Integrating Uncertainty Handling Formalisms in Distributed Artificial Intelligence

Simon Parsons[1,2] and Alessandro Saffiotti[3]

[1] Advanced Computation Laboratory, Imperial Cancer Research Fund, P.O. Box 123, Lincoln's Inn Fields, London WC2A 3PX, United Kingdom.
[2] Department of Electronic Engineering, Queen Mary and Westfield College, Mile End Road, London E1 4NS.
[3] IRIDIA, Université Libre de Bruxelles, 50 Av Rooseveldt, CP 194/6 B-1050 Bruxelles, Belgium

Abstract. In distributed artificial intelligence systems it is important that the constituent intelligent systems communicate. This may be a problem if the systems use different methods to represent uncertain information. This paper presents a method that enables systems that use different uncertainty handling formalisms to qualitatively integrate their uncertain information, and argues that this makes it possible for distributed intelligent systems to achieve tasks that would otherwise be beyond them.

1 Introduction

Distributed artificial intelligence (DAI) is that part of the field of artificial intelligence that deals with problem solving distributed amongst a number of intelligent systems. Thus DAI combines the power of artificial intelligence techniques with the advantages of distributed systems such as robustness and the ability to combine existing systems together in new configurations. One particular advantage of distributed artificial intelligence is the ability to take several existing systems and couple them together [2] perhaps with some new systems, into a community of agents that can tackle problems that are beyond the scope of any individual system. Thus a number of medical expert systems, each specialising in a particular area, can be used together to solve problems that are insoluble by any one system on its own [1].

When building distributed communities of agents, communication between the agents is very important. In order for agents to be able to communicate, they either need to use the same knowledge representation method, or have a means of translating between the methods used by different agents. This is especially true if the various agents are capable of handling uncertainty [11]. If different agents have different means of handling uncertainty, then one agent that, say, uses possibility theory to represent its uncertainty will not be able to understand the results of another agent that employs belief functions unless the agents have a means of translating from one formalism to another.

2 An Example of a Multi-agent System

As an example of the problems that different formalisms can pose consider the following hypothetical situation. We have a community of medical agents which between them have access to a body of medical knowledge similar to that of the Oxford System of Medicine [5]. This knowledge can be summarised in a network, a fragment of which is shown in Figure 1. This fragment encodes the medical information that joint trauma (T) leads to loose knee bodies (K), and that these and arthritis (A) cause pain (P). The incidence of arthritis is influenced by dislocation (D) of the joint in question and by the patient suffering from Sjorgen's syndrome (S). Sjorgen's syndrome affects the incidence of vasculitis (V), and vasculitis leads to vasculitic lesions (L). Consider further that none of the agents

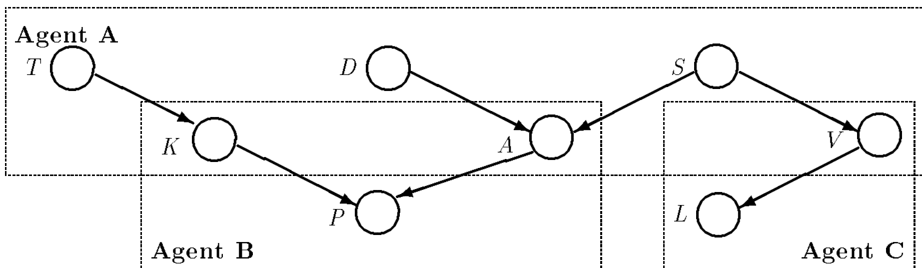

Fig. 1. Medical knowledge about arthritis and associated conditions

has knowledge of the complete network, since, like the modular systems of [1], each is a specialist in a narrow area. In Figure 1, agent A has knowledge of the associations between diseases and their causes; agent B is a specialist in pain; and agent C's main competence is with lesions. Notice that some nodes are shared between different agents: these constitute the communication channels between them. Suppose now that the agents use different methods for representing the uncertainty inherent in the medical knowledge. Thus the strengths of the influences known to agent A are given as probabilities, reflecting the fact that there is good statistical data relating the various complaints. Agent B represents strengths as possibilities [4], with values based on the physical possibility of one condition influencing another. Finally, the dependencies used by agent C are quantified by subjective values of belief strength, expressed using belief functions [9]. All of these numerical values are given in Table 1. The lack of conditional belief values other than those given is a result of ignorance about the incidence of pain under these circumstances. It is important to notice that the heterogeneity in the way uncertainty is represented across different agents is not an oddity of the system, but a consequence of the fact that uncertainty may be present in many different forms, and each form should be represented by the formalism that is most suitable for the job [8].

