

AUTOMATED PLANNING AND ACTING

Autonomous AI systems need complex computational techniques for planning and performing actions. Planning and acting require significant deliberation because an intelligent system must coordinate and integrate these activities in order to act effectively in the real world. This book presents a comprehensive paradigm of planning and acting using the most recent and advanced automated-planning techniques. It explains the computational deliberation capabilities that allow an actor, whether physical or virtual, to reason about its actions, choose them, organize them purposefully, and act deliberately to achieve an objective.

Useful for students, practitioners, and researchers, this book covers state-of-the-art planning techniques, acting techniques, and their integration which will allow readers to design intelligent systems that are able to act effectively in the real world.

Malik Ghallab is a Research Director at the National Center for Scientific Research, France. For most of his career he has been with LAAS-CNRS, University of Toulouse. His research activity is mainly focused on planning, acting and learning in robotics and AI. He has co-authored several books and more than two hundred scientific publications. He has been head of a French research program in Robotics, director of LAAS-CNRS and CEO for Science and Technology of INRIA. He is an ECCAI Fellow.

Dana Nau is a Professor at the University of Maryland, in the Department of Computer Science and the Institute for Systems Research. He has more than 300 refereed technical publications. Some of his best-known research contributions include the discovery of pathological game trees, the strategic planning algorithm that Bridge Baron used to win the 1997 world championship of computer bridge, applications of AI in automated manufacturing, automated planning systems such as SHOP, SHOP2, and Pyhop, and evolutionary game-theoretic studies of cultural evolution. He is an AAAI Fellow and an ACM Fellow.

Paolo Traverso is the Director of FBK ICT IRST, the Research Center at Fondazione Bruno Kessler (FBK). Paolo has worked in the advanced technology groups of companies in Chicago, London, and Milan, leading projects in safety critical systems, data and knowledge management, and service oriented applications. His research activity is mainly focused on planning and acting under uncertainty. His contributions to research in automated planning include the technique called “planning via symbolic model checking.” He has published more than one hundred papers in artificial intelligence. He is an ECCAI Fellow.

AUTOMATED PLANNING AND ACTING

MALIK GHALLAB

LAAS-CNRS, University of Toulouse, France

DANA NAU

University of Maryland, U.S.A.

PAOLO TRAVERSO

FBK ICT IRST, Trento, Italy



CAMBRIDGE
UNIVERSITY PRESS



CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment,
a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of
education, learning and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781107037274

© Malik Ghallab, Dana Nau, and Paolo Traverso 2016

This publication is in copyright. Subject to statutory exception and to the provisions
of relevant collective licensing agreements, no reproduction of any part may take
place without the written permission of Cambridge University Press & Assessment.

First published 2016

A catalogue record for this publication is available from the British Library

Library of Congress Cataloging-in-Publication data

Names: Ghallab, Malik, author. | Nau, Dana S., author. | Traverso, Paolo, author.

Title: Automated planning and acting / Malik Ghallab (Le Centre national de la recherche scientifique,
France), Dana Nau (University of Maryland, U.S.A.), Paolo Traverso (ITC-IRST (Center for Scientific and
Technological Research), Italy).

Description: New York, NY : Cambridge University Press, 2016. | Includes bibliographical references and
index.

Identifiers: LCCN 2016017697 | ISBN 9781107037274 (hardback ; alk. paper) | ISBN 1107037271 (hardback ;
alk. paper)

Subjects: LCSH: Production planning—Data processing.

Classification: LCC TS183.3 .G43 2016 | DDC 658.5/03—dc23

LC record available at <https://lccn.loc.gov/2016017697>

ISBN 978-1-107-03727-4 Hardback

Cambridge University Press & Assessment has no responsibility for the persistence
or accuracy of URLs for external or third-party internet websites referred to in this
publication and does not guarantee that any content on such websites is, or will
remain, accurate or appropriate.

To think is easy, to act is hard.
The hardest is to act in accordance with your thinking.
Elective Affinities, Johann Wolfgang von Goethe

Contents

List of Algorithms	<i>page</i> ix
Table of Notation	xi
Foreword	xiii
Preface	xv
1 Introduction	1
1.1 Purpose and Motivations	1
1.2 Conceptual View of an Actor	5
1.3 Deliberation Models and Functions	8
1.4 Illustrative Examples	14
1.5 Outline of the Book	17
2 Deliberation with Deterministic Models	19
2.1 State-Variable Representation	19
2.2 Forward State-Space Search	30
2.3 Heuristic Functions	36
2.4 Backward Search	46
2.5 Plan-Space Search	48
2.6 Incorporating Planning into an Actor	54
2.7 Discussion and Historical Remarks	59
2.8 Exercises	69
3 Deliberation with Refinement Methods	74
3.1 Operational Models	74
3.2 A Refinement Acting Engine	81
3.3 Refinement Planning	90
3.4 Acting and Refinement Planning	100
3.5 Discussion and Historical Remarks	106
3.6 Exercises	110
4 Deliberation with Temporal Models	114
4.1 Introduction	114
4.2 Temporal Representation	116
	vii

