

Index

- = (equals), 532
- A^* search, 89
- \wedge (and), 163
- \Leftarrow (base-level if), 581
- \Leftarrow (if), 163, 496
- \models (entails), 160, 499
- \models (true in), 222
- \neq (not equal to), 535
- ϕ (denotation of terms), 496, 514
- π (denotation of predicate symbols), 496
- π (meaning of atoms), 159
- \mapsto (rewritten as), 521
- \vdash (derive), 167
- \vee (or), 585
- % (comment), 496
- & (base-level and), 581

- abduction, 199, 200
- abductive diagnosis, 201
- abilities, 11
- absolute error, 290
- absorbing state, 399
- abstractions, 15
- achievement goal, 25, 355
- action, 45, 350
- action constraints, 361
- action features, 360
- action function, 74
- action profile, 425, 435
- activation function, 306
- active learning, 285
- active sensor, 64
- acts, 4
- actuator, 44, 45
- acyclic, 184, 195
- acyclic directed graph, 75
- add list, 354

- additive independence, 378
- additive utility, 378
- admissibility, 91
- aerodynamics, 10
- agent, 4, 10, 43, 45
 - embedded, 59
 - purposive, 44
- agent system, 45
- agent system model, 59
- AI, 3
- algebra, 222
- algebraic variable, 113
- Allais Paradox, 380
- alpha-beta (α - β) pruning, 431
- alternative, 615
- analysis, 4
- Analytical Engine, 7
- ancestor, 184
- annealing schedule, 138
- answer, 166, 171, 187, 504
- answer clause, 171
- answer extraction, 504, 509
- anytime algorithm, 26
- application of substitution, 506
- approximately correct, 332
- approximately optimal, 14
- approximately optimal solution, 14
- arc, 75
- arc consistent, 121
- argument, 197, 494
- Aristotelian definition, 567
- array, 634
- Arrow's impossibility theorem, 446
- artificial intelligence, 3
- ask, 161, 166
- ask-the-user, 175, 576
- askable, 175

- assertional knowledge base, 568
- assignment space, 118
- assumable, 187, 200
- assumption-based truth maintenance system, 207
- asymptotic complexity, 83
- asynchronous value iteration, 406
- ATMS, 207
- atom, 158, 163
- atomic clause, 163, 496
- atomic proposition, 158, 163
- atomic symbol, 494
- attribute, 552
- auction, 448
- autonomous delivery robot, 30
- average reward, 402
- axiom, 160, 161, 501
- axiom schema, 533
- axiomatizing, 161, 501
- axioms
 - of probability, 224
 - of rationality, 373

- back-propagation learning, 317
- background knowledge, 17, 607
- background variables, 206
- backtracking, 80, 120
- backward induction, 430
- bagging, 319
- base language, 580
- base level, 580
- base-level algorithms, 319
- Basic Formal Ontology, 573
- Bayes' rule, 229
- Bayesian classifier, 309
- Bayesian learning, 334
- Bayesian network, *see* belief network
- Bayesian probability, 220
- beam search, 141
- belief, 49
- belief monitoring, 271
- belief network, 235, 236, 458
 - causality, 241
 - inference problem, 252
- belief state, 48, 61
- belief state transition function, 49
- best response, 436
- best-first search, 88
- beta distribution, 337
- BFO, 573
- bias, 286, 321, 331
- bias-free, 331
- bidirectional search, 101
- binary constraint, 116
- binary feature, 112
- biology, 6
- bit, 231
- blame attribution problem, 465
- body, 44
 - agent, 45
 - rule, 163, 496
- Boltzmann distribution, 142, 473
- Boolean property, 553
- Boolean variable, 113, 126, 158

- boosting, 320
- bottom-up proof procedure, 167
- boundary of a process, 576
- boundary of an object, 575
- bounded rationality, 26
- branch-and-bound, 98
- branching factor, 77
- breadth-first search, 84

