Although the title implies a general view of manufacturing applications of AI, the book has a strong bias towards planning and scheduling, and towards architectures for integration (the topics of the workshops). After a brief introduction and overview, it is divided into three parts, but these do not have separate introductory sections, nor do they contain overview or review articles.

The title of the first part, "Design and Planning", is rather misleading, for there is little discussion of design. Instead, the six articles concentrate on process and assembly planning—the identification of the sequences of operations required to manufacture and assemble an engineering product. The papers range widely from detailed discussion of reasoning with engineering part descriptions through application of machine learning in planning to more heuristic approaches in fixture design and operation planning. The computational detail of the approaches is most complete in this part, and the individual papers are each of interest to researchers in the area. However, one struggles to find an overview of the topics, and the papers are not a very representative cross-section of work in the domain.

Part 2 is entitled "Scheduling and Control", and the emphasis here is very much on the former. Production scheduling is closely linked to process planning in that it involves the scheduling of manufacturing operations for all the parts made by a manufacturing system. There are again six chapters, and these concentrate on scheduling using distributed, intelligent agents, on planning under uncertainty, and on reasoning, search and learning strategies. The papers form a coherent and readable collection, and the subject reviews are generally more comprehensive than in Part 1. Conversely, the detail in the articles is sometimes less complete.

The final six chapters form Part 3, which is devoted to "Manufacturing Integration". In this part the book returns to some of themes of Part 1, including the integration of manufacturing applications with product modelling using high-level abstractions known as features, and the building of process knowledge bases. The themes of intelligent agents and distributed problem solving also reappear, within the context of the integration of diverse applications. Once again, the implementational detail is sometimes a little sparse, but overview material is generally quite useful.

The editors note in their introductory remarks the relative lack of penetration of AI techniques into real engineering systems. It is my view that much of the value of work on AI in manufacturing to-date is not so much in the systems which have been delivered, but in the insights which are afforded into the manufacturing process, and in the techniques which are migrating into more mainstream CADCAM and CIM work. In this respect, the work is peppered with good insights and useful techniques.

The book is, however, an edited collection, and it therefore shares the strengths and weaknesses of the genre. It reflects very much the view of its author community. It does not cover the breadth of manufacturing applications of AI. Nor, as noted, does it introduce specific subjects well, because there are no overview sections, and the subject reviews in individual articles are sometimes weak. It is well produced in a consistent style and, usefully for an edited book, it does have an index, although this is little more than an index of words in section headings. It is not, therefore, a book for teachers or undergraduate students, nor indeed for practitioners at the applications end of industry (except perhaps to browse for the ideas and insights noted). But for the researcher in applications of AI in process planning, in production planning and scheduling, or in systems integration, there are a number of competent research papers, and the book is therefore well worth reading.

Reviewed by: Chris McMahon, Department of Mechanical Engineering, University of Bristol, Bristol BS8 1TR, UK.

Advances in cognitive science: Volume 2: theory and applications by G. Tiberghien (ed.), Ellis Horwood, Chichester, 1989, pp 206, £39.95, ISBN 0-7458-0532-9.

This book, from Ellis Horwood's series in cognitive science, claims to represent the state of cognitive science in France in 1989 (though some non-French work is included). It is a deliberate attempt to bridge the gap between the English-speaking and French-speaking cognitive science communities. As such, it is somewhat eclectic: apart from the French connection, there is little coherence between the papers presented. I fear that this severely limits the appeal of the book. It surveys a number of issues, but it is not an introductory text. Consequently, it falls between students of cognitive science and specialists in particular sub-areas. I suspect that few would be interested in more than a couple of the papers it contains. The production is also disappointing, with many typographical errors and some awkward (and I suspect misleading) translation.

An introduction by Guy Tiberghien precedes four theoretical and five applied papers. The introduction tends towards jargon, and presents as fact a somewhat contentious definition of cognitive science. Part A addresses the "new" theoretical issues of concept representation, interpretation biases in understanding, learning arbitrary associations and the computational treatment of context. Part B discusses computer vision models of object recognition, applications of connectionist networks to facial feature extraction and arithmetic, a system for "concept" extraction from verbal characterizations, and cognitive approaches to process control.

