


---

*Readings in*

# PLANNING



*Edited by*

*James Allen, James Hendler, and Austin Tate*

*Readings in*  
PLANNING

*Edited by*

James Allen  
*University of Rochester*

James Hendler  
*University of Maryland*

Austin Tate  
*University of Edinburgh*



# CONTENTS

**Preface** v

**Foreword** by Nils J. Nilsson xi

**Part I Introduction to Planning** 1

**Chapter 1 Introduction** 3

**Planning** 5

Michael P. Georgeff

**A Review of AI Planning Techniques** 26

Austin Tate, James Hendler, Mark Drummond

**Formal Models of Planning** 50

James Allen

**Part II Planning Systems** 55

**Chapter 2 Beginnings** 57

**GPS, A Program that Simulates Human Thought** 59

Allen Newell, H. A. Simon

**Application of Theorem Proving to Problem Solving** 67

Cordell Green

**STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving** 88

Richard E. Fikes, Nils J. Nilsson

**Planning in a Hierarchy of Abstraction Spaces** 98

Earl D. Sacerdoti

<b>Chapter 3 Interactions and Dependencies</b>	<b>109</b>
<b>The Virtuous Nature of Bugs</b>	<b>111</b>
Gerald Jay Sussman	
<b>Achieving Several Goals Simultaneously</b>	<b>118</b>
Richard Waldinger	
<b>Extract for APIC Studies in Data Processing</b>	<b>140</b>
D. H. D. Warren	
<b>A Representation for Robot Plans</b>	<b>154</b>
Philip J. Hayes	
<b>The Nonlinear Nature of Plans</b>	<b>162</b>
Earl D. Sacerdoti	
<b>Planning with Constraints (MOLGEN: Part I)</b>	<b>171</b>
Mark Stefik	
<b>Chapter 4 Planning and Acting</b>	<b>187</b>
<b>Learning and Executing Generalized Robot Plans</b>	<b>189</b>
Richard E. Fikes, Peter E. Hart, Nils J. Nilsson	
<b>Decision Theory and AI II: The Hungry Monkey</b>	<b>207</b>
Jerome A. Feldman and Robert F. Sproull	
<b>Planning and Acting</b>	<b>225</b>
Drew McDermott	
<b>A Cognitive Model of Planning</b>	<b>245</b>
Barbara Hayes-Roth and Frederick Hayes-Roth	
<b>A Model For Planning in Complex Situations</b>	<b>263</b>
Robert Wilensky	
<b>Integrating Marker-Passing and Problem Solving</b>	<b>275</b>
James A. Hendler	
<b>Chapter 5 Integrated Planning Systems</b>	<b>289</b>
<b>Generating Project Networks</b>	<b>291</b>
Austin Tate	
<b>Planning in Time: Windows and Durations for Activities and Goals</b>	<b>297</b>
Steven A. Vere	
<b>Domain-independent Planning: Representation and Plan Generation</b>	<b>319</b>
David E. Wilkins	

- ISIS—A Knowledge-based System for Factory Scheduling** 336  
Mark S. Fox, Stephen F. Smith
- O-Plan—Control in the Open Planning Architecture** 361  
Ken Currie and Austin Tate
- Hierarchical Planning Involving Deadlines, Travel Time,  
and Resources** 369  
Thomas Dean, R. James Firby and David Miller

**Part III Foundations of Planning** 389

**Chapter 6 Formal Models Of Action** 391

- Some Philosophical Problems from the Standpoint of Artificial  
Intelligence** 393  
John McCarthy and Patrick Hayes
- A Temporal Logic For Reasoning About Processes  
and Plans** 436  
Drew McDermott
- Towards a General Theory of Action and Time** 464  
James F. Allen
- A Formal Theory of Knowledge and Action** 480  
Robert C. Moore

**Chapter 7 Formal Models of Planning Systems** 521

- On the Semantics of STRIPS** 523  
Vladimir Lifschitz
- Plan Synthesis: A Logical Perspective** 532  
Stanley J. Rosenchein
- Planning for Conjunctive Goals** 537  
David Chapman
- Planning Using a Temporal World Model** 559  
James F. Allen and Johannes A. Koomen
- Planning as Search: A Quantitative Approach** 566  
Richard E. Korf