Given the information distributed among the agents, it should be possible for the community to reason from what is known about joint trauma and Sjor-

Table 1. Conditional values for the medical example

$p(k \mid t)$	$= 0.6$	$p(v \mid s)$	$= 0.1$
$p(k \mid \neg t)$	$= 0.2$	$p(v \mid \neg s)$	$= 0.3$
$p(a \mid d, s)$	$= 0.9$	$p(a \mid \neg d, s)$	$= 0.6$
$p(a \mid d, \neg s)$	$= 0.6$	$p(a \mid \neg d, \neg s)$	$= 0.4$
$bel(p \mid k, a)$	$= 0.9$	$bel(p \mid k, \neg a)$	$= 0.7$
$bel(p \mid \neg k, a)$	$= 0.7$	$bel(p \mid k \cup \neg k, a)$	$= 0.6$
$bel(p \mid k, a \cup \neg a)$	$= 0.7$	$bel(\neg p \mid \neg k, \neg a)$	$= 0.5$
$bel(\neg p \mid \neg k, a \cup \neg a)$	$= 0.4$		
$\Pi(l \mid v)$	$= 1$	$\Pi(l \mid \neg v)$	$= 1$
$\Pi(\neg l \mid v)$	$= 0.1$	$\Pi(\neg l \mid \neg v)$	$= 0.1$

gen's syndrome to establish something about the pain or vasculitic lesions. For instance, knowing that the patient has joint trauma should make the agents increase their belief in her suffering pain.

However, in order to perform deduction across the boundaries of their individual knowledge, the agents need to exchange information. Agent A can calculate the probability of loose knee bodies, arthritis and vasculitis, but nothing can be said about the belief in pain, or the possibility of vasculitic lesions unless this information is passed on to agents B and C. Moreover, unless the agents exchange information, nothing can be said about how these quantities will change when it is established, for instance, that the patient in question is suffering from Sjorgen's syndrome. Because all the information is represented using different formalisms, it is not possible for the multi-agent system to make these sort of cross-deductions in a direct way; in order for the overall system to perform reasoning we need some kind of integration between the different formalisms.

Two solutions are commonly employed in the DAI literature. The first one is to provide each agent with the ability to translate knowledge (and uncertainty) expressed in its own language into the language of any agent it needs to communicate with, and *vice-versa*. The complexity of this solution quickly becomes prohibitive as the number of agents grows. Moreover, translating from one formalism to another is not always feasible, and may introduce arbitrary assumptions (e.g., assumptions of equiprobability).

The second solution is to develop a common knowledge representation language, or *interlingua*, and require all inter-agent communications to use this interlingua [2]. Although this solution is in general less expensive than the previous one, its cost can still be very high, and developing an interlingua that is powerful enough to subsume each of the individual languages may be unfeasible. An alternative approach is to adopt an interlingua that is weak enough to be subsumed by all of the agents' languages. In the remainder of this paper we outline a technique of this kind that may be used to integrate different uncertainty representation formalisms, and discuss it in the context of the example given above. Space restrictions unfortunately limit the discussion, but more detail may be found in [6] and proofs of all results may be found in [7].

3 Integrating the Formalisms

Our approach to integrate different uncertainty handling formalisms is grounded on the notion of *degrading*: given a representation of uncertainty, we degrade its information content to a level that can be shared between all the different formalisms; this degraded information is then communicated between agents. We represent degraded uncertainty as qualitative changes: given knowledge of conditional probabilities, possibilities and beliefs relating a set of hypotheses, we focus on how the probabilities, possibilities and beliefs will change when we have new evidence. More precisely, we use the conditional values to establish the qualitative values of the derivatives that relate the probabilities, possibilities and beliefs, that is to establish whether the derivatives are positive [+], negative [−] or zero [0]. The derivatives tell us how changes in value move through a network. This method of integration is fully described elsewhere [6, 7]