- candidate elimination algorithm, 329
- canonical representation, 534
- cardinal, 14
- cardinal preference, 25
- case analysis, 125
- case-based reasoning, 324
- causal, 48
- causal link, 364
- causal mechanism, 206
- causal model, 204, 206
- causal network, 241
- causal rule, 353, 601
- causal transduction, 48
- causality, 204, 241
- central limit theorem, 223
- chain rule, 227
- chance node, 384
- characteristic function, 559
- children, 142
- choice space, 615
- choose, 170
- Church-Turing thesis, 7
- clarity principle, 114, 241
- Clark normal form, 538
- Clark's completion, 194, 538
- class, 298, 309, 451, 559, 568
- classification, 284, 298
- classification tree, 298
- clause, 126
 - definite, 163, 494, 496
 - Horn, 185
- closed list, 94
- closed-world assumption, 193
- cluster, 451
- clustering, 451
- cognitive science, 10
- command, 45
- command function, 49
- command trace, 46
- commonsense reasoning, 13
- competitive, 424
- complements, 381
- complete, 131, 167
 - bottom-up derivation, 169
- complete knowledge assumption, 193
- completeness
 - of preferences, 373
- complex preference, 25
- complexity, 83
- compositional measure of belief, 226
- compound proposition, 158
- computational, 4
- computational learning theory, 332
- computational limits dimension, 26
- computational linguistics, 520

- concept, 572
- conceptualization, 494, 563
- conditional effect, 355
- conditional
 - expected value, 231
- conditional probability, 225
- conditional probability distribution, 227, 297
- conditionally independent, 233
- Condorcet paradox, 445
- conflict, 132, 187
- conjunction, 158
- consequence set, 167
- consistency-based diagnosis, 187, 189
- consistent, 328, 358, 608
- constant, 494
- constrained optimization problem, 144
- constraint, 115
- constraint network, 121
- constraint optimization problem, 145
- constraint satisfaction problem, 117
- context-free grammar, 521
- context-specific independence, 250
- contingent attribute, 562
- continuant, 573
- continuous variable, 113
- controller, 45, 48, 426
- cooperate, 444
- cooperative, 424
- cooperative system, 196
- coordinate, 444
- coordination, 437
- cost, 76
- credit-assignment problem, 285
- cross validation, 324, 325
- crossover, 142
- CSP, *see* constraint satisfaction problem
- culture, 6
- cumulative probability distribution, 257
- cumulative reward, 401
- cyc, 564
- cycle, 75
- cycle check, 93
- cyclic, 184

- DAG, 75
- Datalog, 494
- datatype property, 568
- DBN, *see* dynamic belief network
- DCG, *see* definite clause grammar
- dead reckoning, 58
- debugging, 179
 - incorrect answers, 180
 - infinite loops, 184
 - missing answers, 182
- decision
 - sequential, 386
 - single, 384
- decision function, 391
- decision network, 384, 387, 428
- decision node, 384
- decision tree, 298, 382, 426
 - learning, 298, 321
- decision tree learning, 232
- decision variable, 382

- decision-theoretic planning, 409
- deduction, 167, 200
- Deep Space One, 212
- default, 176, 196
- definite clause, 163, 494, 496
- definite clause grammar (DCG), 522
- definite clause resolution, 169
- delay, 536, 590
- delete list, 354
- delivery robot, 30
- DENDRAL, 9
- denote, 496
- dense, 46
- dependent continuant, 573, 575
- depth-bounded meta-interpreter, 587
- depth-first search, 80
- derivation, 167, 171, 510
- derived, 167, 204, 354, 600
 - knowledge, 558
- description logic, 568
- design, 4, 202
- design space, 19
- design time reasoning, 17
- desire, 49
- deterministic, 24
- diagnosis, 201
 - abductive, 201
 - consistency-based, 187
 - decision theoretic, 383
 - probabilistic, 245
- diagnostic assistant, 30, 33
- dictator, 447
- difference list, 522
- differentia, 567
- dimension
 - computational limits, 26
 - effect uncertainty, 24
 - learning, 26
 - modularity, 19
 - number of agents, 25
 - planning horizon, 22
 - preference, 24
 - representation scheme, 20
 - sensing uncertainty, 23
- directed acyclic graph, 75
- directed graph, 75
- Dirichlet distribution, 337
- discount factor, 402
- discounted reward, 402
- discrete, 46
- discrete variable, 113
- disjunction, 158, 185, 585
- disjunctive normal form, 189
- disposition, 575
- distribution
 - normal, 223
- do, 598
- domain, 112, 113, 496, 552, 633
- domain consistent, 121
- domain expert, 63
- domain ontology, 571
- domain splitting, 125
- dominant strategy, 446

- dominates, 439
- don't-care non-determinism, 170
- don't-know non-determinism, 170
- dot product, 271
- DPLL, 127
- dynamic, 599
- dynamic belief network, 272
- dynamic decision network, 409
- dynamic programming, 103

