ADVANCES IN ARTIFICIAL INTELLIGENCE

edited by
TIM O'SHEA

ADVANCES IN ARTIFICIAL INTELLIGENCE

Proceedings of the Sixth European Conference on
Artificial Intelligence, ECAI-84
Pisa, Italy, September 5–7, 1984

Edited by:

Tim O'SHEA
The Open University
Milton Keynes
United Kingdom

1985

NORTH-HOLLAND
AMSTERDAM • NEW YORK • OXFORD

for the European Coordinating Committee for Artificial Intelligence

ISBN: 0 444 87611 1

Published by:
ELSEVIER SCIENCE PUBLISHERS B.V.
P.O. Box 1991
1000 BZ Amsterdam
The Netherlands

Sole distributors for the U.S.A. and Canada:
ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
52 Vanderbilt Avenue
New York, N.Y. 10017
U.S.A.

PRINTED IN THE NETHERLANDS

ADVANCES IN ARTIFICIAL INTELLIGENCE

Sixth European Conference on
Artificial Intelligence, ECAI-84
Pisa, Italy, September 5-7, 1984

Organised under the auspices of
The European Coordinating Committee for
Artificial Intelligence

Sponsored by
AICA and AISB

NORTH-HOLLAND
AMSTERDAM • NEW YORK • OXFORD

PREFACE

This book is an edited version of the Proceedings of the Sixth European Conference on Artificial Intelligence held in Pisa in September, 1984. The conference was organised under the auspices of the European Co-ordinating Committee for Artificial Intelligence and sponsored by AICA and AISB.

251 papers were submitted for consideration in 15 different subfields. Each subfield was chaired by an experienced and respected worker in A.I. who received reports on each submission from two referees with appropriate publication records. Papers on which there was not an initial consensus were sent to additional members of the Programme Committee. This Committee was nominated by the various European A.I. societies. The paper selection process depended on considerable work from the Subfield Chairmen, Referees and Programme Committee.

For the purposes of this volume, I have made a selection from the prize-winning papers and the long papers in the Proceedings volume and have organised them into five main general areas, namely, Expert Systems, Robotics and Vision, Cognitive Modelling and Learning, Natural Language and Knowledge Representation. I regret that limitations of space made it impossible to include more of the papers presented at the conference. This selection consists of about 15% of the papers originally submitted.

On the acknowledgements page I have listed all those who helped in the preparation of the conference programme and proceedings. I would like to single out two people, the Programme Secretary, Diane Mason for her patient and heroic contributions to the conference, the proceedings and this edited book, and the General Chairman, Stefano Cerri who ensured that the whole event was carried out with tremendous flair, efficiency and style.

Tim O'Shea
Programme Chairman,
SIXTH European Artificial Intelligence Conference.

ACKNOWLEDGEMENTS

Subfield Chairmen

L. Aiello (Rome) - Theorem Proving
G. Albers (Karlsruhe) - Industrial Applications
P. Ambler (Edinburgh) - Robotics
W. Bibel (Munich) - Automatic Programming
M. Boden (Brighton) - Philosophical Implications
J. Breuker (Amsterdam) - Cognitive Modelling
H. Coelho (Lisbon) - Knowledge Representation
P. Greussay (Paris) - System Support
G. Hagert (Uppsala) - Learning
G. Koch (Copenhagen) - Logic Programming
J. Laubsch (Stuttgart) - Lisp Systems
C. Mellish (Brighton) - Natural Language
B. Neumann (Hamburg) - Vision
M. Somalvico (Milan) - Planning and Search
L. Steels (Brussels) - Expert Systems

Other Members of Programme Committee

B. Bara (Milan), J. de Bruin (Amsterdam), I. Futo (Budapest), W. Horn (Vienna) J-P. Jounnaud (Nancy), D. Kayser (Paris), R. Mantaras (Barcelona), R. Mohr (Paris), S. Ohlsson (Uppsala), G. Prini (Pisa), P. Ross (Edinburgh), P. Szerdi (Budapest), H. Trost (Vienna), K. Vershinin (Kiev), W. Wahlster (Saarbruecken), Y. Wilks (Clacton).

