducing as little imprecision as possible is chosen. Step 1.b ensures the normality condition, in case of ineffectiveness of step 1.a. Step 2 gives to the remaining atoms (if any) as little possibility as possible (with the exception of ties, see below) while obeying condition (8), as it can be easily seen using (10).

An interesting question is: to what extent does this procedure tend to minimize the additional imprecision ? Ties can be a first source of imprecision. Steps 1.b and 2 treat differently upper probability ties only if the corresponding lower probabilities are not equal. The aim is both to translate correctly cases like the one mentioned after (9) and to limit the (usually large) imprecision due to the translation into possibilities (for instance, not differentiating two ties $\overline{P}(e_i) = \overline{P}(e_{i+1})$ in 1.b adds $2^{m-i+1} \cdot (\overline{P}(A_i) - \overline{P}(A_{i+1}))$ to the transformation imprecision). In other words, undifferentiated ties increase imprecision, but it does not seem reasonable to order strictly all ties.

Suppose now that no ties arise: then the procedure introduces as little imprecision as possible at each step, but still there may exist, subject to certain conditions, alternative procedures which achieve a smaller S in (9). This second source of imprecision is typical of imprecise probability – possibility transformations. To exemplify, consider the following proposition, whose proof is not difficult and will be omitted.

Proposition Let $1/2 > \overline{P}(e_1) > \overline{P}(e_2) > \ldots > \overline{P}(e_m)$. Consider two possibility measures Π_1 and Π_2, where Π_1 is assigned applying the procedure described above:

(11) $\Pi_1(e_i) = \overline{P}(A_i)$, *for* $i = 1, \ldots, n$;

Π_2 is obtained exchanging the role of two contiguous (with respect to the ordering induced by $\overline{P}$) atoms, e_h, e_{h+1}, in the procedure:

$\Pi_2(e_i) = \overline{P}(A_i)$, *for* $i = 1, \ldots, h-1, h+2, \ldots, n$,
$\Pi_2(e_{h+1}) = \overline{P}(e_h \vee e_{h+1} \vee A_{h+2})$,
$\Pi_2(e_h) = \overline{P}(e_h \vee A_{h+2})$. Then it is:

(12) $S_1 > S_2$ iff $\overline{P}(e_{h+1} \vee A_{h+2}) > \overline{P}(e_h \vee A_{h+2})$.

This kind of situation *cannot* occur if $\overline{P}$ is a precise probability P: in fact, the condition in the right-hand member of (12) is then always false, by the additivity of P (and since $P(e_{h+1}) < P(e_h)$). On the other hand, if $\overline{P}$ is a (coherent) upper probability, it can be verified that the condition in (12) is not vacuous. This is due to the subadditivity of upper probabilities, which makes 'preference inversions' possible, in the sense that conditions like the following may coexist:
$\overline{P}(E_1) < \overline{P}(E_2)$ and $\overline{P}(E_1 \vee F) > \overline{P}(E_2 \vee F)$.
Conditions similar to (12) can be obtained for more complex alternative procedures, thereby suggesting that, apart from ties, the procedure tends to minimize S, if $\overline{P}$ does not differ too much from a precise probability (in the sense that it does not allow inversions).

5 Conclusions

In this paper we addressed the problem of defining an interchange format for uncertain information exchange