**Chapter 8 Time and The Frame Problem 579**

**Problems in Formal Temporal Reasoning 581**

Yoav Shoham and Drew McDermott

**The Frame Problem and Related Problems in Artificial Intelligence 588**

Patrick J. Hayes

**Temporal Data Base Management 596**

Thomas Dean and Drew McDermott

**Nonmonotonic Logic and Temporal Projection 624**

Steve Hanks and Drew McDermott

**Why Things Go Wrong: A Formal Theory of Causal Reasoning 641**

Leora Morgenstern and Lynn Andrea Stein

**Part IV New Directions in Planning Systems 647**

**Chapter 9 Learning and Reuse 649**

**Selectively Generalizing Plans for Problem-Solving 651**

Steven Minton

**CHEF: A Model of Case-based Planning 655**

Kristian J. Hammond

**An Adaptive Planner 660**

Richard Alterman

**Chapter 10 Extending the Classical Framework 655**

**Refining and Extending the Procedural Net 667**

Mark E. Drummond

**Localized Representation and Planning 670**

Amy L. Lansky

**Formulating Multiagent, Dynamic-World Problems in the Classical Planning Framework 675**

Edwin P. D. Pednault

<b>Chapter 11 Planning and Execution</b>	<b>711</b>
<b>An Architecture for Intelligent Reactive Systems</b>	<b>713</b>
Leslie Pack Kaelbling	
<b>Reactive Reasoning and Planning</b>	<b>729</b>
Michael P. Georgeff and Amy L. Lansky	
<b>Integrating Planning, Execution, and Monitoring</b>	<b>735</b>
Jose A. Ambros-Ingerson and Sam Steel	



## FOREWORD

In the early days of artificial intelligence, the word “planning” meant something different than what it means today. In their summary description of the problem-solving behavior of humans, Feigenbaum and Feldman<sup>1</sup> wrote: “A subject using a *planning* method abstracts or simplifies a complex problem. He then solves the simpler problem and uses the information obtained in the solution of the simpler problem in the solution of the original complex problem. This technique has also been included in GPS...”

Around 1970, my colleagues and I at SRI began using the word “plan” in its more usual sense: *a detailed formulation of a program of action*.<sup>2</sup> This volume contains many of the important papers that propose and explore techniques for automatically formulating and using such programs of action. And, recalling the earlier definition, many of the papers are concerned with *hierarchical planning*, in which the high level actions that achieve a simpler, more abstract goal are used to guide the search for primitive actions to achieve a more detailed one.

*Action*. How closely associated it is with intelligence! Animals act; plants do not (except in a strained and limited sense). Animals are intelligent; plants are not. And the most intelligent animals (I suppose that’s us) are able to think before acting and condition actions on thoughts. Thinking about the consequences of actions before being forced to experience them is the essential part of planning; the result of such thinking is a plan, and the most important part of a plan is the part that specifies what to do next. Since all animals act, *something* tells them what to do next, but we really don’t know to what extent animals other than ourselves make and use plans.

We, the planning animals, know that planning is necessary for effective action, and therefore most artificial intelligence researchers want their computer systems to plan. The enterprise of attempting to build such systems has uncovered a multitude of questions about how to represent and reason about the effects of actions. We have the *frame problem*, the *qualification problem*, and the *ramification problem*—to name a few.

Because there are animals that seem too simple to be capable of planning but are nevertheless capable of quite complex behavior (consider the bee, for example), some artificial intelligence researchers<sup>3</sup> have decided to concentrate first on systems that act but that do not plan. It is as if those researchers believed that the ontogeny of artificial intelligence must recapitulate the phylogeny of natural intelligence.

The present volume contains most of the important work of researchers (myself included) who predict that we can bypass those eons of evolutionary history that produced only dull animals which made no plans. We want to build machines straightaway that would rank high

---

<sup>1</sup>Feigenbaum, E. A., and Feldman, J., *Computers and Thought*, McGraw-Hill, New York, 1963, page 278.

<sup>2</sup>*Webster’s Ninth New Collegiate Dictionary*, First Digital Edition, NeXT Inc. and Merriam-Webster, Inc., Merriam-Webster, Springfield, MA, 1988.

<sup>3</sup>See, for example, Brooks, Rodney A., “A Robust Layered Control System for a Mobile Robot,” *IEEE Journal of Robotics and Automation*, March 1986.



on the evolutionary scale, perhaps machines as good as or better than humans are in terms of thinking about what we are going to do before doing it. It's worth a try! And judging from the results reported here, we are making good progress.

Can we combine our work on planning with the results of the "action-without-planning" school? In Chapter 11, we begin to see a synthesis emerging. Certainly, some systems (even those that ordinarily base their actions on plans) will sometimes have to act quickly when there is no time to make a plan. A frontier problem, only now beginning to receive much attention, is how to build autonomous agents that have the ability to base actions both on complex and deep reasoning, as well as on more "instinctive," built-in patterns—using each method appropriately depending on the situation.

Other frontier problems are suggested by the analogy between planning and learning. Learning occurs when an agent modifies its behavior based on experiences in the world it inhabits. Planning occurs when an agent selects behavior based on "experiences" in an internal model of its world. Some initial attempts to link these two activities are reported in Chapter 9. It is also intriguing to speculate about whether or not recent work in "genetic algorithms" (which model evolutionary processes) might inform both learning and planning research. In a certain sense, learning is simulated evolution, and planning is simulated learning.

Pioneers who want to explore these and other frontiers of artificial intelligence research can do no better than to outfit themselves first with what is already known about planning, and the best of that knowledge is contained in this volume.

