1 Introduction

Voices blurred soothingly in his ears. He cracked one puffy lid; everything indistinct, trimmed with foggy white; happy fungus growing everywhere. Ah, surfacing from anesthesia in Sick Bay . . .

Seth Morgan, *Homeboy*

The discipline of artificial intelligence has as its aim the construction of computer-based systems that exhibit intelligence. However, the exact meaning of “intelligence” in this context is somewhat unclear. When work on artificial intelligence first began, the standard definition of “intelligence” in a computer program was that given by Alan Turing. Turing’s proposal was that a computer program whose intelligence was to be gauged should be connected up to a terminal at which a person was sitting so that questions and answers could be exchanged between the two. If, after extensive questioning, the person was completely unsure whether or not the entity with which she was communicating was human or a computer program, then the program would be deemed to have intelligence. Sadly, despite the early optimism of workers in the field (McCorduck, 1979), it has become clear that no program is going to pass the so-called “Turing test” in the near future. Furthermore, the recent Leobner Prize Competitions, which offer cash rewards for programmers who are capable of building systems that can make a fair attempt at passing a restricted version of the Turing test (Russell and Norvig, 1995, page 5) seem to indicate, if anything, that systems which are good at fooling people into thinking that they are human are not those that show most intelligence. For instance, the winner of the first competition in 1991 (described in (Epstein, 1992)), engaged the judges in rather random whimsical conversation, while the winner of the 1994 competition (described in (Platt, 1995)) was a system that offered 380 stock pieces of advice about sex. Neither appear to be particularly smart, and neither comes close to being the kind of system popularly associated with the idea of computer intelligence, whether it is HAL from *2001: A Space Odyssey*, Wintermute (Gibson, 1993), or the *Arbitrary* (Banks, 1993). What characterises such systems is not the fact that they can hold conversations (though, of course, they do) but that they are capable of carrying out complex tasks like flying a spaceship, organising an armed robbery, or surveying a planet,[1] and it seems to me that it is this side of their cognitive ability—the ability to solve hard problems in a complex world—that marks them out as intelligent.

Now, by definition, solving such problems is a difficult matter, and the business of building a system to do so is far from trivial. There are difficulties, for instance, in

1 Though it could be argued that the *Arbitrary*’s fondness for collecting snowflakes or its ability to cheat amusingly when playing word games are even more telling signs of intelligence.

sensing the environment in which the task is to be carried out, and in planning how the task may be achieved. Even moving around is far from easy. However, one of the most fundamental problems is the handling of imperfect information. For those unfamiliar with the field, it is worth elaborating this point. A truely intelligent system needs to operate in the real world—it needs to be situated (Russell and Norvig, 1995, page 403). To interact successfully in the world the intelligent system needs some kind of representation of that world, a representation that may be explicit or implicit, and a means of reasoning with that representation. Now, when dealing with toy problems in a laboratory, all information about the world may be assumed to be present and correct, and providing a means of representing and reasoning with it is relatively simple. However, the real world is much more complicated. A moment's thought is sufficient to reveal the extent to which imperfect information is present in daily life and, as Morgan and Henrion (1990) point out,

> we have evolved cognitive heuristics and developed strategies, technologies and institutions such as weather reports, pocket-sized raincoats, and insurance to accommodate or compensate for the effects of uncertainty.

It is exactly because information about the world is imperfect, that intelligent computer systems operating in the real world have to be able to represent and reason with imperfect information. Of course, there are strategies for "engineering out" (Cohen, 1985) the imperfections in a given situation to reduce the situation to one of perfect knowledge. This is the kind of approach adopted by Clark et al. (1994) to reduce the imperfections in genetic map data to manageable proportions, and advocated by Castro and Trillas (1993) for dealing with inconsistent information. However, while this approach may be suitable in some situations, in others ignoring or attempting to smooth out the imperfections can severely degrade the performance of the system (Uhrik, 1982). Thus there is often a need to deal with imperfect information when building intelligent systems. This need rules out the use of classical logic as the basis for these systems (Israel, 1987), and has stimulated the development of a number of so-called non-standard methods. Indeed, over the past twenty years a large number of formal methods have been proposed for the management of uncertainty in artificial intelligence systems including evidence theory (Shafer, 1976; Smets and Kennes, 1994), possibility theory (Dubois and Prade, 1988f; Zadeh, 1978), probabilistic networks (Jensen, 1996; Pearl, 1988b), and a host of other approaches (Smets et al., 1988).

These methods, by and large, involve attaching a numerical measure to items of information whose certainty is not known. The measure represents the degree to which it is certain, suspected, or believed that the item is true, and as different items are combined to reach a conclusion, the associated measures are combined to derive the degree to which the conclusion is certain, suspected, or believed, to be true. The methods are often specified in a precise mathematical way. As a result it is possible to precisely determine the meaning of the derived measure. It is also possible to precisely determine the conditions under which the measure is valid. These methods also require specific types of information to be known, and place specific conditions on how the measures are to be allocated—for instance the use of one method might require the attachment of a measure to each of a set of facts while another might only require a measure to be attached to at least one. Such conditions on the measures lead to quite stringent constraints on the situations in which models may be built using the methods. This, in turn, means that it is often the case that, when modelling a specific situation, no method exactly fits the data that is available. This raises a question that is largely avoided by people carrying out research into reasoning with imperfect information, which is:

How does one choose which method to adopt for handling imperfect information?