viii	Contents
4.3	Planning with Temporal Refinement Methods 127
4.4	Constraint Management in Temporal Planning 133
4.5	Acting with Temporal Models 141
4.6	Discussion and Historical Remarks 149
4.7	Exercises 153
5	Deliberation with Nondeterministic Models 155
5.1	Introduction and Motivation 155
5.2	The Planning Problem 157
5.3	And/Or Graph Search 165
5.4	Symbolic Model Checking Techniques 170
5.5	Determinization Techniques 179
5.6	Online Approaches 182
5.7	Refinement Methods with Nondeterministic Models 187
5.8	Acting with Input/Output Automata 198
5.9	Discussion and Historical Remarks 210
5.10	Exercises 214
6	Deliberation with Probabilistic Models 217
6.1	Introduction 217
6.2	Stochastic Shortest Path Problems 218
6.3	Heuristic Search Algorithms 232
6.4	Online Probabilistic Approaches 244
6.5	Acting with Probabilistic Models 253
6.6	Representations of Probabilistic Models 255
6.7	Domain Modeling and Practical Issues 261
6.8	Discussion and Historical Remarks 268
6.9	Exercises 273
7	Other Deliberation Functions 276
7.1	Perceiving 276
7.2	Monitoring and Goal Reasoning 280
7.3	Learning and Model Acquisition 287
7.4	Hybrid Models 294
7.5	Ontologies for Planning and Acting 304
8	Concluding Remarks 310
Appendix A	Search Algorithms 313
A.1	Nondeterministic State-Space Search 313
A.2	And/Or Search 315
Appendix B	Strongly Connected Components of a Graph 318
Bibliography	321
Index	351

List of Algorithms

2.1	Forward-search planning schema	<i>page</i> 30
2.2	Deterministic-Search, a deterministic version of Forward-search	30
2.3	H ^{FF} , an algorithm to compute the Fast-Forward heuristic	43
2.4	Backward-search planning schema	46
2.5	PSP, plan-space planning	50
2.6	Run-Lookahead replans before every action	56
2.7	Run-Lazy-Lookahead replans only when necessary	56
2.8	Run-Concurrent-Lookahead does acting and replanning concurrently	57
3.1	Main procedure of the Refinement Acting Engine (RAE)	84
3.2	RAE: progressing a refinement stack	85
3.3	RAE: trying alternative methods for a task	85
3.4	SeRPE, the Sequential Refinement Planning Engine	92
3.5	IRT, a refinement-planning algorithm that can do interleaving	98
3.6	Subroutine of IRT to simulate the next step in a method	99
3.7	Replanning before every action	100
3.8	Replanning only when necessary	101
3.9	Replanning concurrently with acting	101
3.10	Main procedure of REAP	103
3.11	REAP's procedure for progressing a refinement stack	104
3.12	REAP's version of RAE's Retry subroutine	105
4.1	TemPlan, a temporal planner	128
4.2	Path consistency algorithm for simple constraint networks	135
4.3	A dispatching function for eRAE	144
5.1	Procedure for performing the actions of a policy	160
5.2	Planning for solutions by forward-search	166
5.3	Planning for safe solutions by forward-search	166
5.4	Planning for safe acyclic solutions by forward-search	167
5.5	Planning for safe acyclic solutions by MinMax Search	168
5.6	The policy with minimal cost over actions	169
5.7	Planning for safe solutions by symbolic model checking	174
5.8	PruneUnconnected: Removing unconnected states	175

x	List of Algorithms
5.9	RemoveNonProgress: Removing states/actions 176
5.10	Planning for safe acyclic solutions by symbolic model checking 177
5.11	Guided planning for safe solutions 180
5.12	Planning for safe solutions by determinization 182
5.13	Transformation of a sequential plan into a corresponding policy 182
5.14	Interleaving planning and execution by lookahead 184
5.15	Online determinization planning and acting algorithm 186
5.16	MinMax Learning Real Time A* 186
5.17	Planning based on search automata 195
5.18	Associating states to contexts 196
6.1	A simple procedure to run a policy 221
6.2	Policy Iteration algorithm 225
6.3	Synchronous Value Iteration algorithm 226
6.4	Asynchronous Value Iteration algorithm 226
6.5	Bellman update procedure 227
6.6	A guaranteed approximation procedure for VI 229
6.7	Find&Revise schema 233
6.8	Best-first search algorithm AO* and LAO* 235
6.9	Bottom-up update for AO* 236
6.10	A “VI-like” update for LAO* 238
6.11	ILAO*, a variant of LAO* 239
6.12	A heuristic depth-first search algorithm for SSPs 239
6.13	Solved-SCC: labeling solved strongly connected components 240
6.14	Algorithm LDFS _a 241
6.15	Acting with the guidance of lookahead search 245
6.16	A determinization planning algorithm 246
6.17	Algorithm LRTDP 248
6.18	Procedure to check and label solve states for LRTDP 248
6.19	Sampling lookahead Tree to Estimate 251
6.20	A recursive UCT procedure 252
6.21	Compression framework for sparse probabilistic domains 262
7.1	Monitoring of the progression of a plan 283
7.2	Q-learning, a reinforcement learning algorithm 289
A.1	Equivalent iterative and recursive versions of nondeterministic search 314
A.2	A deterministic counterpart to Nondeterministic-Search 315
A.3	A generic nondeterministic And/Or search algorithm 316
B.1	Tarjan’s algorithm for finding strongly connected components 319