# Preface

Early in the growth of the field of artificial intelligence it was recognized that an important behavior for any intelligent agent was the ability to plan a set of actions to accomplish its goals. The attempt to realize programs with this ability has resulted in one of AI's main subdisciplines—the field of planning.

Since the first published papers on planning in the late 1950s, the field has grown and papers have appeared in virtually all of the major AI conferences and journals. Although numerous review articles on planning have appeared over the years, no systematic attempt to collect the major papers in this field into one volume has been made previously. Courses on planning have been difficult to teach due to the need to gather a large number of papers from disparate sources. In addition, although most AI texts have an introduction to planning, those wishing to begin applying the planning technology have not been able to find a volume pointing out the seminal papers in the field.

The goal of this book is to remedy this situation by bringing together in one place a set of readings that can be used to develop a familiarity with the planning literature, with the major AI theory underlying planning, and with recent papers in the field indicating several directions of current research.

This volume comprises four sections. The first section is a collection of papers introducing the field of planning. These papers are intended for those familiar with AI concepts, but lacking some familiarity with the research in planning. The second section presents work describing important advances in planning systems developed over the past 30 years. The third section concerns AI research in the areas underlying planning: representation of time and coping with the frame problem. This section also includes some papers that formalize planning systems described in the second section. A final section is included that presents some short papers discussing several directions in which current planning research is moving.

In putting together this volume several hard decisions needed to be made. The AI literature in planning is quite diverse and to include all the strands in a single volume would be impossible for the following reasons:

1. Planning systems have traditionally fallen into two major classes: domain-independent planners, which concentrate on the control mechanisms of planning, and domain-dependent planners which concentrate on using domain heuristics to encourage efficient search. The editors decided to concentrate on domain-independent planning in the current volume as it was felt that the more general design principles described herein would be applicable to a wider range of planning work.
2. Planning is related to many other subareas of AI. The underlying theoretical material in this volume is restricted to papers that were deemed to have an explicit relation to an understanding of planning systems. This necessarily meant leaving out much important work in logic and knowledge representation (much of which is available in other issues of the Readings series).
3. Many realistic applications of AI planning methods require extensive reasoning about time, limited resources and other constraints. Although this is an exciting area of current research, space limitations constrain us to including only a few early papers addressing this theme. Pointers to further reading on this subject can be found in the introductory material and in the included papers.
4. The area of plan recognition is a rich literature that is clearly related to the planning area. Unfortunately, to include enough papers on recognition to provide a fair treatment would require significantly more space than was available.

Putting together this volume would have been a far greater undertaking without the efforts of the people who contributed articles, many of whom also provided early feedback on the purpose of the volume or reviewed the proposed set of papers prior to solicitation of the manuscripts. In addition, support from Mike Morgan and Sharon Montooth of Morgan-Kaufmann Publishers was invaluable in bringing this volume to fruition. Finally, on a more personal note, we thank our family and friends who supported us during this effort.



MORGAN KAUFMANN SERIES IN REPRESENTATION AND REASONING  
Ronald J. Brachman, Series Editor

---

## *Readings in* PLANNING

---

Edited by James Allen (University of Rochester),  
James Hendler (University of Maryland)  
and Austin Tate (University of Edinburgh)

Although numerous review articles on planning have appeared over the years, no systematic attempt has been made to collect the major papers in this field into one volume. The goal of this book is to remedy this situation by bringing together in one place a set of readings that can be used to develop a familiarity with the planning literature, with the major Artificial Intelligence (AI) theory underlying planning, and with the exciting directions of current research.

The first section of the book introduces the techniques and terminology of the AI planning community. The second section collects papers describing planning systems developed over the past 30 years. A third section presents research in the area of formal models of planning. Finally, a fourth section collects recent work representative of the field's current direction. The editors, three accomplished researchers, provide introductory material surveying the field of planning as well as introductions to groups of papers placing the work in perspective.

Useful as a text for a planning course, a supplement to a more general AI course, and as a reference for AI system developers and researchers, this book is certain to be among the most popular of the Morgan Kaufmann Readings.

---

### OTHER TITLES OF INTEREST FROM MORGAN KAUFMANN

REASONING ABOUT PLANS by James Allen, Henry Kautz, Richard Pelavin  
and Joshua Tenenber

READINGS IN UNCERTAIN REASONING edited by Glenn Shafer and  
Judea Pearl

READINGS IN QUALITATIVE REASONING ABOUT PHYSICAL SYSTEMS  
edited by Daniel S. Weld and Johan de Kleer

PRACTICAL PLANNING: EXTENDING THE CLASSICAL AI PLANNING  
PARADIGM by David Wilkins

---

MORGAN KAUFMANN PUBLISHERS  
P.O. BOX 50490  
PALO ALTO, CA 94303-9953

ISSN 1046-9567  
ISBN 1-55860-130-9  
ARTIFICIAL INTELLIGENCE