6 Reference

[1] Z. An, D.A. Bell, and J.G. Hughes. Res: A relative method for evidential reasoning. In D. Dubois et al., editor, 8th Conference on Uncertainty in Artificial Intelligence, pages 1-8. Morgan Kaufmann, San Mateo, California, 1992.

[2] P.R. Cohen, D. Day, J.D. Lisio, M. Greenberg, R. Kjeldsen, D. Suthers, and P. Berman. Management of uncertainty in medicine. Int. J. of Approximate Reasoning, 1:103-116, 1987.

[3] J. Fraser, and I. Harrison. Modelling Expertise Using Belief Networks, in Bramer A and Macintosh A L (eds), Research and Development of Expert Systems X, Proceedings of Expert Systems 93, Robinson College, Cambridge, UK, December 1993. Also available as AIAI-TR-134, October 1993.

[4] S.L. Lauritzen and D.J. Spiegelhalter. Local computations with probabilities on graphical structures and their applications to expert systems. In G. Shafer and J. Pearl, editors, Reading in Uncertain Reasoning, pages 415-448. Morgan Kaufmann, San Mateo, California, 1990.

[5] T.Y. Leong. Representation requirements for supporting decision model formulation. In B.D. D'Ambrosio et al., editor, 7th Conference on Uncertainty in Artificial Intelligence, pages 212-219. Morgan Kaufmann, San Mateo, California, 1991.

[6] J. Pearl. Fusion, propagation and structuring in belief networks. Artificial Intelligence, 29(3):241-288, 1986.

[7] R. D. Shachter, B. D'Ambrosio, and B. A. Del Favero. Symbolic probabilistic inference in belief networks. In AAAI-90, volume 1, pages 126-131, 1990

[8] G. Shafer. Perspectives on the theory and practice of belief functions. Int. J. of Approximate Reasoning, 4:323-362, 1990.

Enterprise

John Fraser

1 Introduction

The Enterprise project is about halfway through its 3-year life. Along with our partners, IBM, Lloyd's Register, Logica and Unilever, we have produced a state of the art survey, built two demonstrator programs and built an enterprise ontology. We are in the process of defining exemplar enterprise modelling scenarios and designing a tool set to support these and other scenarios. Later this year we will be implementing the tool set and testing it against our exemplar scenarios. Next year we will produce a further, freely available scenario, which will be used to demonstrate the full power of the Enterprise tool set.

2 The Context: Management of Change

You may recall that Enterprise is the UK Department of Trade and Industry's major initiative to promote the use of knowledge-based systems in enterprise modelling. The project, part of the DTI Intelligent Systems Integration Programme (ISIP), is budgeted at over £2.6 million and is scheduled to run from August 1993 to October 1996.

The focus for our interest in enterprise modelling is the management of change. Understanding and visualising complex businesses enables one to identify and address areas that might be constraining business performance. Enterprise modelling helps one focus on those areas which can be changed, how these areas are currently functioning, how they might be optimised and how any changes might impact other areas.

3 The Issues: Generic Models, Reuse and Integration

Generic models are those that are not built for a specific purpose: the implication is that they can be used for different purposes at different times. There are two ways in which the Enterprise project is building re-usable models.

First, we recognise the power of having information structures which people can share: for communication, for consistency and for understanding. Second, we recognise the power of having shared methods for tackling tasks. Together, the shared information structures and shared methods constitute a common view or perspective of a particular subject or "domain". The development of shared mental models can be enhanced by such a common view. We have been devoting much of our efforts recently to the derivation of the Enterprise ontology, which we hope provides a common view of the information and activities involved in enterprise modelling.

The Enterprise ontology helps provide a gel for integrating what is currently a disparate set of modelling techniques and tools. This gel is "semantic": it helps clarify the meaning of the terms used. The existence of a gel with clearly specified properties also make it easier to design new techniques and tools with the intention of integrating them easily with existing ones.

Integration of computer tools also requires an infrastructure for communicating messages between tools - ours is based on an agent model - and a means of translating a user's needs into demands on the tools - ours is supplied by what we call our task management.