Referees

M. Allard, H. Alshawi, P.B. Andersen, G. Attardi, D.P. Barnes, M. Bidoit, B. Boguraev, M. Borillo, M.A. Bramer, I. Bratko, M. Bruynooghe, B. Buchberger, A. Bundy, S. Busemann, J. Campbell, J. Castaing, R. Chatila, E. Chouraqui, M. Cordier, E.J.F. Costa, D. Davies, J. Doran, W. Dzida, E. Eder, M. Eisenstadt, H. Farreny, A. Feldmann, L. Ferinas del Cerro, R. Fisher, W. Frey, L. Fribourg, M. Gehrke, M. Ghallab, A. Giordana, G. Goerz, P.M.D. Gray, H. Grumbach, G. Guida, J. Hallam, N.W. Hardy, T.A. Harley, J-P. Haton, U. Hein, E. Hollnagel, J. Howe, H. Jappinen, H. Kaindl, Y. Kodratoff, L. Konst, K. Koskenniemi, P. Laface, M. Latombe, J-P. Laurent, M. Lee, A. Lux, R.L. de Mantaras, H. Marburger, A. Martelli, B. Mayoh, L. Mero, K. Morik, J. Mudler, F. Nef. R. Nossum, H-J.Novak, L. Ny, J.J. Ohlbach, H. Pain, G. Papon, J. Pitrat, G.S. Pospelow, A. Ramsay, U. Reyle, C. Rohrer, M. Rosner, A. Schubert, J. Siekmann, B. Silver, F. Sirovich, A. Sloman, I. Steinacker, H. Tropf, H. Trost, J. Virbel, M. Yazdani, R. Young, G.P. Zarri, Z. Zdrahal.

Logistic Support and Practical Help

Olwyn Wilson, Coral Long, Frances Thomas, Carole Fulcher, Debbie Skeats, Kim Curtis, Christine Love, Jonathan Davies, Tim Rajan, John Domingue, Tony Hasemer, Rick Evertsz, Joanne Mason, Neil Mason, Eileen Scanlon, Alistair Edwards, Stephanie Smit and John Butterfield.

CONTENTS

ORIGINAL PAGE IS
OF POOR QUALITY

EXPERT SYSTEMS

Advances in Artificial Intelligence, T. O'Shea (Ed.)
Elsevier Science Publishers B.V. (North-Holland)

Interpretation of verbal data for knowledge acquisition

B.J. Wielinga and J.A. Breuker

Department of Social Science Informatics
University of Amsterdam, Weesperplein 8
1018 XA Amsterdam, The Netherlands

Abstract

Knowledge acquisition for knowledge based systems is a difficult task. In particular the mapping of verbal data obtained in interviews with an expert to knowledge structures for an expert system, is tedious and prone to errors. In this paper we investigate the problems occurring in this interpretation process, and we show how abstract models of problem solving processes can guide the interpretation. An outline of a methodology for knowledge acquisition based on such models, is presented.

1. The Nature of Verbal Data

A major problem in the development of knowledge based ('expert') systems is the acquisition of knowledge from domain experts. The prevailing method for such knowledge acquisition is the interpretation of verbal data, usually obtained during interviews with the expert, in terms of a formalism suitable for implementation (e.g. production rules). Although verbal data seem to be the most convenient source of information for knowledge acquisition -due to their richness, expressiveness and the natural way in which they are used to communicate knowledge in general-, there are a number of serious problem with the elicitation, interpretation and quality assessment of verbal data.

There is a number of methods for the elicitation of verbal data, ranging from rather open interviews to self report data obtained in highly controlled experimental situations. We have found 5 basic methods (Breuker & Wielinga, 1983). In the traditional interview a number of topics is addressed (focussed interview), or a number of concepts is explicated by deep probing (structured interview). Introspection refers to a situation in which the expert gives an account of how he would solve an imaginary, but typical case. In self report the expert produces an on-line thinking aloud protocol while solving a real problem. Such problem solving can be performed in interaction with a user (via teletypes), thus simulating interactions of the prospective expert system: user dialogues. Finally, the expert may be asked to review protocols obtained earlier. Within each method, various strategies may be employed. In table 1 the basic methods with some strategies, are presented. The kind of data that can be obtained by each method, differ widely, as is summarised in this table as well.

Although verbal data are in principle an ideal source for knowledge acquisition, in practice their interpretation is often problematic. It is well known that verbal data can be interpreted in a variety of ways, depending on the viewpoints of the speaker and listener, the assumed background knowledge and possible social effects. Besides the fact that

1) This research was supported in part by the ESPRIT programme of the Commission of the European Community, by contract 3.1/12.

verbal data are hard to interpret in a consistent way, these data are always incomplete (Ericsson & Simon, 1980; Breuker, 1981).

method/strategy	**data on**
focussed interview	factual knowledge
probing	types of problems
critical incident	functions of expertise
reclassification	environment (objects, agents)
	user characteristics
structured interview	structure of concepts
socratic dialogue	(part of) mental model
20 questions	reasoning/explanation
introspection	global strategies
hypothetical case	justification
forward scenario	evaluation
self report	use of knowledge sources
secondary task	heuristics
selective report	reasoning strategies
user dialogues	user-expert interaction
real life	problem 'negotiation'
via teletype	
review	repair of gaps in data
of data	interpretation of data
of prototype	justification

Table 1. Knowledge elicitation methods and nature of data.

