In addition to the numerous line drawings there are fourteen glossy colour charts. Eleven feature vertebrate phylogenies in considerable detail, with many genera mentioned. In contrast two charts of plant and invertebrate phylogenies are very poor and largely at phylum level; the latter lacks a time-scale.

The over-emphasis on vertebrates throughout the book reaches ludicrous proportions when, under the heading "Zone fossils", vertebrate groups receive more space than invertebrates.

There is a glossary which is again heavily vertebrate-oriented. The geological entries are sometimes appallingly incorrect and could not have been written by a professional geologist; the authorship of this section is not given.

The book is presumably aimed at the layman who wants to learn something about life through Earth's history. If he seeks information on extinct vertebrates, then this may be the book for him, but it will otherwise have a limited appeal. In spite of the excellence of the great majority of the drawings and its wealth of detailed information on vertebrate fossils, its inaccuracies and inconsistencies together with, for most readers, an unacceptable imbalance in its content, render the title *Encyclopaedia* a pretentious one. Though it is attractively laid out and handsomely bound, at £12 for 218 pages it is an expensive way to cover your coffee-table.

J.C.W. Cope

J.C.W. Cope is Senior Lecturer in Geology at University College, Swansea, UK.

Artificial intelligence at MIT

Artificial Intelligence. (Two volumes). Edited by P.H. Winston and R.H. Brown. Pp.492 and 486. (MIT Press: Cambridge, Massachusetts, and London, UK, 1979.) £16.25; $26 each volume.

AMONG the new scientific arts which have arisen since World War II, artificial intelligence must surely be among the most consequential and the most widely misconstrued. Yet its practitioners are neither diffident nor ambiguous in defining their field of study. "Its central goal", say P.H. Winston and M. Brady in their foreword to a new series from MIT Press, "is to make computers intelligent, both to make them more useful and to understand the principles that make intelligence possible".

The first fruit of the MIT Press series is a two-volume collection of 29 papers by graduate students, young post-doctoral workers and senior staff at the Massachusetts Institute of Technology. They are scattered over a variety of AI-related topics and first appeared, usually in longer form, as technical memos. The editors, P.H. Winston (who now directs MIT's AI Laboratory) and R.H. Brown, have grouped them into sections as follows: in Volume 1 — expert problem solving, natural language understanding and intelligent computer coaches, representation and learning; in Volume 2 — understanding vision, manipulation and productivity technology, computer design and symbol manipulation.

The reader who hopes for a true perspective view, a shining framework of deductively related principles and laws, will recoil from this congeries. Thus some pre-Newtonian seeker after a science of mechanics might have recoiled from Leonardo's notes on motion, weights, impact and moments of force. Technology's most thrilling advances commonly take place through spasms of untidy thought and activity. While waiting for the orderly abstraction of the theoretician, we may have to quarry general principles from the disorder of notes made in the heat of the struggle by the protagonists themselves. The editors of this book do, however, give conscientious aid to the quarrying process. Very full summarising prefaces precede each section, with an additional shorter note on each chapter. Newcomers to the field will be particularly grateful where specific AI techniques are listed and sign-posted to individual papers.

Some contributions, such as Stallman and Sussman on mechanised circuit analysis, represent deep and thorough finished studies. Others record brave first skirmishes with problems which will not fall before many such assaults have spent themselves. Johan de Kleer's study of qualitative and quantitative reasoning about roller-coaster mechanics is a stimulating example, and Winston's exercise in machine learning through simile raises hope of good eventual solutions in an extremely difficult area. In Marvin Minsky's "The Society Theory of Thinking", in which "the mind is viewed as an organised society of intercommunicating 'agents'", bravery shades into audacity, and the lines of contact with neurobiological fact are stretched exceedingly thin. Yet speculation may still in this field engender aids for the practitioner, as Minsky's earlier "frames" proposal has already amply shown.

Of all the difficult problems addressed by artificial intelligence workers, common consent gives the computational analysis of vision pride of place. Among the six vision papers of the second volume, David Marr's is outstanding. From the strenuous detail and breadth of this attack, methodological lessons can be learned. One of these is summarised by Winston and Brown: ". . . it is wrong to begin with devotion to some particular type of algorithm with a view toward finding some problem that it will solve". These words reflect AI's recent change of emphasis for which the MIT school can claim chief credit, namely that intelligent behaviour can only be conjured from a sufficiently rich and massive base of empirical knowledge about the specific domain appropriately organised in machine memory.

Artificial intelligence has thus followed an opposite course to that of the mechanical crafts, which in their early days lost themselves in the properties of specific domains. It was left to Leonardo to make the all-important step of abstracting the concept of a generalised mechanism — for example, for converting between rotary and reciprocating motion — which could then be made to function in any of a variety of different specialised machines. With AI it went the other way, and for 15 years most of us were hypnotised by the promise of general mechanisms — for heuristic search, for deductive reasoning, for machine learning, and the like — as though these universals could accomplish every problem-solving task without the additional need for building in carefully articulated representations of knowledge.

Indeed for a time the MIT school went to the opposite extreme and maintained the exuberant position that domain-specific knowledge should be inextricably interwoven ("procedural embedding") into the substance of the algorithmic mechanisms themselves, instead of maintaining, as we do today, separate "knowledge bases" to which these more general mechanisms can have access. It has now been sufficiently demonstrated that the resulting loss of modularity is fatal to incremental construction and flexible adjustment of an intelligent program's packet of skills. Incremental properties are quintessential. It is refreshing to note from these volumes that the point is now taken, including at MIT.

Space does not allow mention of many other rewarding views to be had of AI people at work. Lozano-Perez's discussion of the problem of telling a robot how to move its hand illuminates the intrinsic complexity of a task at which humans are hereditary Grandmasters. The software technologist will find Hewitt on control as message-passing a source of provoking novelty. Hewitt's chapter closes a book which will irritate purists and antagonise some professionals whose orientation is to Meccas other than MIT. But to those who wish to take a dip into the fast-flowing stream of an active and authoritative laboratory, the publication of this book offers a unique opportunity.

Donald Michie

Donald Michie is Professor of Machine Intelligence at the University of Edinburgh, UK.