Part A begins with a paper by Gregory Murphy and Edward Wisniewski on "Feature correlations in conceptual representations". The authors are concerned with the relationship between features and concepts involving those features. They describe a number of experiments which suggest that feature correlations as such are not learnt, but that feature-concept associations are.

In the second paper, "Interpretation biases in problem solving", Jean-Francois Richard argues that misinterpretations of instructions in problem solving situations are due to interpretation biases, and these biases give evidence of the processes involved in comprehending instructions. Three sources of bias are identified: typicality effects (where atypical interpretations are avoided); analogical transfer (where the problem is incorrectly reduced to a known one); and general metarules (where general rules are employed even though they may not be applicable). Several experiments are cited in support of each source of bias. Richard does not provide precise criteria for distinguishing these three sources; nor does he draw general conclusions concerning interpretation. Nevertheless, the phenomenon of interpretation biases seems undeniable.

Bernard Ans, in the third paper, "Learning arbitrary associations: a neuromimetic model", considers the general problem of learning arbitrary associations with non-linearly separable patterns. (By "neuromimetic" Ans appears to mean "connectionist".) He presents a three-stage model in which successive layers essentially refine the mappings learnt by previous layers. The model is claimed to be biologically plausible, though this is undermined by the requirement of three distinct phases for learning.

The last of the new theoretical issues to be broached concerns the treatment of context in cognitive processing. In "Contextual cognitive machines", Bernard Amy argues that whilst context is central to cognition, it is typically treated very naively. This is probably true, though some reference to other literature on the treatment of context (e.g., situation semantics, situated automata) would be helpful. Amy claims that context may be distinguished from other inputs because though possibly necessary for the functioning of a machine, contexts are not sufficient, whereas other inputs are both necessary and sufficient. I remain to be convinced of this distinction. Amy continues by treating context as a parameter, claiming that this parameter cannot be treated appropriately by standard symbolic computing, which, he claims, is unsuited to processing parameters which may vary while the machine functions. The final (and unsubstantiated) claim is that neural networks are able to treat context appropriately (as further inputs to a network).

The applied section of the book begins with a survey of "Computer vision models of object recognition" by Augustin Lux. Lux presents a number of extant computer vision models within a single framework, and examines each for computational and phenomenological adequacy. He concludes, from the fact that almost any knowledge may influence recognition, that "the theoretical foundations for recognition must be layed in the context of the whole of a cognitive system."

Following this, Alice O'Toole and Herve Abdi consider "Connectionist approaches to visually-based facial feature extraction". In particular, they construct a "memory matrix" according to an autoassociative model. Qualitative analysis of the eigenvectors of this matrix suggests that they typically correspond to macro-features, such as head/hair shape or sex of the owner of the face. It remains for future work to determine the psychological relevance of such features for face discrimination and characterization.

Simple arithmetic poses a number of problems for standard neural networks. Not the least of these is the representation of number. In "Representing simple arithmetic in neural networks", Susan Viscuso, James Anderson and Kathryn Spoehr propose that magnitude information must be central to any such representation. On the basis of this, they describe a network for "qualitative multiplication". Whilst the network exhibits some psychologically interesting behaviours (such as associative inference, practice effects, and the symbolic distance effect), the authors admit that the neural network approach remains unsatisfactory.

Jean-Francois Le Ny discusses a system for "Automatic induction of concept contents from verbal characterizations". Within the system, which is used for evaluating students understanding of concepts, input is parsed by ATN and converted into an initial frame/script like structure representing its semantics. This structure is then refined by a "summarizer/schematizer" before being compared to a representation of the target concept.

In the final paper, "Cognitive approaches to process control", Jean-Michel Hoc reviews information processing approaches to process control. He proposes a "general cognitive characterization of process control", together with a typology of process control situations, and argues that extent theoretical frameworks require enrichment in three areas: the representation of dynamic environments involving continuous variables; planning in dynamic environments; and cooperative problem solving.

Reviewed by: Richard Cooper, Department of Psychology, University College London, UK.