Some of the reasons for incompleteness are:

* Omissions. In recalling some cases, or cases in general the expert may forget to mention many essential features or special conditions.
* The knowledge states may be hard to express in language, because they are very rich, or require drastic transformations (e.g. the description of a scene).
* Many knowledge states are not accessible for inspection by the minds eye; the knowledge is 'compiled'.
* In language use, much information is communicated by simple reference to knowledge that is assumed to be known by the receiver (pragmatics). The receiver has to account for such 'gaps'. In an interview, an expert may not further elaborate some issue, assuming that the rest is known, but often the interviewer has no means to evaluate whether more is involved.
* Experts may not be motivated to reveal their inner thoughts. There are many reasons to believe that an expert is a priori uncooperative.
* Most experts have little or no experience in giving report of their thinking. Particularly, presenting on-line self report requires skill, analogous to on-line translation.

Apart from being incomplete, verbal data are often -though not inherently- inaccurate. Subjects, when asked to explain their behaviour, often 'fill the gaps' by sensible guesses, rather than accurate data (Ericsson & Simon, 1980). Self report data suffer least from this problem, but are seldomly used in knowledge acquisition (Welbank, 1983, Grover, 1983, Hayes-Roth, 1983). The reasons for not using self report data become clear if we consider the requirements for use and interpretation of the different types of data. With the order presented in table 1, the following requirements for the use of a method become more severe:

* The amount of acquaintance the knowledge engineer has with the basic concepts in a domain.
* The amount of cooperation that is required from the expert. Experts prefer interviews to self report.
* The amount of interpretation tools to process the data in a consistent way.

Although self report data provide the most reliable information, planning self report sessions requires considerable knowledge of the domain and the types of problems that the expert normally solves. Further, the interpretation of self report data requires a much more powerful model, than interview data (Welbank, 1983). In the next section we will further probe into this interpretation problem. Section 5 will discuss recommendations on the use of the various data acquisition methods in different stages of the knowledge acquisition process.

2. The interpretation process

The purpose of the interpretation process is to establish a mapping between verbal data and knowledge structures. This mapping can be performed on different levels, depending on the types of constructs that are used to express the knowledge. Sloman (1979) analyses the different degrees of depth at which questions about knowledge can be asked: questions at the individual level, at the conceptual level, at the level of a particular formalism and at the level of mechanisms that implement a formalism. These levels closely correspond to the levels of representational primitives for different kinds of knowledge that Brachman (1978) distinguishes: **implementational**, **logical**, **epistemological**, **conceptual** and **linguistic** knowledge.

For the purpose of mapping verbal data onto knowledge we propose five levels representing a synthesis between Sloman's classification and Brachman's representational levels.

knowledge identification
This level of analysis corresponds to simply recording what one or more experts report on their knowledge. Although the result may be in a formalised form, the representational primitives on which this formalisation is based are linguistic (in the sense that Brachman uses this term). The same knowledge of different experts may have to be represented differently, because they use different terminology, or because their knowledge is structured in a different way.

knowledge conceptualisation
aims at the formalisation of knowledge in terms of **conceptual relations**, **primitive concepts** and **conceptual models.** The knowledge of different experts, and possibly of different subdomains, is unified within one conceptual framework.

epistemological analysis
At the epistemological level the analysis uncovers **structural** properties of the conceptual knowledge, formalised in an epistemological framework. Such a framework is based on **epistemological primitives** representing types of concepts, types of knowledge sources, structuring relations (such as hierarchical relations, inheritance), and types of strategies.

logical analysis
This level of analysis applies to the **formalisms** in which the knowledge on higher levels is expressed and which is responsible for inference making.

implementational analysis
At this level of analysis, **mechanisms** are uncovered on which higher levels are based. The representational primitives are the ones which are normally used when an implementation of an AI program is described (e.g. matching, testing, slot-filling, etc.).

Most research on knowledge-based ("expert") systems, has been mainly concerned with the problem of mapping knowledge at the first (linguistic) level onto the implementation level. Only recently an appreciation of the importance of other levels is beginning to emerge.

In particular the work of Clancey (1983) shows the importance of analysis of knowledge on the epistemological level both for eliciting proper explanations and for structuring knowledge sources and uncovering strategies of knowledge use. Recent work on deep knowledge in expert systems (Davis, 1983) also indicates the need for analysis at an epistemological level. Here we take the stand that interpretation of verbal data on the epistemological level is a crucial step in the design of knowledge based systems, not only because the gap between the data and the implementation level is generally too wide to be bridged in one step, but also because there is less danger of loosing information in the process. In addition we will show how the use of interpretation techniques on the epistemological level will guide the planning and conducting of verbal data acquisition. The next section will describe template models that can be used for the interpretation on the epistemological level.

