

Contents

	Preface	xi
1	Planning and Autonomous Behavior	1
1.1	Autonomous Behavior: Hardwired, Learned, and Model-based	1
1.2	Planning Models and Languages	3
1.3	Generality, Complexity, and Scalability	6
1.4	Examples	8
1.5	Generalized Planning: Plans vs. General Strategies	11
1.6	History	12
2	Classical Planning: Full Information and Deterministic Actions	15
2.1	Classical Planning Model	15
2.2	Classical Planning as Path Finding	16
2.3	Search Algorithms: Blind and Heuristic	16
2.4	Online Search: Thinking and Acting Interleaved	20
2.5	Where do Heuristics come from?	23
2.6	Languages for Classical Planning	24
2.7	Domain-Independent Heuristics and Relaxations	27
2.8	Heuristic Search Planning	33
2.9	Decomposition and Goal Serialization	34
2.10	Structure, Width, and Complexity	34
3	Classical Planning: Variations and Extensions	37
3.1	Relaxed Plans and Helpful Actions	37
3.2	Multi-Queue Best-First Search	38
3.3	Implicit Subgoals: Landmarks	38
3.4	State-of-the-Art Classical Planners	39
3.5	Optimal Planning and Admissible Heuristics	41
3.6	Branching Schemes and Problem Spaces	42
3.7	Regression Planning	43
3.8	Planning as SAT and Constraint Satisfaction	45
3.9	Partial-Order Causal Link Planning	46
3.10	Cost, Metric, and Temporal Planning	47
3.11	Hierarchical Task Networks	49

4	Beyond Classical Planning: Transformations	51
4.1	Soft Goals and Rewards	51
4.2	Incomplete Information	53
4.3	Plan and Goal Recognition	57
4.4	Finite-State Controllers	60
4.5	Temporally Extended Goals	62
5	Planning with Sensing: Logical Models	65
5.1	Model and Language	65
5.2	Solutions and Solution Forms	67
5.3	Offline Solution Methods	69
5.4	Online Solution Methods	72
5.5	Belief Tracking: Width and Complexity	73
5.6	Strong vs. Strong Cyclic Solutions	76
6	MDP Planning: Stochastic Actions and Full Feedback	79
6.1	Goal, Shortest-Path, and Discounted Models	79
6.2	Dynamic Programming Algorithms	84
6.3	Heuristic Search Algorithms	86
6.4	Online MDP Planning	92
6.5	Reinforcement Learning, Model-based RL, and Planning	95
7	POMDP Planning: Stochastic Actions and Partial Feedback	97
7.1	Goal, Shortest-Path, and Discounted POMDPs	97
7.2	Exact Offline Algorithms	99
7.3	Approximate and Online Algorithms	102
7.4	Belief Tracking in POMDPs	105
7.5	Other MDP and POMDP Solution Methods	107
8	Discussion	109
8.1	Challenges and Open Problems	109
8.2	Planning, Scalability, and Cognition	111
	Bibliography	113
	Author's Biography	129