Current trends in AI planning: EWSP 93—2nd European Workshop on Planning edited by C. Backstrom and E. Sandwall, IOS Press, Amsterdam, 1994.

AI planning is an important domain of research and applications in artificial intelligence. It aims at formally constructing computationally adequate representations and mechanisms to model reasoning about actions and plans which, of course, represents an important activity in the behaviour of any intelligent agent. What AI planning tackles can thus be regarded as fundamental to the development of intelligent machines that are able to deal effectively with real-world problems. And it also has been widely realized that this is an extremely difficult field.

Until the mid-1980s, AI planning had been constrained by what was called the classical problem framework, that is, at any instant of time the world is considered to be in one of a potentially infinite number of states; various properties remain unchanged in each of those states; the world changes from one state to the next only as a result of the occurrence of actions. AI planning thus restricted itself to only considering a single agent operating in a potentially static world model. However, the classical problem framework inspired an important approach, deductive planning. Considering an initial world model, deductive planning attempts to solve a problem by deductively proving that there exists such a state in which all goals logically become true. As a result of the proof, a sequence of actions is synthesized, which can transform the world from initial state to the goal state. In the past decade or so, the classical problem framework has been expanded, and many novel planning theories have been proposed to redefine *worlds, world models, actions* and *plans.* As a result, various new paradigms of AI planning have been established, in which concepts such as *simultaneous actions, external events, temporal constraints, reactivity, causality* and *uncertainty* play a more and more important role. However, deductive planning still amazingly remains a dominant approach in the majority of AI planning paradigms.

The book under review is a collection of articles based upon papers presented at the Second European Workshop on Planning, organized by the European planning community. After a brief preface, the book is divided into two parts, respectively titled *Invited Lectures* and *Selected Papers.* However, the articles can be roughly classified into several categories in terms of their subjects.

Although the title implies a general view of current trends of AI planning, the book has a strong bias towards logic-based deductive planning. A brief review of the domain can be found in *Invited Lectures.* This subject is covered by nine articles, and the emphasis is very much upon deductive planning with temporal, abductive and probabilistic reasoning. As a formal foundation, these patterns of reasoning are introduced into deductive planning on the basis of either modal temporal logic or probabilities logic to formalize time constraints, plan hypotheses, and world models with uncertainty. Detailed discussions about problem formalizations are provided. In addition, there are also some papers proposing algorithmic approaches to deductive planning, for instance, plan recognition with decision-theoretic models, planning under structural and temporal constraints, etc. These paradigms are also treated as deductive planning since the logical formalisms still play a key role in plan representation. Some of the papers in this category are excellent.

Multi- and autonomous-agent planning, which is another important subject of AI planning, is addressed in only three articles. They cover, however, some important issues in the area, such as the generation and execution of plans in a distributed environment in which communicating multi-agents operate, the emergent properties of autonomous systems, and building intelligent autonomous agents based upon the integration of planning and learning. The papers in this category present some significant aspects of progress in the area.

A number of papers carry out some in-depth discussions about classical planning which still remains one of the most important issues in AI planning. These articles attempt to address several hard classical problems by using new representations and mechanisms. The topics include the least-commitment criteria for converting a total-ordered plan into a least-constrained plan, hierarchical planning with a generalized operator abstraction, the reconstruction of the STRIPS-style formalisms based upon temporal reasoning, etc.

Finally, the book presents a number of articles addressing the issues that are closely related to planning, for example, plan control, plan execution, resource allocation and scheduling.

The book is a collection of research outcomes that present the significant aspects of progress in current AI planning. The emphasis is upon various logical formalisms and reasoning mechanisms to support deductive planning, and some of the approaches are certainly novel and promising. Another feature of the book is that many of the authors carry out their discussions in a very formal, concise way, such that main theoretical and technical issues in the articles are accurately addressed and expressed. As a workshop proceedings, however, the book can only portray the special interests of those who attended. Although deductive planning is a central subject of the book, the papers do not provide a very representative cross-section of work in the domain. With the introduction of a variety of non-classical logics, for example, a central issue for deductive planning is to develop sound proof theories supporting these reasoning patterns for plan generation. However, most of the articles in the book emphasize logical formalisms for plan representation without providing sufficient discussions of deductive reasoning procedures for plan generation. Besides, none of the papers addresses non-monotonic reasoning which is one of the most important logical formalizations in AI planning. Finally, it is my view that the book covers too few aspects of AI planning to match the title.

The book is not organized into subject-based sections. Nor does it have an index. It is not, therefore, a book for teachers or undergraduate students, nor indeed for practitioners at the practical application end. But for the researchers in AI planning, in automated reasoning, or in knowledge processing, the book is well worth reading.

Reviewed by Dr Huaming Lee and Mr Jon Sims Williams, Department of Engineering Math, University of Bristol, Bristol BS8 1TR, UK

Logic programming: operational semantics and proof theory by J. H. Andrews, Cambridge University Press, 1992, pp 104, £25.00, ISBN 0-521-43219-7.

This book is published by Cambridge University Press in its distinguished dissertation series in computer science. It does not qualify as a textbook, but researchers interested in the semantics of both sequential and parallel logic programming will find it an important contribution to the field.

The thesis provides operational semantics and proof theoretic characterizations for systems having each of the possible combinations of parallel or sequential "and" and parallel or sequential "or". A proof theoretic characterization associates a program with a proof system with respect to which the operational semantics is sound and complete.

The proposed operational semantics, Stack-of-Stacks (SOS), operates on programs consisting of definitions of predicates by goal-formula bodies, as opposed to SLD-resolution which operates on programs consisting of Horn clauses. Thus the proposed operational semantics manipulates goal-formulae as its basic units. The rules for SOS are simple and very predictable. The variants of SOS correspond to the parallel "and" and "or", sequential "and", sequential "or", and sequential "and" and "or" control disciplines.

The proof theoretic characterization is in the form of sequent calculi. The core part, LKE, of this calculi is basically a Gentzen-style sequent calculus for first-order logic with equality. LKE augmented with a set of axioms describes the success and failure behaviour of queries under certain assumptions. For example, LKE augmented with a set of PAR axioms characterizes the queries which succeed in parallel-or systems and those which fail in parallel-and systems. This means that if a query A succeeds (resp. fails) in a parallel-or (resp. parallel-and) system, then the sequent $[\rightarrow S(\exists[A])]$ resp. $[\rightarrow F(\exists[A])]$ is derivable from LKE + PAR. Similarly, LKE + SEQ characterizes those queries which succeed in the sequential-and, sequential-or system, and those which fail in sequential-and system; LKE + PASO characterizes those queries which succeed in the parallel-and, sequential-or system.