

Contents

<i>Credits and Acknowledgments</i>	page xv
<i>Introduction</i>	xvii
1 Distributed Constraint Satisfaction	1
1.1 Defining distributed constraint satisfaction problems	2
1.2 Domain-pruning algorithms	4
1.3 Heuristic search algorithms	8
1.3.1 The asynchronous backtracking algorithm	9
1.3.2 A simple example	11
1.3.3 An extended example: the four queens problem	13
1.3.4 Beyond the ABT algorithm	15
1.4 History and references	17
2 Distributed Optimization	19
2.1 Distributed dynamic programming for path planning	19
2.1.1 Asynchronous dynamic programming	19
2.1.2 Learning real-time A*	21
2.2 Action selection in multiagent MDPs	23
2.3 Negotiation, auctions, and optimization	27
2.3.1 From contract nets to auction-like optimization	27
2.3.2 The assignment problem and linear programming	29
2.3.3 The scheduling problem and integer programming	36
2.4 Social laws and conventions	43
2.5 History and references	45
3 Introduction to Noncooperative Game Theory: Games in Normal Form	47
3.1 Self-interested agents	47
3.1.1 Example: friends and enemies	48
3.1.2 Preferences and utility	49
3.2 Games in normal form	53
3.2.1 Example: the TCP user's game	54
3.2.2 Definition of games in normal form	55
3.2.3 More examples of normal-form games	56
3.2.4 Strategies in normal-form games	58
3.3 Analyzing games: from optimality to equilibrium	60

3.3.1	Pareto optimality	60
3.3.2	Defining best response and Nash equilibrium	61
3.3.3	Finding Nash equilibria	62
3.3.4	Nash's theorem: proving the existence of Nash equilibria	64
3.4	Further solution concepts for normal-form games	71
3.4.1	Maxmin and minmax strategies	72
3.4.2	Minimax regret	75
3.4.3	Removal of dominated strategies	77
3.4.4	Rationalizability	79
3.4.5	Correlated equilibrium	81
3.4.6	Trembling-hand perfect equilibrium	83
3.4.7	ϵ -Nash equilibrium	83
3.5	History and references	85
4	Computing Solution Concepts of Normal-Form Games	87
4.1	Computing Nash equilibria of two-player, zero-sum games	87
4.2	Computing Nash equilibria of two-player, general-sum games	89
4.2.1	Complexity of computing a sample Nash equilibrium	89
4.2.2	An LCP formulation and the Lemke–Howson algorithm	91
4.2.3	Searching the space of supports	99
4.2.4	Beyond sample equilibrium computation	101
4.3	Computing Nash equilibria of n -player, general-sum games	102
4.4	Computing maxmin and minmax strategies for two-player, general-sum games	105
4.5	Identifying dominated strategies	106
4.5.1	Domination by a pure strategy	106
4.5.2	Domination by a mixed strategy	107
4.5.3	Iterated dominance	109
4.6	Computing correlated equilibria	110
4.7	History and references	111
5	Games with Sequential Actions: Reasoning and Computing with the Extensive Form	113
5.1	Perfect-information extensive-form games	113
5.1.1	Definition	113
5.1.2	Strategies and equilibria	115
5.1.3	Subgame-perfect equilibrium	117
5.1.4	Computing equilibria: backward induction	119
5.2	Imperfect-information extensive-form games	125
5.2.1	Definition	125
5.2.2	Strategies and equilibria	126
5.2.3	Computing equilibria: the sequence form	129
5.2.4	Sequential equilibrium	136
5.3	History and references	139