- economically efficient, 446
- effect, 350, 354
- effect constraints, 361
- effect uncertainty dimension, 24
- effectively computable, 7
- effector, 44
- efficient, 446
- eligibility trace, 479
- elimination ordering, 252
- EM, 452, 455, 460
- embedded agent, 59
- embodied, 44
- empirical frequency, 295
- endogenous variables, 206
- engineering goal, 4
- ensemble learning, 319
- entails, 160
- entity, 573
- entropy, 232, 292
- environment, 10
- epistemological, 221
- epistemology, 9
- equal, 532
- error, 288
 - absolute, 290
 - sum squares, 290
 - worst case, 291
- error of hypothesis, 332
- Euclidean distance, 95, 325
- evaluation, 145
- evaluation function, 132, 433
- event calculus, 604
- evidence, 225
- evidential model, 205
- evolutionary algorithm, 466
- exclusive-or, 308
- existence uncertainty, 619
- existentially quantified variable, 500
- exogenous variables, 206
- expanding, 78
- expectation maximization (EM), 452
- expected monetary value, 376
- expected utility, 386, 392
- expected value, 230
 - of utility of a policy, 392
- experience, 468
- expert opinion, 297
- expert system, 9, 10, 61
- explained away, 201, 243, 245
- explanation, 201
- exploit, 472
- explore, 472
- expression, 496
- extensional, 559
- extensionally, 116
- extensive form, 426
- external knowledge source, 65
- extrapolation, 286

- $f(p)$, 89
- factor, 249, 634
- factored finite state machine, 50
- factored representation, 50
- factorization, 236
- fail, 170
- failing naturally, 96
- failing unnaturally, 96
- failure, 197
- fair, 170, 518
- fairness, 19
- false, 159, 498
- false-negative error, 14, 182, 292
- false-positive error, 14, 180, 292, 293
- fault, 187
- feature engineering, 484
- feature-based representation of actions, 353
- features, 21, 112
- feed-forward neural network, 316
- fiat part of an object, 575
- fiat part of process, 576
- filler, 556
- filtering, 58, 267, 271
- finite failure, 199
- finite horizon, 23
- finite state controller, 50
- finite state machine, 50
- first-order predicate calculus, 517
- fixed point, 169
- flat, 20
- floundering goal, 541
- fluent, 599
- flying machines, 9
- for all, 500
- forward chaining, 167
- forward planner, 356
- frame, 556
- frame problem, 618
- frame rule, 353, 601
- framing effect, 380
- frequentist, 220
- fringe, 77
- frontier, 77
- fully observable, 23
- fully observable dynamic decision network, 411
- fully observable Markov decision process, 401
- function, 575, 633
 - symbol, 514
- fuzzy terms, 52

- gambling, 220
- game tree, 426
- Gaussian distribution, 223
- general boundary, 329
- M sized additive independence, 381
- generalize, 285
- generalized answer clause, 509

- generalized arc consistency algorithm, 122
- generate and test, 118, 146
- genetic algorithms, 142
- genus, 567
- Gibbs distribution, 142, 473
- Gini index, 302, 345
- global minimum, 149
- goal, 11, 25, 49, 171, 356
 - node, 75
- goal constraints, 361
- gradient descent, 149, 304
- grammar, 521
 - context-free, 521
 - definite clause, 522
- graph, 75
- greedy, 23, 609
- greedy ascent, 132
- greedy descent, 132
- ground, 496
- ground instance, 506, 507
- ground representation, 580

- $h(n)$, 87
- hard clustering, 452
- hard constraint, 111, 115
- head, 163, 496
- help system, 246, 312
- Herbrand interpretation, 508
- heuristic
 - function, 87
 - knowledge, 72
 - search, 87
- heuristic depth-first search, 88
- hidden Markov model (HMM), 267
- hidden variable, 244, 460
- hierarchical, 20
- hierarchical Bayesian model, 338
- hierarchical control, 50
- hill climbing, 132
- history, 10, 48, 468
- HMM, *see* hidden Markov model
- Hoeffding's inequality, 259
- horizon, 361
- Horn clause, 185
- how, 17
- how question, 177
- human-computer interaction (HCI), 578
- hyperparameters, 339
- hypothesis, 288
- hypothesis space, 288, 328

