

## Preface

*He did swim in deep waters, on his belly, on his back, sideways with all his body, with his feet only, with one hand in the air.*

Rabelais, *Gargantua*

The problem of reasoning with imperfect information is widely recognised as being of great importance within the field of artificial intelligence, and a wide range of formal methods has been proposed for dealing with the problem. Many of these methods rely on the use of numerical information to handle the imperfections. This, of course, is quite a natural way to address the problem since to handle information about random events, for instance, it is useful to have numerical information about the frequency with which the events occur. However, taking quantitative information as a base for models that solve the problem of handling imperfect information merely creates another problem—how to provide the quantitative information. In some situations this second problem can be easily solved, in others it proves more resistant to solution, and in recent years several authors have looked at a particularly ingenious way of tackling it. The methods that they suggest, while drawing inspiration from quantitative methods, in particular probability theory, do not require the kind of complete numerical information required by quantitative methods. Instead they provide information that, though less precise than that provided by quantitative techniques, is often sufficient.

This book proposes a number of new techniques with similar aims. In particular, it suggests means of handling imperfect information in accordance with the rules of well-founded methods such as probability theory, possibility theory, and the Dempster-Shafer theory of evidence, while not requiring the kind of numerical information normally assumed by those applying them. This is achieved through the use of qualitative methods, by which I mean methods that abstract away from numerical information, sometimes just retaining the most critical distinctions such as those between positive and negative information, and sometimes by allowing values to range across intervals. This book also provides support for another recent development—the recognition that rather than searching for the best method for handling all imperfect information, one should instead use whichever method best fits the problem at hand. Such an eclectic position leads naturally to the use of several techniques in the solution of a single problem, and hence to the problem of combining the results of applying these techniques. This then is the other focus of this book, to provide a solution to the problem of integrating different formalisms, and it turns out that qualitative methods provide such a solution.

The structure of the book is as follows. Chapter 1 discusses the handling of imperfect information in the context of artificial intelligence, and introduces the particular aspects of the problem that are addressed by this book. Chapter 2 then provides a survey of the

literature of imperfect information, covering ground from economics via psychology to information theory. This is intended to give a flavour of the philosophical background to the problems of modelling imperfect information, and to provide support for some of the positions adopted later on. Chapter 3 builds upon this with a review of a selection of the formal techniques that have been developed for dealing with imperfect information, including the theories that form the basis of the qualitative methods developed in later chapters. The basic ideas behind each of these formalisms are discussed, and the essential axioms laid down. This is complemented by Chapter 4, which reviews previous work on qualitative approaches to reasoning with imperfect information, from naïve physics to models of argumentation, with particular emphasis on qualitative versions of quantitative formalisms. Chapter 5 then sets the scene for the more technical part of the book by describing the framework in which the work is carried out and providing examples that are used to illustrate the methods as they are developed. This framework borrows heavily from existing work on network representations in which models take the form of directed acyclic graphs. In such graphs nodes represent variables and arcs represent the dependencies between variables.

The technical work begins in Chapter 6, which introduces the idea of a qualitative algebra using an example from the literature, and then expands this and gives it a formal definition. Perceived limitations of this algebra lead to the definition of another that can handle interval values, and the use of both algebras to handle both integration and the problems caused by incomplete quantitative information about a situation is discussed. Chapter 7 presents a qualitative analysis of the use of probability, possibility, and evidence theories that provides a means of assessing the qualitative behaviour of networks from the conditional values controlling the interactions between nodes. The integration of formalisms and a solution to the problem of incomplete models using the results of the qualitative analysis is discussed. Chapter 8 extends the qualitative analysis to cover other classes of network and a number of different scenarios, including that of “explaining away.” Chapter 9 describes the system MummU (Multiple Uncertainty Management techniques), which implements the ideas of Chapters 6, 7, and 8. Chapter 10 deals with the application of the techniques to the problem of protein topology prediction, and Chapter 11 concludes. There are also two appendices and a glossary. The first appendix contains the proofs of all new results stated in the book, which I wanted to keep separate in order to avoid interrupting what narrative thread there is. I trust that I have been successful in doing this in such a way that it is possible to read the book and understand it without having to refer to the proofs, though the reader is encouraged to examine the proofs in order to really understand what is going on. The second appendix contains some calculations that I felt rather disrupted the text but which needed to be included somewhere. The glossary was included to give a brief description of a number of commonly used terms which it did not seem appropriate to give

in the main text, but which, if undefined, might prevent a full understanding. Most terms, including all those which are used at all extensively, are defined in the main text and will not be found in the glossary.

The focus on eclecticism in handling imperfect information and the need to integrate different approaches betrays the origin of the book. It started life as my PhD thesis, and the preparation of the thesis was funded by the European Commission as part of Esprit Basic Research Action 3085, Defeasible Reasoning and Uncertainty Management Systems (DRUMS). This project had as its aim the study of different formalisms for handling imperfect information and, in particular, their integration. However, the book is more than just a nicely bound edition of my thesis. In fact, the relation between my thesis and this book is rather similar to that between Boolos's *The Unprovability of Consistency* and *The Logic of Provability*, which he neatly summarised in the preface to the latter (Boolos, 1993):

“Is it a new book or [just] a second edition of your other book?” my colleagues have asked me.