4 The State-of-the-Art Survey

We have evaluated current and emerging enterprise modelling methods and technologies, and documented them in a State of the Art Survey, available from AIAI at a cost of £25.

The five-part survey covers the following:

Part 1 - A Business Perspective on Enterprise Modelling

Part 2 - Contexts in which Enterprise Modelling is Used

Part 3 - Enterprise Modelling Methods

Part 4 - Reuse Applied to Enterprise Modelling

Part 5 - Technologies Supporting Enterprise Modelling

The Survey has been described by AI Watch, the Newsletter of AI Intelligence as:

> "... an excellent overview of many of the major advanced technologies within IT, providing definitions, historical perspectives and brief explanation. They are also placed in context, pointers are often given to relevant commercial suppliers, and a list of further, more detailed, reading material is always provided. We would strongly recommend anyone with an interest in advanced technologies getting hold of a copy of this Survey - even if they have no interest in enterprise modelling."

5 The *i*think® and Demos™ Demonstrators

We have taken two off-the-shelf software packages and used them to demonstrate aspects of enterprise modelling which might well be incorporated within a tool set.

***i*think®** is a tool for business modelling, based on the "accumulator and flow" paradigm of systems dynamics. Its strength lies in its availability to simulate complex relationships, particularly those incorporating feedback. We used it to model a generic business supply chain, based on IBM's Retail Application Architecture Enterprise Model.

Demos™ is a tool for decision modelling, based on the influence diagram paradigm of decision analysis. Its strengths are similar to those of *i*think®. We used it to model the building and evaluation of a portfolio of projects, taking into account factors such as the match with company strategy and the degree of balance of risk within the portfolio. HyperCard was used to provide non-experts with an easy interface to the Demos™ decision model.

Information sheets are available from AIAI for both the demonstrators.

6 The Enterprise Ontology

We have explored five areas pertinent to enterprise modelling (activity, organisation, strategy, marketing and time) and, after discussion, we have "bought in" to the definition of over 100 terms. We acknowledge the time, expertise and material of various other people and bodies who have contributed, directly or indirectly, to the development of the ontology, which we believe is one of the first attempts of its kind anywhere in the world. These contributors include:

TOVE -	Toronto Virtual Enterprise project, University of Toronto
ARPI -	ARPA/Rome Laboratory Planning and Scheduling Initiative
ORDIT -	Esprit project No. 2301
CYC -	The common sense knowledge base project at MCC and Cycorp
KSE -	The ARPA Knowledge Sharing Effort
WfMC -	The Workflow Management Coalition
PIF -	The Process Handbook and the Process Interchange Format at MIT and the University of Hawaii

The first version is paper based and we are now working on encoding it in KIF, the Knowledge Interchange Format developed under the Knowledge Sharing Effort (see the article in this issue of *airing)*. We know that the ontology will evolve as the project proceeds. We envisage the Enterprise tool set providing facilities to maintain and expand the ontology.

7 The Enterprise Application Scenarios

IBM, Lloyd's Register and Unilever are contributing to the project as tool set users, and each is focusing on a particular enterprise modelling scenario of relevance to themselves. These scenarios are currently being described in process and object notations in, as far as possible, the terms of the ontology. The scenarios will then be used as testbeds for both the ontology and for the Enterprise tool set when it has been implemented.

8 The Enterprise Tool Set

We have mentioned the tool set several times already. It will support various roles of user, from ontology maintainer, through method modeller (who describes generic enterprise modelling scenarios) and tool administrator (who defines which tools are applicable at certain stages in the scenario), to business analysts (who are guided in the use of the common methods as they follow the scenario).

We are using available technology wherever possible, for instance the Hardy tool and wxWindows library which have been developed here at AIAI. The CommonKADS method for the development of knowledge based systems has been used for analysis and design, and the STUDIO method is being referred to for user interface design.

9 The Final Demonstrator

Towards the end of the project we will be using the complete Enterprise tool set to define and build a general demonstrator which will embody, as far as possible, common features of typical enterprise modelling scenarios. We welcome any suggestions which you may think would interest and benefit a wide audience.