Table of Notation

Notation	Meaning
a, A	action, set of actions
α, \mathcal{A}	action template, set of action templates
$\text{cost}(a), \text{cost}(s, a)$	cost of a , cost of a in state s
$\text{cost}(\pi), \text{cost}(s, \pi)$	cost of π , cost of π in state s
$\text{Dom}(f), \text{Dom}(\pi)$	domain of a function or plan
$\text{eff}(a)$	effects of action a
g, S_g	goal conditions, goal states
$\gamma(s, a)$	progression, i.e., predicted result of applying a in s
$\gamma^{-1}(g, a)$	regression, i.e., conditions needed for a to produce s
$\widehat{\gamma}(s_0, \pi)$	{all states reachable from s_0 using π }, if π is a policy
$\widehat{\gamma}(s_0, \pi)$	sequence of states π produces from s_0 , if π is a plan
h	heuristic function
m, M	refinement method, set of methods
$\text{head}(a)$	a 's name and argument list
$\text{Pr}(s' s, a)$	probability of transition to s' if a is used in s
P, \mathcal{P}	planning problem
π, Π	plan or policy, set of plans or policies
$\text{pre}(a)$	preconditions of action a
$\text{Range}(v)$	range of a function or variable
s, S	predicted state, set of states for the planner
ξ, Ξ	actual state, set of states for the actor
s_0, S_0	initial state, set of initial states
$\Sigma = (S, A, \gamma)$	planning domain
$\Sigma = (B, R, X, \mathcal{I}, \mathcal{A})$	state-variable representation of a planning domain:
$a.\pi, \pi.a, \pi.\pi'$	concatenation of actions and plans
(a, b, \dots, u)	tuple
$\langle a, b, \dots, u \rangle$	sequence
$s \models g, s \not\models g$	s satisfies g , s does not satisfy g

Foreword

Over ten years ago, Malik Ghallab, Dana Nau, and Paolo Traverso gave us the first – and to date only – comprehensive textbook dedicated to the field of Automated Planning, providing a much needed resource for students, researchers and practitioners. Since then, this rich field has continued to evolve rapidly. There is now a unified understanding of what once seemed disparate work on classical planning. Models and methods to deal with time, resources, continuous change, multiple agents, and uncertainty have substantially matured. Cross-fertilization with other fields such as software verification, optimization, machine learning, and robotics has become the rule rather than the exception. A phenomenal range of applications could soon be within reach – given the right future emphasis for the field.

Today, the authors are back with a new book, *Automated Planning and Acting*. As the title indicates, this is not a mere second edition of the older book. In line with the authors’ analysis of where the future emphasis should lie for the field to realize its full impact, the book covers deliberative computational techniques for both planning *and acting*, that is, for deciding *which* actions to perform and also *how* to perform them. *Automated Planning and Acting* is more than a graduate textbook or a reference book. Not only do the authors outstandingly discharge their duties of educating the reader about the basics and much of the recent progress in the field, but they also propose a new framework from which the community can start to intensify research on deliberative acting and its integration with planning.

These aims are reflected in the book’s content. The authors put the integration of planning and acting at the forefront by dedicating an entire chapter to a unified hierarchical model and refinement procedures that suit the needs of both planning and acting functions. Each chapter devoted to a particular class of representations also includes significant material on the integration of planning and acting using these representations. Overall, the book is more focused than its predecessor, and explores in even greater depth models and approaches motivated by the needs of planning and acting in the real world, such as handling time and uncertainty. At the same time, the authors successfully balance breadth and depth by providing an elegant, concise synthesis of a larger body of work than in their earlier text.

There is no doubt that *Automated Planning and Acting* will be the text I require my students to read when they first start, and the goto book on my shelf for my own reference. As a timely source of motivation for game-changing research on the integration of planning and acting, it will also help shape the field for the next decade.