New book.

All right, there are borderline cases, aren't there? (Hmm, maybe books don't exist since we lack strict criteria of book identity.) And cases over the border, but near it, too. That's what I think the situation is with *The Logic of Provability*: there's enough new material in it for it to be a new book. Of course, when I thought I couldn't do better than I had done in *The Unprovability of Consistency*, I unashamedly copied sections from it. But I think I may call this work the “successor” to that.

In others words, this book is a substantially revised version of my thesis.

There are a number of reasons for this. The first is that I was unhappy with some of the limitations that time had imposed upon the results in my thesis, and wanted to rework them to put the proposed methods on a stronger footing. The second is that the anonymous reviewers of the original manuscript made a number of suggestions that I wanted to incorporate. The third, and possibly the most important, is that the field has moved on a good deal in the time that it has taken me to put this material together so that there are now so many other qualitative and eclectic approaches that the whole shape of this work has changed. Of course, as with any such work, there is more that could be done—more results that could be established and more approaches that could be compared and contrasted with mine.

To paraphrase Umberto Eco (1998), the process of writing academic books is the following. You spend many years thinking, getting your ideas straight in your own head, and writing them down. You then find out what other people think of your ideas. Sometimes you find you were wrong, and sometimes what other people think about your ideas leads you to change your mind. You then revisit what you wrote originally, making corrections and alterations, some of which contradict what you said earlier. This process then repeats. This book marks the second cycle of this process with respect to the work it contains. It

has changed hugely since the first iteration, and doubtless it will change in the future, but for now it is the best I can do.

This seems an appropriate place to make a few remarks about what this book is and is not. In particular, I feel it is worth stressing that this is a research monograph rather than a textbook, so that the various methods considered within and the views expressed about them more reflect my personal opinions than the measured views of the field as a whole. The fact that this is a monograph does not mean that it turns its back upon all the other work on reasoning under uncertainty that already exists. Indeed, it contains a rather large amount of introductory and survey material (much of which was included at the suggestion of the reviewers) so that not only does it introduce many of the ideas that it uses by discussing them in the context of the work in which they were developed, but in doing so covers a large amount of material that is never used in the theoretical parts. The reason for doing this is to give a sound introduction to the state of the art in handling uncertainty at the time of writing. Actually, to be frank, it is really a snapshot of the field since, like snapshots of steppes, mountain ranges, and other landscapes with wide horizons, it only covers a part of the whole. As I sit here writing (in early 2001) the field seems sufficiently large that it can only be reasonably surveyed by a lengthy book and so, as it will no doubt continue to expand, I do not attempt to be comprehensive. Instead my aim is to focus on those developments that are most relevant to the work in the rest of the book. As a result the survey I give is rather uneven, covering some areas in a good deal of detail, but skimming over others. In particular, the detail is more often about older work since this is more foundational, and less often about more recent work, since the latter is frequently on particular topics (such as learning probabilistic networks) which are rather out of the scope of this book. With respect to this issue it should also be noted that throughout the book little space is given to discussion of decision making. While I acknowledge that making decisions is an important activity, I think of it, perhaps wrongly, as a topic that is distinct from reasoning under uncertainty and, certainly correctly, as a topic on which I currently have little to say.

As a research monograph this book has a two main aims. First, but not foremost, it is intended as a manifesto for the use of qualitative approaches to reasoning with imperfect information, which argues that they are useful and well-developed. It does this partly by highlighting the various qualitative systems that have already been developed, and partly by developing new ones and showing how they may be used. Second, this book is intended as a manifesto for the eclectic approach. As such it both argues that the eclectic approach is a natural way of modelling the various different forms of imperfection that may be present in models of the world when these may not be easily captured in a single model, and provides some practical assistance for those who wish to build eclectic systems by furnishing a theoretical means of combining information expressed in different formalisms. It should be

noted that these two aims are, essentially, independent since qualitative methods are useful in building non-eclectic systems, and eclectic systems can be built (at least in principle) without resorting to qualitative methods.

Finally, a few words on the vexed question of personal pronouns. The use of the word “his” to denote an individual who may either be male or female will no doubt continue to be controversial for years to come. The use of “his/her” may well be correct, but, like Dawkins (1988) I find this too ugly to bear. However, unlike Dawkins I tend to use “her” in those cases where it is impossible to find a means of phrasing things to avoid those pesky pronouns. Like Hofstadter (1986) I feel that it is better to err on the side of the under-used alternative rather than the over-used one. Of course this choice is still going to make some people unhappy because they feel that it is silly, patronising, hypocritical, or just incorrect. Well, that’s tough. Of course, if enough people object then maybe I’ll change things next time I write a book. Which leads, naturally enough, to the usual statement that all comments on any aspect of this book, whether brickbats or accolades, are most welcome—comments via email especially so. The relevant address, for the foreseeable future, is `s.d.parsons@csc.liv.ac.uk`.