3. Interpretation Models for Knowledge Acquisition

An interpretation model consists of a typology of basic elements and structuring relations for a certain class of domains. The basic elements that we distinguish (cf. Clancey, 1983) are: objects, knowledge sources, models, and strategies.

An object typology for a class of domains characterises the types of objects that have to be identified during the knowledge acquisition process. For example, many domains deal with quantificational objects. A typology for such objects will have to include concepts such as: dimension, quantity, measure, constraint, range etc. A diagnostic domain (or rather a domain of type "Recommend-Actions", Bennett, 1983)), will contain object types such as: evidence, hypothesis, cause, action. Typical object types in design tasks are: plan, specification structure, constraint.

Following Clancey we define a Knowledge Source (KS) to be a piece of knowledge that derives (infers) new information from existing data. KS's may be formalised in terms of rules (e.g. MYCIN), frametype structures (e.g. Friedland, 1981) or complex processes (like in Hearsay). Several types of generic KS's can be identified. **Identification KS's** classify an object in terms of its properties. Typical examples of identification KS's are rules in diagnostic systems which classify objects on the basis of sets of features (Clancey, 1983): 'if the organism is gram-negative, anaerobic rod, its genus may be bacteroides (0.6)'. **Causal KS's** relate information on the basis of empirical observations or causal models. **Factual** KS's represent common sense or domain specific relations, which are based on empirical facts ('grass is green') or on definitions within the domain ('adiabatic implies that there is no exchange of heat'). **Transformation KS's** transform representational objects into new objects on the basis of some constraining model. A typical example is knowledge about how to transform an algorithmic description of a device to a description which is implementable in hardware (e.g. Mostov, 1983). **Planning KS's** generate new information on the basis of a goal and a number of refinement constraints. Typical examples are decomposition rules in design tasks. **Information management KS's** specify how and when certain information should be stored for later use. Some of MYCIN's mapping rules are examples of this type. **Algorithmic KS's** specify in detailed steps how a particular goal can be achieved, or how a task should be performed.

Models are knowledge structures which represent a set of complex relationships in a coherent structure, which can be used to predict new information. Models are distinguished from KS's here, because they are of a more declarative and generally more complex structure than KS's. In many domains models play an important role in reasoning. Most current expert systems do not explicitly represent these models in the knowledge base, but use KS's of some sort as a 'compiled out' representation of an underlying model (cf. Clancey, 1983). During knowledge acquisition it is however important to explicate the models that underly the expert's reasoning, not just for a better understanding of the expert's knowledge but also for documentation, maintenance and justification purposes.

Types of models include: **causal** models, in which causal relations are explicitly represented, **process** models, in which relations between events or operations are represented, but not necessarily in a causal form, **formal** models based on a mathematical or physical theory, **empirical** models, based on empirical (e.g. statistical) evidence and **geometric** models, representing the geometric structure of a configuration.

Both objects and KS's can be organised in a structural way. The most common structures are based on hierarchies organised along a particular relational dimension such as sub/superclass, part/whole, refinement/broadening. Alternative organisations include: family classes organised around prototypes, relational networks (such as used in INTERNIST, Pople, 1982).

The structure of the knowledge base can support particular problem solving strategies. Different ways of structuring the knowledge base to support different strategies have been discussed by Clancey (1983) and will not be repeated here.

Table 2 shows a summary of an interpretation model for a diagnostic task (cf. Clancey, 1983; Szolovits, 1981; Bennet, 1983).

object types:
- symptoms
- evidence on basis of tests
- problems
- causes

knowledge source types:
- identification (data --> hypotheses)
- causal (causes --> data; causes --> problems)
- factual (facts --> facts)
- complicational (data+unusual situation --> hypotheses)

strategic knowledge types:
- hypothetico-deductive
- backward chaining
- successive refinement of hypotheses

structures:
- hierarchy of fault classes/causes
- hierarchy of symptoms & evidence factors
- families centered around prototypes
- relational networks

models:
- causal models
- process models
- formal (mathematical/physical/statistical) models
- geometric models

Table 2. Sample interpretation model for diagnostic tasks

4. A Sample Interpretation

As an illustration of the use of interpretation models, we present here a summary of an analysis of a three hour interview on the control of the vinification process (the fermentation of grapes into wine), in professional wine making. The expert was a professional wine maker in the south of France, producing a relatively high quality wine.

The domain is that of process control, but has aspects of a diagnostic domain as well, since many of the actions that influence the vinification process are preventive or remedial. In addition, a planning process which establishes the type of wine and projected time of sale, provides constraints on the decision making processes of the expert.