10 More Information

You can get more information from:

David Parry, Project Manager
(D.Parry@ed.ac.uk)

John Fraser, Technical Project Manager
(J.Fraser@ed.ac.uk)

or by viewing our WWW pages at URL

http://www.aiai.ed.ac.uk/~entprise/enterprise/

The Enterprise State of the Art Survey can be obtained from The Publications Secretary (AIAI@ed.ac.uk).

UK Planning and Scheduling SIG Work Plan

Howard Beck

1 Introduction

For some years, the AI Planning and Scheduling community in the UK has met through the Planning and Scheduling Special Interest Group (SIG) to discuss technical topics and to hear about application areas and product information. With the rapidly increasing maturity of planning and scheduling technologies, it was felt that it would be timely to expand the activities of the SIG to increase general awareness of these developments especially in the commercial sector. A proposal to this effect was accepted by the Department of Trade and Industry, and a two-year initiative is now being supported through their Intelligent Systems Integration Programme. This article describes the current activities on this initiative.

2 Background

The UK Planning and Scheduling SIG was first formed in June 1984 during the UK Alvey Programme. The interest was such that the SIG continued beyond Alvey and has enjoyed continuing support from its industrial and academic membership. At its 12th Annual Meeting, prior to receiving support from the DTI, the SIG decided to extend its activities with the aim of increasing awareness and interest in planning and scheduling in the UK. To this end, a widely-distributed newsletter would be produced.

3. Update on Activities Under the DTI-Supported Initiative

The 13th Annual Meeting of the UK Planning and Scheduling Group was held at the University of Strathclyde on the 14th and 15th September 1994. The meeting was very well attended and attracted a number of papers and visitors from Europe. Dr Patrick Prosser of the University of Strathclyde was Programme Chairman for the event, and at the annual meeting Dr Tim Grant took over as SIG Chairman from Prof Austin Tate who had held the chairmanship for more than 10 years. The papers for the SIG meeting have been printed as a set of proceedings which are available for £10 from the University of Strathclyde's Department of Computer Science.

The first of a series of Planning and Scheduling newsletters has been produced and distributed to more than 4000 people. The newsletter provides case studies of planning and scheduling applications across a range of sectors employing a variety of planning and scheduling techniques. In each of the newsletters, two case studies are included, as well as current news, events and conferences. In the first issue, two systems which are in regular use are described:

- a knowledge-based production scheduler development at Digital Equipment Corporation;
- a Genetic Algorithms based exam timetabler developed in the Department of Artificial Intelligence at the University of Edinburgh.

Another of the activities which have been undertaken as part of the ISIP initiative is the development of an online information resource on planning and scheduling covering a variety of topics: R&D groups, bibliographies, ontologies, benchmarks, products, applications and technologies. This resource, which is being extended as new relevant web pages become known, is available through the World Wide Web (URL *http: //www.aiai.ed.ac.uk/~pasg/isip/pas-sig.html*)

Over the period of the last year, there have been expressions of interest from a number of people not previously involved with the SIG. The mailing list for the SIG has been updated and extended by Dr Barbara Smith of the University of Leeds: the total number of people on the mailing list is now about 200. If any of the readers of *airing* are interested in being added to the Planning and Scheduling SIG mailing list, please contact Howard Beck at AIAI.

BOOK REVIEW: Automated Knowledge Acquisition

Publisher: Prentice Hall

Authors: Sabrina Sestito and Tharam S. Dillon

John Kingston

This textbook looks at automated approaches to knowledge acquisition, also known as machine learning approaches. Machine learning approaches include neural networks, rule induction, case based reasoning and (arguably) genetic algorithms. The book gives a fairly good introduction to induction algorithms, and learning using neural networks; indeed, four chapters are devoted to a learning technique based on multi-layered neural networks known as BRAINNE. Genetic algorithms are also discussed briefly, though helpfully; case based reasoning is not discussed. The last chapter gives a short introduction to the KADS methodology, and a token description of some manual knowledge acquisition techniques.