Consider three variables A, B, and C related such that if $p(A)$ increases, $p(B)$ decreases, and if $bel(B)$ increases, $bel(C)$ increases. In other words, $\left[\frac{\partial p(B)}{\partial p(A)}\right] = [-]$ and $\left[\frac{\partial bel(C)}{\partial bel(B)}\right] = [+]$. These derivatives allow us to propagate changes in value from node to node so that given $\Delta p(A)$ we can establish $\Delta p(B)$, and if we know $\Delta bel(B)$ we can establish $\Delta bel(C)$. If we accept the following *monotonicity assumption*:

> If the value of a hypothesis in one formalism increases, the value of the same hypothesis in any other formalism does not decrease,

then we can use the qualitative changes to integrate different formalisms as follows. If we have $\Delta p(A) = [+]$ then $\Delta p(B) = [-]$. Now, from $\Delta p(B) = [-]$, we know that $\Delta bel(B) = [-]$ or $[0]$, and so we can establish that $bel(C) = [-]$ or $[0]$. Other translations may be carried out in exactly the same manner.

So to come back to our example, the agent that deals with probabilistic knowledge knows that $\Delta p(s) = [+]$, $\Delta p(t) = [0]$ and $\Delta p(d) = [0]$ since the only change that it knows about is that the patient is now observed to be suffering from Sjorgen's syndrome. Since a change of [0] can never become a change of [+] or [−] [6] it can ignore the latter changes, and using the fact that $p(x)+p(\neg x) = 1$ for all x it knows that $\Delta p(\neg s) = [-]$. Now, Theorems 3.1 and 5.1 from [7] establish the qualitative derivatives that link V and A to S from the values of Table 1. These results are that:

$$\left[\frac{\partial p(v)}{\partial p(s)}\right] = [-]\,, \left[\frac{\partial p(v)}{\partial p(\neg s)}\right] = [+]\,, \left[\frac{\partial p(a)}{\partial p(s)}\right] = [+]\,, \text{and} \left[\frac{\partial p(a)}{\partial p(\neg s)}\right] = [-]\,,$$

so that $\Delta p(a) = [+]$, and $\Delta p(v) = [-]$ from which it is possible to deduce that $\Delta p(\neg a) = [-]$ and $\Delta p(\neg v) = [+]$. These results may then be passed to the agents that deal with possibilities and beliefs. Using the monotonicity assumption these agents know that if the probability of a hypothesis increases then both the possibility of that hypothesis and the belief in it do not decrease. Similarly if the probability decreases then the possibility and belief do not increase. Thus the agent that handles beliefs knows that $\Delta bel(a) = [+]$ or $[0]$, $\Delta bel(\neg a) = [-]$ or $[0]$ while the agent that deals with possibilities knows that $\Delta\Pi(v) = [-]$ or $[0]$ and $\Delta\Pi(\neg v) = [+]$ or $[0]$. Now, Theorem 5.3 in [7], when applied to the values in Table 1. gives:

$$\left[\frac{\partial bel(p)}{\partial bel(a)}\right] = [+]\,, \left[\frac{\partial bel(p)}{\partial bel(\neg a)}\right] = [0]\,, \left[\frac{\partial bel(\neg p)}{\partial bel(a)}\right] = [-]\,, \text{and} \left[\frac{\partial bel(\neg p)}{\partial bel(\neg a)}\right] = [+]\,,$$

Thus the agent that deals with beliefs can tell that $\Delta bel(p) = [+]$ or $[0]$ and $\Delta bel(\neg p) = [-]$ or $[0]$. Since the agent that deals with possibility is initially ignorant about the possibility of vasculitis, we have $\Pi(v) = \Pi(\neg v) = 1$, and Theorem 3.2 in [7] gives:

$$\left[\frac{\partial\Pi(l)}{\partial\Pi(v)}\right] = [0]\,, \left[\frac{\partial\Pi(l)}{\partial\Pi(\neg v)}\right] = [0]\,, \left[\frac{\partial\Pi(\neg l)}{\partial\Pi(v)}\right] = [0]\,, \text{and} \left[\frac{\partial\Pi(\neg l)}{\partial\Pi(\neg v)}\right] = [0]\,,$$

with the result that $\Delta\Pi(v) = \Delta\Pi(\neg v) = [0]$. Thus the result of the new evidence about the patient's suffering from Sjorgen's syndrome is that belief in the patient's pain may increase, while the possibility of the patient having vasculitic lesions is unaffected. Notice that the achievement of this conclusion has required the integration of the knowledge available to different agents, and would otherwise have been beyond the system. As discussed in [6] it is also possible to integrate while performing diagnostic reasoning from observations about P and L to learn something about $p(t)$ and $p(s)$