between distinct and independently developed software components, such as agents in a multi-agent system. As to our knowledge, this issue has received limited attention in past years, in spite of its importance, witnessed by the growing diffusion of multi-agent applications.
Similar issues are raised in [18] [19], which stress the importance of interoperability between heterogeneous expert systems and considers the problem of defining translation methods among different uncertainty representation approaches. However this work deals only with the uncertainty models used in the EMYCIN, PROSPECTOR and MYCIN systems and does not consider more general theories. Moreover it fails to introduce the notion of a common interchange format and therefore considers direct inter-formalism transformation, which is disadvantageous in many respects and in particular does not allow indirect communication.
Some of the translation methods we consider might be related to the approximation of non-additive measures by k-additive measures proposed in [9]: this will be explored in future work.
A related research direction concerns the definition of general formalisms able to include various existing theories as particular cases (see [10] [15] [17]). This is a complementary area which might provide suggestions for the definition of the interchange format, but leaves the translation problem open.
Proposing an uncertainty interchange format is made especially complex by the variety of existing uncertainty theories and by the differences existing among them both at semantic and syntactic level. The main questions related to this problem have been analyzed and a preliminary approach has been sketched. Thanks to its ability to include several well-known theories as special cases, imprecise probability theory has been identified as a suitable basis for our proposal. We then examined the issue of translation from imprecise probabilities into precise probabilities, belief functions, and possibilities, and defined the relevant transformation methods, which generalize previous similar proposals in the literature, where they exist. The proposed methods should be regarded as a first useful result: we are currently working on an extended version generalizing the assumptions of sect. 3.1.

Acknowledgements

We thank the referees for their helpful comments.

References

[1] M.S. Bazaraa, H.D. Sherali, C.M. Shetty, *Nonlinear Programming: Theory and Algorithms*, Wiley, 1993.

[2] B. de Finetti, *Theory of Probability*, Vol. 1, Wiley, London, 1974.

[3] D. Dubois and H. Prade, *Possibility Theory: An Approach to Computerized Processing of Uncertainty*, Plenum Press, New York, 1988.

[4] D. Dubois, H. Prade, S. Sandri, On possibility/probability transformations, in R. Lowen and M. Roubens (Eds.) *Fuzzy Logic – State of the art*, Kluwer, 1993, 103-112.

[5] F. Esteva, P. Garcia-Calvés, L. Godo, Relating and extending semantical approaches to possibilistic reasoning, *Int. J. Approx. Reason.*, 10, 1994, 311-344.

[6] T. Finin, Y. Labrou, J. Mayfield, KQML as an agent communication language, in J. Bradshaw (Ed.), *Software Agents*, MIT Press, 1995.

[7] FIPA 97 Specification, *Part 2: Agent Communication Language*, http://www.fipa.org

[8] M. R. Genesereth, Knowledge Interchange Format, draft proposed American National Standard (dpANS), NCITS.T2/98-004, http://logic.stanford.edu/kif

[9] M. Grabisch, Upper approximation of non-additive measures by k-additive measures - The case of belief functions, *Proc. ISIPTA 99*, Ghent, B, 1999, 158-164.

[10] C. Joslyn, L. Rocha, Towards a formal taxonomy of hybrid uncertainty representations, *Information Sciences*, 110 (3-4), 1998, 255-277.

[11] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ, 1976.

[12] P. Smets, Belief functions, in P. Smets, E. Mamdani, D. Dubois, H. Prade (Eds.), *Non-Standard Logics for Automated Reasoning*, Academic Press, London, UK, 1988, 253-277.

[13] P. Walley, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, 1991.

[14] P. Walley, Measures of uncertainty in expert systems, *Artificial Intelligence*, 83, 1996, 1-58.

[15] E. Weydert, General Belief Measures, *Proc. 10th Conf. Uncertainty in AI*, 1994, 575-582.

[16] M. Wolfenson, T.L. Fine, Bayes-like decision making with upper and lower probabilities, *JASA*, 77, 1982, 80-88.

[17] S.K.M. Wong, L.S. Wang, Y.Y. Yao, Interval Structure: A Framework for Representing Uncertain Information, *Proc. 8th Conf. Uncertainty in AI*, 1992, 336-343.

[18] C. Zhang, Cooperation under uncertainty in distributed expert systems, *Artificial Intelligence*, 56, 1992, 21-69.

[19] C. Zhang, Heterogeneous transformation of uncertainties of propositions among inexact reasoning models, *IEEE Trans. on Knowledge and Data Engineering*, 6 (3) , 1994, 353 –360.