

Lecture Notes in Computer Science

Edited by G. Goos and J. Hartmanis

232

Fundamentals of Artificial Intelligence

An Advanced Course

Edited by W. Bibel and Ph. Jorrand



Springer-Verlag

Berlin Heidelberg New York London Paris Tokyo

Editorial Board

D. Barstow W. Brauer P. Brinch Hansen D. Gries D. Luckham
C. Moler A. Pnueli G. Seegmüller J. Stoer N. Wirth

Editors

Wolfgang Bibel
Institut für Informatik, Technische Universität München
Postfach 202420, D-8000 München 2

Philippe Jorrand
LIFIA-IMAG
BP 68, F-38402 St. Martin d'Hères Cedex

CR Subject Classifications (1985): F.4.1, I.2.3, I.2.4, I.2.6, D.1.3

ISBN 3-540-16782-X Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-16782-X Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically those of translation, reprinting, re-use of illustrations, broadcasting, reproduction by photocopying machine or similar means, and storage in data banks. Under § 54 of the German Copyright Law where copies are made for other than private use, a fee is payable to "Verwertungsgesellschaft Wort", Munich.

© Springer-Verlag Berlin Heidelberg 1986
Printed in Germany

Printing and binding: Beltz Offsetdruck, Hemsbach/Bergstr.
2145/3140-543210

PREFACE

The expectations in Artificial Intelligence – or Intellectics – have never been as high as they are today. Clearly, they are too high given the current state of the art in this fascinating field. Although we have seen some remarkable systems performing extremely well, there is no doubt for those who understand how they work that this performance reflects just a beginning in our understanding of the fundamentals that might be required in systems to perform in a truly intelligent way.

One of the basic paradigms in Artificial Intelligence has always been that experiment needs to be complemented with theory, or vice versa. During a wave of experimenting throughout the world in AI we feel that some emphasis on the more theoretical side might be appropriate – and necessary for progress in this understanding of the fundamentals of AI.

This sort of reflection motivated us when we took the initiative in organizing the first Advanced Course in Artificial Intelligence that was held in Vignieu, France, in July 1985. Seven well-known AI researchers were asked to cover basic topics that might be of relevance for the fundamentals of our field. The present volume comprises the elaborated and harmonized versions of their lectures. Most of them have been written in the form of a tutorial, so that the book provides a most valuable guide into the more fundamental aspects of AI.

One might be inclined to say that intelligence is the capability to acquire, memorize, and process knowledge in a way that appropriate accommodation in a changing world is achieved. In any case, the concept of knowledge and its representation clearly are among the fundamental issues in AI. The book begins with this topic discussed in the contribution by Delgrande and Mylopoulos in order to give the reader a feel for the variety of aspects that have to be taken into account for the later issues raised in the book.

The subsequent four articles are dealing with the second focus in this book which is the processing of knowledge. Obviously, knowledge cannot be processed unless it is adequately represented; on the other hand, an appropriate representation is determined by the way of processing needed. This is to say that these two issues are intimately related with each other. It is therefore not accidental that some of the aspects raised in PART ONE reoccur in different context here in PART TWO (such as in the contribution by Bibel).

For some processing of knowledge would be synonymous with computation, for others with inference or deduction. We see in the remarkable contribution by Huét that both are just different aspects of the same phenomenon which is studied in great depth in this contribution. The reader may find it rewarding to overcome the difficulty of studying such a formal and concise text.

One of the most successful tools for processing knowledge represented on a logical level of language is resolution. The chapter written by Stickel gives a thorough and knowledgeable introduction into the most important aspects of deduction by some form of resolution. It may be regarded as the basis for the more advanced forms of reasoning discussed in the subsequent two papers, but also for the programming language PROLOG.

The deductive forms of reasoning captured by resolution in its pure form do by no means exhaust the kinds of knowledge processing known from human experience and studied in Artificial Intelligence. A particularly important one is the kind of reasoning that is involved in inductive inferencing, problem solving (or programming) from examples, and in learning. It is covered in the contribution by Biermann that the reader will like also for its style of presentation.

There are many more forms of knowledge processing and inferencing than those discussed in the previous three papers. The tutorial by Bibel covers the more important among those that might play a significant role in common-sense reasoning. It takes, however, the position that the basic deductive tools like resolution or the connection method are essential for these forms of inference as well.

The third part of the book focuses on the more advanced programming tools for implementing the kind of systems that are envisaged with the topics discussed before. Logic programming and functional programming are known as the main styles for AI programming. Both lend themselves to a parallel treatment.

In this part of the book, Jorrand takes the approach that the semantic elegance and the mathematical properties of functional programming languages can be preserved within a language where computations or inferences can also be described as networks of cooperating parallel processes. This is shown in the language FP2, where term rewriting forms the basis for the semantics of both functional and parallel programming.

By its nature, PROLOG lends itself to parallel processing. To some extent the programmer might wish to control such parallel processes in PROLOG programs without compromising PROLOG's elegance as a descriptive language. Concurrent PROLOG developed by Shapiro and his group provides such features. His introduction to this language stages the adequate finale of the whole book.

There is an obvious lack of more good textbooks in Artificial Intelligence. One reason is that in a rapidly progressing field like AI it is nearly impossible for a researcher to actively contribute to the field's progress and at the same time be able to overlook large portions of this developing area, not mentioning the time needed for working out an appropriate presentation. A volume like the present one certainly cannot be a substitute for a textbook. But nevertheless we feel that it may be regarded as a good compromise for the time being.

In this sense we hope that the book not only refreshes the memories of those who attended the course, but serves as a unique source of valuable information for graduate classes and individuals with interest in the more fundamental and advanced topics of this exciting area of research.

München and Grenoble, April 1986

W. Bibel and Ph. Jorrand

CONTENTS

PART ONE: KNOWLEDGE REPRESENTATION

| | |
|---|---|
| Knowledge Representation: Features of Knowledge J.P. Delgrande and J. Mylopoulos | 3 |
|---|---|

PART TWO: KNOWLEDGE PROCESSING

| | |
|---|-----|
| Deduction and Computation G. Huet | 39 |
| An Introduction to Automated Deduction M.E. Stickel | 75 |
| Fundamental Mechanisms in Machine Learning and Inductive Inference A.W. Biermann | 133 |
| Methods of Automated Reasoning W. Bibel | 171 |

PART THREE: KNOWLEDGE PROGRAMMING

| | |
|--|-----|
| Term Rewriting as a Basis for the Design of a Functional and Parallel Programming Language. A case study: the language FP2 Ph. Jorrand | 221 |
| Concurrent PROLOG: A Progress Report E. Shapiro | 277 |