4 Discussion

We have outlined a scheme for integrating uncertainty handling formalisms, in the setting of distributed artificial intelligence, as a means of enabling communication between intelligent agents. The integration is qualitative and is based upon the analysis of the relationships between variables in probability, possibility and evidence theories, and what we have called the monotonicity assumption. Two important issues arise from this integration: the validity of the monotonicity assumption, and the usefulness of the qualitative results.

There are several informal arguments that may be made in favour of the monotonicity assumption (as well as more formal ones — see [6]). Firstly, this assumption seems to be intuitively acceptable as a principle of coherence. Similar principles have been proposed for the relation between probability and possibility [10], and between subjective and objective probability [3]. Secondly, and perhaps more importantly, this seems to be the weakest assumption that allows some form of integration. In fact, relaxing this assumption would eliminate any constraint between values of belief in different formalism, and render any communication content free, and strengthening it would lead to the introduction of spurious information in certain cases.

As for the usefulness of the qualitative results, it is clear that the kind of qualitative integration introduced above is extremely weak, and will never produce as accurate results as a complete numerical integration because the method suffers from the usual problems of qualitative reasoning. The qualitative integration can only say that a value increases, decreases or does not change, and so cannot distinguish an increase of 0.01 from an increase of 0.99. Furthermore, the combination of an increase and decrease in value of a single variable due to two pieces of evidence will always give [?] indicating that it is impossible to say what the overall change is.

However, despite these disadvantages, qualitative integration is worthwhile. It provides information that could not otherwise be obtained, allowing co-operation between intelligent systems that would otherwise be impossible. Furthermore, the kind of qualitative indications provided by the system are easy to understand and so can be extremely helpful in assisting human decision makers, for instance in analyzing the possible consequences of a therapeutic treatment. Finally, it is possible to improve the precision of the results generated by the qualitative integration, by combining qualitative and quantitative information. In particular it is possible to use interval analysis to handle the bounds on the values of probability, possibility and belief values after integration [6].

Acknowledgements: the work described in this paper was partly supported by a grant from ESPRIT Basic Research Action 3085, DRUMS.

References

1. Agustí-Cullell, J., Esteva, F., García, P., Godó, L., Sierra, C. Combining Multiple-valued logics in modular expert systems, *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, (1991).
2. Avouris, N. M. and Gasser, L. (eds.) *Distributed Artificial Intelligence; Theory and Praxis*, Kluwer Academic Press, (1992).
3. De Finetti, B. Sul significato soggettivo della probabilità, *Fundamenta Mathematica* **17** (1931) 298–329.
4. Dubois, D. and Prade, H. Possibility Theory: An Approach to Computerized Processing of Uncertainty, Plenum Press, New York, (1988).
5. Glowinski, A. O'Neil, M., Fox, J. Design of a generic information system and its application to Primary Care, *Proceedings of the European Conference on Artificial Intelligence in Medicine*, (1989).
6. Parsons, S. Qualitative methods for reasoning under uncertainty, PhD. Thesis, Department of Electronic Engineering, Queen Mary and Westfield College, (1993)
7. Parsons, S. and Mamdani, E. H. On reasoning in networks with qualitative uncertainty, *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, (1993).
8. Saffiotti, A. An AI view of the treatment of uncertainty, *The Knowledge Engineering Review*, **2**(2) (1987) 75–97.
9. Shafer, G. *A mathematical theory of evidence*, Princeton University Press, (1976).
10. Zadeh, L. A. Fuzzy sets as the basis for a theory of possibility, *Fuzzy sets and systems*, **1** (1978) 1–28.
11. Zhang, C. Cooperation under uncertainty in distributed expert systems, *Artificial Intelligence*, **56** (1992) 21–69.

This article was processed using the LaTeX macro package with LLNCS style