There are some exceptions to this avoidance of the question, including Heckerman (1990a), Heckerman and Shwe (1993), and Saffiotti et al. (1994), but the more usual position is to assume that the different methods are completely exclusive and that only the best method is worth bothering with. As a result the question that has been extensively explored is:

Which is the best method?

and much time and energy has been expended in assessing their relative worths.

Simplifying greatly, the argument has progressed along the following lines. Initially the mainstream view within the community of people doing research into artificial intelligence was that, despite its mathematical pedigree, probability theory was an inappropriate method because it both required infeasibly large amounts of information and did not handle imperfect information in the same way that people did. This led to the development of a number of models of reasoning under uncertainty that were supposed to be more easily applicable and more able to model the way that people deal with imperfect information. Three such models are certainty factors (Shortliffe, 1976), fuzzy sets (Zadeh, 1965), and the theory of evidence (Shafer, 1976). This provoked a number of people who were interested in the use of probability theory to argue that not only was it possible to apply probability theory when the right models were used (Cheeseman, 1985), but that it also provided the only properly rational means of handling uncertainty (Horvitz et al., 1986). This cause was greatly helped by the development of probabilistic causal networks (Pearl, 1986b) which provided a simple and efficient way to build probabilistic models, but this has not prevented an intense and, at times, rather acrimonious debate attempting to determine the best system. Notable contributions have been made by Pearl (1985), Cheeseman (1986; 1988b), Zadeh (1986), and Smets (1988b).

The debate now seems to have tailed off, and it seems to me that this is because the proponents of the different approaches no longer care if the supporters of other approaches agree with them or not. Each of the main methods has its own set of adherents who talk to each other, go to the same conferences, and tolerate (but are not necessarily convinced by) the point of view of other groups.

In recent years, however, an eclectic school of thought has emerged with authors such as Fox (1986), Saffiotti (1987), and Krause and Clark (1993) espousing the idea that the various methods are largely complementary. The argument runs along the lines that the different methods are often designed to model different aspects of imperfect information, require different preconditions, and provide different types of solution. This suggests not only that no one method can ever be shown to be better than every other under every condition, since each will be best under the conditions for which it was designed to operate. Such a position is reinforced by the work of Léa Sombé[2] (1990) who attempted to express some simple pieces of imperfect information in a number of different formal models and found that none was able to represent them all, and seems to be becoming well accepted in certain circles even if it is not directly stated as such. For instance, it is possible to read Darwiche's (1993b) position—that one should examine probability to see which features are desirable and adopt these as the basis of new methods—as being supportive of the eclectic viewpoint. Taking this viewpoint, of course, shifts the emphasis from answering the second question back to answering the first. Many clues as to the answer can be gleaned from existing work, including that of Léa Sombé, Kruse and colleagues (Kruse et al., 1991), Krause and Clark (1993), and Hunter (1996), but this remains an area in which much work still needs to be done.

Although this work is of major importance if the eclectic position is to be of practical use, there are other, equally important, issues. It is these other issues that are addressed in this book. The first of these stems from a more extreme version of the eclectic position, which is developed at length in Chapter 5. Briefly the argument is as follows. If the eclectic position is valid, and I strongly believe that it is, then no single model for handling imperfect information can handle every type of imperfect information that will be encountered. Therefore, there will be scenarios in which no single model will be able to handle the imperfect information that is present. Thus the best possible treatment of imperfect information in those cases will be achieved by using several formalisms in combination, and the pertinent question is:

How may the different models be combined together in a principled way?

Answering this question was one of the main aims of the research that led to this book, and results suggesting how this might be done are given in later chapters and constitute

2 The name is a pseudonym for a group of French researchers interested in the handling of imperfect information, and is an approximation to the expression "Les As sont Bs" ("The As are Bs"), one of the expressions that they attempted to capture in their work.

one of the key parts of the book. The other aims of the research described here are best summarised by considering another question about models of handling imperfect information. Since the models often require a good deal of numerical information to quantify the imperfections in the data, what must be done when the available data is not rich enough to satisfy the axioms of any model? In other words:

What must be done when data expressed in a particular formalism has missing values?

Well, one answer to this question is to make up the relevant data using some suitable assumption about the missing values. Another, and to my mind more reasonable, solution is to take whatever information is available and see what can be deduced from it, usually sacrificing some of the precision of the original methods. The development of such an approach constitutes the other main part of this book. Now this kind of approach, which has widely been termed "qualitative" is becoming quite widespread and, by happy coincidence, also provides a means of answering the previous question. This then summarises the remainder of the book—the development of a number of methods (well, two or three depending on what is counted as a method) that answer the questions of how to handle missing values in models for reasoning with imperfect information, and how to enable different models to be used in combination. There is a good way to go before such models can be defined, and the next chapter starts by considering the different types of imperfect information that these models have to handle.