Any description of neural network techniques involves considerable discussion of mathematics; however, these authors have succeeded in keeping this book at a level which is fairly easy to read and also useful as a reference book. This is an impressive achievement for a book which is probably heavily based on the first author's recent Ph.D thesis. However, the book also suffers from the focused approach of a typical thesis; it concentrates on a particular technique which appears to be home-grown rather than widely available, and it lacks information on certain topics (notably case-based reasoning). The book is recommended for anyone interested in developing neural networks and/or induction systems; it should also be consulted by anyone with an interest in machine learning, or in knowledge acquisition in general.

BOOK REVIEW: Natural Language Processing for Prolog Programmers

Publisher: Prentice Hall

Author: Michael A. Covington

Nicolas Nicolov, Department of AI

This textbook aims to introduce natural language processing to students with working knowledge of Prolog. Language processing (on a computer and as done by humans) involves manipulating certain structures. This book focuses on the computer representation of these structures and their processing.

Chapter 1 sets the scene. It defines the interdisciplinary field of Natural Language Processing (NLP) as an area of applied linguistics. It is a concise introduction to what natural language is all about, set in terms that are easily accessible to non-linguists.

Chapter 2 introduces the need for NLP from a practical perspective. Two small applications are developed:

- understanding operating system commands;
- a database query system.

The first application uses template matching while the second application relies on keyword analysis which is a slightly more sophisticated method and is used in some commercial products.

Chapter 3 examines the standard Prolog approach to processing natural language: the formalism of Definite Clause Grammars. Many linguistic concepts are gradually introduced and the established Prolog approaches to these phenomena are reviewed. The chapter concludes by looking at an older parsing technique - that of augmented

transition networks.

If any of the previous chapters have led the reader to believe that NLP is a trivial matter, Chapter 4 on English phrase structure shows that syntax is indeed a complex issue.

Chapter 5 introduces unification-based grammar formalisms which are extensively used in modern systems. The chapter starts slightly misleadingly as the motivation for developing unification-based grammars is not immediately highlighted, but quickly improves with a good discussion of the building of a syntactic tree in a bottom-up fashion, and of an implementation technique whereby feature structures are compiled into representations which subsequently allow Prolog term unification to be used to model feature structure unification.

Chapter 6 discusses and compares well-established parsing algorithms: shift reduce, left-corner, chart parsing, and Earley's algorithm. The abstract mechanism of the parsing techniques is well-articulated and presented in a way that facilitates the transition to the implementational level which is also described. It would seem that there are didactic advantages in such a separation of levels.

Chapters 7 and 8 dive into the sea of natural language semantics. Chapter 7 examines how English input can be processed in order to answer database queries. The techniques are somewhat *ad hoc*, but there is a strong emphasis on the underlying logical theory. Chapter 8 gives a sketch of language translation, word-sense disambiguation and understanding of events. The book is restricted here by the immaturity of the field of computational semantics.

Chapter 9 returns to the first task which is necessary for text understanding - morphological analysis. While English morphology is less complicated than that of many other languages the author acknowledges that other languages place heavier emphasis on morphology, to determine syntactic and even semantic relationships. A popular approach to morphological analysis (two-level morphology) is discussed.

Appendix A reviews some topics in Prolog which tend not to be fully understood by students that have just taken a course in Prolog, which makes the book of potential interest to linguistics students who have taken an introductory Prolog course. Appendix B deals with the question of how text input can be made more user-friendly within Prolog.

The book as a whole covers a wider range of topics than is usual for books on NLP in Prolog. The book would have been even more useful if the author had devoted some space to natural language generation, especially since it is either not dealt with at all in most current textbooks or given very short shrift.

It's worth mentioning that a lot of effort has been put in making sure that the code used as examples in the book runs under a wide a range of Prolog platforms. This continues the tradition set in "Prolog Programming in Depth" by the same author.

This book is not meant to be an original contribution to the field, but as a textbook it is very valuable and recommendable for use in computational linguistics courses at the undergraduate or even graduate level - it is

meant to serve as the basis for a one semester course. The structure of the text makes it good reference material for anyone with an interest in natural language processing techniques.