

document processing fall into a large number of different categories and not all of them support the notion of integration in the same way.

The book consists of 12 chapters, and for reviewing purposes, we can usefully divide the wide range of topics covered into three broad classes: text processing including hypertext (five chapters), storage media and applications (four chapters), and information retrieval (two chapters). In addition, an introductory overview chapter is provided as well as a subject index.

The text processing theme is introduced early to familiarize the reader with the idea of language as data and the notion of text representation. It also provides a focus for corpus-based studies such as concordances and the use of computers in lexicography, taking the Cobuild project as an example. The text processing theme is taken up again in the concluding chapters which deal with hypertext, document description and mark-up languages (e.g. SGML; ODA), desk-top publishing (Ventura) and formatting systems (LaTeX and troff), and finally, Postscript. Taken together, I found these parts of the book the most substantial and informative, although I missed a systematic discussion of editing: a mention of Emacs would not have been out of place, even if only for historical reasons.

The storage media and applications theme includes optical storage, CD ROM, worm disk and video disk, and the author has made some attempt to pair up each kind of hardware with associated applications. The treatment of applications is not particularly homogeneous, consisting sometimes of descriptions of rather specialized examples (e.g. CD-ROM and British Airways Technical Publication, a set of manuals for aircraft maintenance); sometimes of very general application domains (e.g. worm disk and document image processing, video disc and computer-based training).

The initial chapter on information retrieval serves to introduce some underlying concepts such as query languages, database design as well as extensions to the basic techniques to handle morphological variation, output ranking, and so on. The second chapter considers issues associated with the implementation of full text retrieval systems (such as Reuters Newsbank) using commercially available software (e.g. the ICL CAFS extension) and the use of relational models based on SQL). Finally, a number of well-known techniques for text compression, such as Huffman coding, are described.

All in all, the book is certainly informative, being packed with information and carefully selected sets of examples. It also includes references on a chapter-by-chapter basis. I also liked the fact that an attempt has been made to impose a certain level of uniformity on the structure of each chapter: for each ends with a chapter summary and a set of investigations intended for the curious student. Only the chapter devoted to optical storage, which is far shorter than all the others, does not conform to this pattern.

My main worry about the book concerns its intended readership, since I get the feeling that it has been written very much in the wake of a particular course rather than than being designed with a particular kind of reader in mind. The content is, for reasons presumably having to do with the nature of the course, pitched at that uneasy level that sits between that of a simple guide intended to attract and maintain the attention of the uninitiated, and that of an undergraduate text which must, amongst other things, attempt to instill some expertise in the subject matter. There is a danger that there is too much detail to please the first kind of reader but not enough focus to satisfy the second.

Reviewed by Michael Rosner, IDSIA, Corso Elvezia 36, 6900 Lugano, Italy.

Knowledge engineering toolkits edited by C. J. Price, Ellis Horwood, Chichester, UK, 1990, pp 261, £44.95, ISBN 0-13-517178-4.

The stated goal of this book is to provide a foundation for determining the suitability of an application for implementation with knowledge-based systems (KBS) technology, as well as serve as a guide to selecting an appropriate KBS toolkit for this task. This is certainly a worthwhile topic,

and an authoritative guide on this subject would be a welcome addition to the library of anybody interested in developing knowledge-based systems. Unfortunately, the book falls short of this rather ambitious goal, providing readers instead with, at best, an informal introduction to the field of knowledge-based toolkits.

The book is split into three major sections, each section containing a number of essays written by different contributors. Part 1 provides the reader with a brief overview of the components of a KBS: Basic reasoning strategies, knowledge representation methodologies, and graphical user interfaces. Part 2 reviews a number of commercially available KBS toolkits including POPLOG, Knowledge Craft, ART, KEE, Nexpert, Goldworks, VP-Expert and LEVEL5. Furthermore, the section introduces some background information about typical KBS applications such as classification, planning, design and diagnosis. The concluding part describes four different applications that were built using the various KBS toolkits presented in the previous chapters: Two classification systems using LEVEL5 and Nexpert, respectively, a process control application using POPLOG, and a very simple tile-configuration task using the ART KBS toolkit. The section concludes with guidelines for selecting an appropriate KBS toolkit, and a description of research activities in the field of KBS.

Clearly, the authors of this book possess a significant amount of practical experience in building expert systems using KBS toolkits. Unfortunately, the book never establishes a consistent framework into which they could fit this expertise. Instead, the descriptions of the various KBS toolkits consist mostly of collections of casual descriptions of the various systems, making it difficult for a reader to compare the various tools, or utilize the information to select an appropriate tool for their application. For example, Chapter 6 states that the KEE tool is integrated with an ATMS (Assumption-based Truth-Maintenance System). No mention is made, however, of how such a subsystem could be utilized, what kinds of applications could benefit from using it, or why an ATMS might justify the price of such a high-end KBS toolkit. A similar lack of analysis plagues the four case studies, which may inform the reader that the development of a system was started "in version 9 of POPLOG (on VMS 3.7) and later moved over to version 11 (on VMS 4.2)", but has very little useful information about the appropriateness of using the particular toolkit, nor does it offer any comparisons to using an alternative toolkit to achieve a similar result.

A second weakness of the book is its lack of direction. For example, the chapter entitled "Improving present-day toolkits" actually *precedes* the chapter introducing the various KBS toolkits. Despite the fact that this is an interesting chapter describing issues in knowledge acquisition and validation, it certainly does not belong there, and even its inclusion in such a book is rather questionable. In addition, if the book is really meant to address an audience of starting KBS application developers, the authors digress too often into research topics, and rely too much on references for readers to acquire more information. Finally, the overall typographical quality of the book is poor and it was clearly written in a hurry, as evidenced by the lack of standard examples, inconsistent graphics, and sparse index.

Notwithstanding the above problems, this book does contain some interesting material. It is very readable and covers a wide range of topics, bringing together the authors' experiences with using KBS tools to develop expert system applications. In particular, the chapters written by the editor provide some practical insights and guidelines for developers interested in building knowledge-based applications. At the same time, readers who are not familiar with the ongoing work in academia may be interested in the numerous references to research projects and software tools that are commonly used by KBS researchers.

In conclusion, if you are just getting started in the KBS field and want to get an overview of the capabilities of a number of KBS toolkits, or are not engaged in AI research and would like to get a very gentle introduction to some of the various KBS research activities, you might want to take a look at this book. If, however, you are looking for a guide for selecting an appropriate KBS toolkit for your work, you will have to continue looking.