

# Uncertainty in artificial intelligence

SIMON PARSONS

*Imperial Cancer Research Fund, P.O. Box 123, Lincoln's Inn Fields, London WC2A 3PX, UK*

## 1 History

The first conference on Uncertainty in Artificial Intelligence was held in 1985 by a group of people who felt that their views on the use of probability theory were not receiving a fair hearing from the rest of the AI community. At the time, mainstream opinion held that computational complexity of, and the amount of data required by, probabilistic methods made them inappropriate for realistic applications. As a result, those who claimed that probability theory was an adequate, if not the only adequate, method of handling uncertainty received a somewhat frosty reception.

Faced with this situation the proponents of probability theory decided to organize their own meeting, to be held in conjunction with AAAI, and it has been running ever since. The early years saw a deal of controversy and a number of heated debates since some of the attendees believed that there was more to the handling of uncertainty than probability. However, by the end of the 1980s the meetings calmed down, with the probabilists and their erstwhile opponents deciding to respect each others' points of view, although between sessions it is still possible to get into a heated argument about the pros and cons of probability and its competitors.

At least that is my feeling, as a non-probabilist. But eavesdropping on other conversations reveals a different point of view. It seems that the proponents of probability feel that probability theory is now universally accepted as the only real method for handling uncertainty since, thanks to much hard work over the last nine years, it is both computationally cheap for real problems, and expressive enough to handle all situations. Strangely, both viewpoints seem to coexist quite peacefully.

## 2 Introduction

The 1993 conference followed tradition in being located in the same city as AAAI, Washington DC, on the campus of the Catholic University of America. 125 papers were submitted to the conference, the bulk being from researchers in the US (since many Europeans prefer to send their work to the *Information Processing and Management of Uncertainty Conference* which is held in Europe every two years), although a fifth were from Europe. Around half of the papers were accepted, of which 29 papers were presented in sequential plenary sessions (one of the great strengths of UAI is the fact that all plenary papers are presented to all attendees) and 27 in two poster sessions. The conference also included panel sessions on Applications and Applications to AI Tasks. There were 151 attendees from 16 countries. This report describes some of the themes that emerged during the meeting.

## 3 Tutorials

The first morning of the conference was given over to tutorials, to facilitate the exchange of information by ensuring that everybody understands everybody else's formalisms. Didier Dubois discussed how fuzzy sets and possibility theory could be used to represent and reason with imprecise and vague information, thus laying out one of the main alternatives to the use of probability theory. Ross Shachter, on the other hand, provided a contrastingly orthodox (for this conference) presentation, an introduction to belief networks and influence diagrams, the two main representations adopted for handling uncertainty probabilistically. The final tutorial was on probabilistic methods for machine learning, and was presented by Padhraic Smyth and Wray Buntine. The latter was interesting because it was less concerned with explaining the technology than explaining how the technology could be applied to solving a given problem.

#### 4 Applications

One plenary session and one panel were given over to applications. The applications in the plenary session included a military application (ship classification) and an equally traditional medical diagnosis study. Less traditional were a paper on using bayesian methods to handle uncertainty in low-level vision (the distinguished student paper) and a paper on using probability to predict the structure of RNA. These discussed problems ranging from handling scaling issues to the representation of diagnoses changing over time. The panel addressed the wider issue of the role of applications in testing the theoretical methods proposed by the conference papers. There was general agreement that the need to build applications should drive the development of theories, and that theories should not be developed in order to form the basis of interesting publications. However, it was also argued that theory can drive applications, by making it possible to carry out previously impossible tasks.

#### 5 Theoretical advances

The remainder of the papers, and a second panel, were concerned with theoretical advances in the handling of uncertainty. This is the traditional preoccupation of the UAI conference and where most of the papers were concentrated. There were a number of papers discussing advances in probabilistic reasoning, including discussions of the meaning of causality and how it may be modelled, incremental computation of probabilities, and refining decision models. Other papers in a similar vein analysed the Noisy-Or model for approximating joint probability distributions, the diagnosis of multiple faults using a variety of probabilistic methods, and the combination of belief networks.

A number of papers were concerned with non-probabilistic methods, in particular possibilistic evidence, diagnosis based upon fuzzy relations, the assumptions behind Dempster's rule, conditional independence in valuation networks, and the generalization of Jeffrey's rule to belief functions. There was also considerable interest in non-numerical methods, with a session being devoted to qualitative methods, with contributions on argumentation and qualitative decision theory, as well as more technical issues such as intercausal reasoning.

After years of being frowned upon by other artificial intelligence researchers, workers in the area of uncertainty now feel that, in probabilistic networks and various other numerical techniques, they have tools with which they can represent and reason with uncertain knowledge correctly and efficiently. As a result, research interests are shifting to other matters, one of which is the use of qualitative abstractions which circumvent the need to elicit large amounts of numerical information. In addition there was a lot of interest (as there has been for the last few years) in building probabilistic networks from raw uncertain data, with two of the three distinguished student papers being on this subject, and there was a paper on performing sensitivity analyses to ensure that a probabilistic network does what it is supposed to. This is encouraging, as is the fact that other people seem to be interested in modelling uncertainty in the context of operations such as planning and scheduling that are close to mainstream AI. This suggests that, at last, the methods that have been under development for so long will soon be applied in all AI systems. Sound methods for handling uncertainty are needed by AI systems if they are to successfully interact with a noisy, dirty and imperfect but real world.