6	Richer Representations: Beyond the Normal and Extensive Forms	141
6.1	Repeated games	142
6.1.1	Finitely repeated games	143
6.1.2	Infinitely repeated games	144
6.1.3	“Bounded rationality”: repeated games played by automata	147
6.2	Stochastic games	153
6.2.1	Definition	153
6.2.2	Strategies and equilibria	154
6.2.3	Computing equilibria	155
6.3	Bayesian games	156
6.3.1	Definition	157
6.3.2	Strategies and equilibria	160
6.3.3	Computing equilibria	163
6.3.4	<i>Ex post</i> equilibrium	165
6.4	Congestion games	166
6.4.1	Definition	166
6.4.2	Computing equilibria	167
6.4.3	Potential games	168
6.4.4	Nonatomic congestion games	170
6.4.5	Selfish routing and the price of anarchy	172
6.5	Computationally motivated compact representations	176
6.5.1	The expected utility problem	177
6.5.2	Graphical games	179
6.5.3	Action-graph games	181
6.5.4	Multiagent influence diagrams	183
6.5.5	GALA	186
6.6	History and references	187
7	Learning and Teaching	189
7.1	Why the subject of “learning” is complex	189
7.1.1	The interaction between learning and teaching	189
7.1.2	What constitutes learning?	190
7.1.3	If learning is the answer, what is the question?	192
7.2	Fictitious play	195
7.3	Rational learning	200
7.4	Reinforcement learning	204
7.4.1	Learning in unknown MDPs	205
7.4.2	Reinforcement learning in zero-sum stochastic games	205
7.4.3	Beyond zero-sum stochastic games	208
7.4.4	Belief-based reinforcement learning	209
7.5	No-regret learning and universal consistency	209
7.6	Targeted learning	211
7.7	Evolutionary learning and other large-population models	212
7.7.1	The replicator dynamic	213

7.7.2	Evolutionarily stable strategies	216
7.7.3	Agent-based simulation and emergent conventions	219
7.8	History and references	221
8	Communication	223
8.1	“Doing by talking” I: cheap talk	223
8.2	“Talking by doing”: signaling games	227
8.3	“Doing by talking” II: speech-act theory	229
8.3.1	Speech acts	229
8.3.2	Rules of conversation	230
8.3.3	A game-theoretic view of speech acts	232
8.3.4	Applications	235
8.4	History and references	238
9	Aggregating Preferences: Social Choice	241
9.1	Introduction	241
9.1.1	Example: plurality voting	241
9.2	A formal model	242
9.3	Voting	244
9.3.1	Voting methods	244
9.3.2	Voting paradoxes	246
9.4	Existence of social functions	247
9.4.1	Social welfare functions	248
9.4.2	Social choice functions	251
9.5	Ranking systems	255
9.6	History and references	258
10	Protocols for Strategic Agents: Mechanism Design	261
10.1	Introduction	261
10.1.1	Example: strategic voting	261
10.1.2	Example: buying a shortest path	262
10.2	Mechanism design with unrestricted preferences	263
10.2.1	Implementation	264
10.2.2	The revelation principle	265
10.2.3	Impossibility of general, dominant-strategy implementation	267
10.3	Quasilinear preferences	268
10.3.1	Risk attitudes	269
10.3.2	Mechanism design in the quasilinear setting	271
10.4	Efficient mechanisms	276
10.4.1	Groves mechanisms	276
10.4.2	The VCG mechanism	280
10.4.3	VCG and individual rationality	282
10.4.4	VCG and weak budget balance	283
10.4.5	Drawbacks of VCG	284
10.4.6	Budget balance and efficiency	288