Sylvie Thiébaux
The Australian National University

Preface

This book is about methods and techniques that a computational agent can use for deliberative planning and acting, that is, for deciding both *which* actions to perform and *how* to perform them, to achieve some objective. The study of deliberation has several scientific and engineering motivations.

Understanding deliberation is an objective for most cognitive sciences. In artificial intelligence research, this is done by modeling deliberation through computational approaches both to enable it and to allow it to be explained. Furthermore, the investigated capabilities are better understood by mapping concepts and theories into designed systems and experiments to test empirically, measure, and qualify the proposed models.

The engineering motivation for studying deliberation is to build systems that exhibit deliberation capabilities and develop technologies that address socially useful needs. A technological system needs deliberation capabilities if it must autonomously perform a set of tasks that are too diverse – or must be done in environments that are too diverse – to engineer those tasks into innate behaviors. Autonomy and diversity of tasks and environments is a critical feature in many applications, including robotics (e.g., service and personal robots; rescue and exploration robots; autonomous space stations, satellites, or vehicles), complex simulation systems (e.g., tutoring, training or entertainment), or complex infrastructure management (e.g., industrial or energy plants, transportation networks, urban facilities).

MOTIVATION AND COVERAGE

The coverage of this book derives from the view we advocated for in our previous work [231], which we now briefly summarize.

Automated planning is a rich technical field, which benefits from the work of an active and growing research community. Some areas in this field are extensively explored and correspond to a number of already mature techniques. However, there are other areas in which further investigation is critically needed if automated planning is to have a wider impact on a broader set of applications. One of the most important

such areas, in our view is the *integration of planning and acting*. This book covers several different kinds of models and approaches – deterministic, hierarchical, temporal, nondeterministic and probabilistic – and for each of them, we discuss not only the techniques themselves but also how to use them in the integration of planning and acting.

The published literature on automated planning is large, and it is not feasible to cover all of it in detail in a single book. Hence our choice of what to cover was motivated by putting the integration of planning and acting at the forefront. The bulk of research on automated planning is focused on a restricted form called classical planning, an understanding of which is prerequisite introductory material, and we cover it in part of Chapter 2. But we have devoted large parts of the book to extended classes of automated planning and acting that relax the various restrictions required by classical planning.

There are several other kind of deliberation functions, such as monitoring, reasoning about one’s goals, reasoning about sensing and information-gathering actions, and learning and otherwise acquiring deliberation models. Although these are not our focus, we cover them briefly in Chapter 7.

The technical material in this book is illustrated with examples inspired from concrete applications. However, most of the technical material is theoretical. Case studies and application-oriented work would certainly enrich the integration of planning and acting view developed in here. We plan to devote a forthcoming volume to automated planning and acting applications.

USING THIS BOOK

This work started as a textbook project, to update our previous textbook on automated planning [230]. Our analysis of the state of the art led us quickly to embrace the objective of covering planning and acting and their integration and, consequently, to face two obstacles:

- The first problem was how to cover a domain whose scope is not easily amenable to a sharp definition and that requires integrating conceptually heterogenous models and approaches. In contrast to our previous book, which was focused on planning, this one proved harder to converge into a reasonably united perspective.
- The second problem was how to combine a textbook approach, i.e., a coherent synthesis of the state of the art, with the development of new material. Most of this new material is presented in comprehensive detail (e.g., in Chapter 3) consistent with a textbook use. In a few parts (e.g., Section 4.5.3), this new material is in preliminary form and serves as an invitation for further research.

This book can be used as a graduate-level textbook and as an information source for scientists and professionals in the field. We assume the reader to be familiar with the basic concepts of algorithms and data structures at the level that one might get in an undergraduate-level computer science curriculum. Prior knowledge of heuristic search techniques would also be helpful, but is not strictly necessary because the appendices provide overviews of needed tools.

A complete set of lecture slides for this book and other auxiliary materials are available online.¹

ACKNOWLEDGMENTS

We are thankful to several friends and colleagues who gave us very valuable feedback on parts of this book. Among these are Hector Geffner, Robert Goldman, Patrik Haslum, Joachim Hertzberg, Jörg Hoffmann, Felix Ingrand, Ugur Kuter, Marco Pistore, Mak Roberts, Vikas Shivashankar, Sylvie Thiébaux, and Qiang Yang.

We also wish to acknowledge the support of our respective organizations, which provided the support and facilities that helped to make this work possible: LAAS-CNRS in Toulouse, France; the University of Maryland in College Park, Maryland; and FBK, ICT-IRST in Trento, Italy. Dana Nau thanks ONR for their support of his planning work, and the students who took courses from rough drafts of this book.

Finally, we wish to acknowledge the support of our families, who remained patient during a project that consumed much more of our time and attention than we had originally anticipated.

¹ <http://www.laas.fr/planning>