10.4.7	The AGV mechanism	288
10.5	Beyond efficiency	289
10.5.1	What else can be implemented in dominant strategies?	290
10.5.2	Tractable Groves mechanisms	292
10.6	Computational applications of mechanism design	293
10.6.1	Task scheduling	294
10.6.2	Bandwidth allocation in computer networks	296
10.6.3	Multicast cost sharing	298
10.6.4	Two-sided matching	301
10.7	Constrained mechanism design	307
10.7.1	Contracts	308
10.7.2	Bribes	309
10.7.3	Mediators	310
10.8	History and references	311
11	Protocols for Multiagent Resource Allocation: Auctions	315
11.1	Single-good auctions	315
11.1.1	Canonical auction families	316
11.1.2	Auctions as Bayesian mechanisms	318
11.1.3	Second-price, Japanese, and English auctions	319
11.1.4	First-price and Dutch auctions	321
11.1.5	Revenue equivalence	323
11.1.6	Risk attitudes	326
11.1.7	Auction variations	327
11.1.8	"Optimal" (revenue-maximizing) auctions	328
11.1.9	Collusion	330
11.1.10	Interdependent values	333
11.2	Multiunit auctions	336
11.2.1	Canonical auction families	336
11.2.2	Single-unit demand	337
11.2.3	Beyond single-unit demand	340
11.2.4	Unlimited supply: random sampling auctions	342
11.2.5	Position auctions	344
11.3	Combinatorial auctions	346
11.3.1	Simple combinatorial auction mechanisms	348
11.3.2	The winner determination problem	349
11.3.3	Expressing a bid: bidding languages	352
11.3.4	Iterative mechanisms	357
11.3.5	A tractable mechanism	359
11.4	Exchanges	361
11.4.1	Two-sided auctions	361
11.4.2	Prediction markets	362
11.5	History and references	364

12 Teams of Selfish Agents: An Introduction to Coalitional Game Theory	367
12.1 Coalitional games with transferable utility	367
12.1.1 Definition	368
12.1.2 Examples	368
12.1.3 Classes of coalitional games	370
12.2 Analyzing coalitional games	371
12.2.1 The Shapley value	372
12.2.2 The core	374
12.2.3 Refining the core: ϵ -core, least core, and nucleolus	378
12.3 Compact representations of coalitional games	381
12.3.1 Weighted majority games and weighted voting games	381
12.3.2 Weighted graph games	382
12.3.3 Capturing synergies: a representation for superadditive games	384
12.3.4 A decomposition approach: multi-issue representation	385
12.3.5 A logical approach: marginal contribution nets	386
12.4 Further directions	388
12.4.1 Alternative coalitional game models	388
12.4.2 Advanced solution concepts	389
12.5 History and references	390
13 Logics of Knowledge and Belief	393
13.1 The partition model of knowledge	393
13.1.1 Muddy children and warring generals	393
13.1.2 Formalizing intuitions about the partition model	394
13.2 A detour to modal logic	396
13.2.1 Syntax	398
13.2.2 Semantics	398
13.2.3 Axiomatics	399
13.2.4 Modal logics with multiple modal operators	399
13.2.5 Remarks about first-order modal logic	400
13.3 S5: An axiomatic theory of the partition model	400
13.4 Common knowledge, and an application to distributed systems	403
13.5 Doing time, and an application to robotics	406
13.5.1 Termination conditions for motion planning	407
13.5.2 Coordinating robots	410
13.6 From knowledge to belief	412
13.7 Combining knowledge and belief (and revisiting knowledge)	413
13.8 History and references	418
14 Beyond Belief: Probability, Dynamics, and Intention	421
14.1 Knowledge and probability	421
14.2 Dynamics of knowledge and belief	425
14.2.1 Belief revision	426

14.2.2	Beyond AGM: update, arbitration, fusion, and friends	430
14.2.3	Theories of belief change: a summary	436
14.3	Logic, games, and coalition logic	436
14.4	Towards a logic of "intention"	438
14.4.1	Some preformal intuitions	438
14.4.2	The road to hell: elements of a formal theory of intention	440
14.4.3	Group intentions	443
14.5	History and references	445
Appendices: Technical Background		447
A	Probability Theory	449
A.1	Probabilistic models	449
A.2	Axioms of probability theory	449
A.3	Marginal probabilities	450
A.4	Conditional probabilities	450
B	Linear and Integer Programming	451
B.1	Linear programs	451
B.2	Integer programs	453
C	Markov Decision Problems (MDPs)	455
C.1	The model	455
C.2	Solving known MDPs via value iteration	455
D	Classical Logic	457
D.1	Propositional calculus	457
D.2	First-order logic	458
<i>Bibliography</i>		459
<i>Index</i>		473