- identity uncertainty, 619
- imperfect information game, 428
- implication, 158
- incoming arc, 75
- inconsistent, 186
- incorrect answer, 180
- incremental gradient descent, 305
- indefinite horizon, 23, 399
- independent, 233
- independent and identically distributed (i.i.d.), 335
- independent choice logic (ICL), 615
- independent continuant, 573, 575
- indicator variable, 141, 290
- indifferent, 373
- individual, 22, 141, 492, 496, 568
- individual-property-value, 552
- induction, 200, 284
- inductive logic programming, 606
- inference, 167
- infinite horizon, 23, 399
- influence diagram, 387
- information content, 232
- information gain, 232, 302
- information set, 428
- information theory, 231, 323
- init, 598
- initial-state constraints, 361
- input features, 288
- insects, 18
- instance, 504, 506
 - ground, 507
- instance space, 328
- insurance, 377
- integrity constraint, 185
- intelligent action, 4
- intelligent tutoring system, 35
- intended interpretation, 161, 501
- intensional, 559
- intensionally, 116
- intention, 49
- interpolation, 286
- interpretation, 159, 496
- intersection, 636
- intervention, 204, 206, 241
- involve, 115
- island, 102
- island-driven search, 101
- iterative best improvement, 132
- iterative deepening, 95
 - A^* , 98

- Java, 495, 558
- join, 636
- joint probability distribution, 236, 252

- k -fold cross validation, 324
- k -means algorithm, 452
- kd -tree, 326
- kernel functions, 314
- knowledge, 12, 60
- knowledge base, 12, 17, 60, 160, 163, 494, 496
- knowledge engineers, 63
- knowledge is given, 26
- knowledge is learned, 26
- knowledge level, 16, 176
- knowledge-based system, 60
- knowledge-level debugging, 180

- landmarks, 52
- language, 521
 - natural, 520
- language bias, 331
- latent tree model, 314
- latent variable, 244, 309, 460

- learning, 61, 283–347, 441–445, 451–488, 606–611
 - Bayesian, 334
 - bias, 331
 - case-based, 324
 - decision trees, 232, 298
 - decision tree, 321
 - multiagent, 441
 - neural network, 315
 - PAC, 332, 340
 - relational, 606
 - supervised, 284
 - to coordinate, 441
 - unsupervised, 285, 451
 - version space, 329
- learning bias, 321
- learning dimension, 26
- learning rate, 304
- leave-one-out cross-validation, 325
- leaves, 75
- level of abstraction, 15
- life-long learning, 6
- likelihood, 229
- likelihood of the data, 292
- linear function, 304
- linear regression, 304, 484
- linearly separable, 308
- Linnaean taxonomy, 567
- lists, 516
- literal, 126, 194
- liveness, 18
- local minimum, 149
- local optimum, 132
- local search, 130, 131, 609
- localization, 268
- logic, 495
- logic program, 514
- logic programming, 207
- logical consequence, 160, 499
- logical variable, 494
- logically follows, 160
- logistic function, 307
- long-term memory, 61
- loop check, 93
- lottery, 373
- lowest-cost-first search, 86

- M features, 360
- M system, 51
- machine learning, *see* learning
- maintenance goal, 25, 355
- MAP model, 321
- mapping, 633
- margin, 314
- Markov assumption, 266
- Markov blanket, 235
- Markov chain, 266
- Markov decision process, 399, 463
- matrix multiplication, 271
- maximum a posteriori probability, 321
- maximum entropy, 233, 234
- maximum likelihood model, 292, 321
- maximum-likelihood estimate, 295
- MDL principle, 323
- MDP, *see* Markov decision process

- measure, 221
- mechanism, 424, 446
- mechanism design, 446
- meta-interpreter, 579, 580, 582
 - ask the user, 589
 - build proof tree, 587
 - built-in predicates, 586
 - delayed goals, 590
 - depth-bounded, 587
 - disjunction, 586
 - vanilla, 583
- meta-level, 580
- metalanguage, 580
- MGU, 507
- min-factor elimination ordering, 130
- minimal
 - conflict, 187
 - model, 169
- minimal diagnosis, 189
- minimal explanation, 201
- minimal model, 509
- minimax, 430
- minimum deficiency elimination ordering, 130
- minimum description length (MDL), 323
- minimum fixed point, 169
- missing at random, 461
- model, 15, 57, 116, 160
- modular, 20
- modularity, 19
- modularity dimension, 20
- modus ponens, 167
- money pump, 374
- monitoring, 267, 271
- monotone restriction, 95
- monotonic, 196
- more general hypothesis, 329
- more specific hypothesis, 329
- most general unifier, 507, 511
- multiagent decision network, 428
- multiple agent, 25
- multiple-path pruning, 93, 94
- mutex constraint, 361
- MYCIN, 9
- myopic, 23
- myopically optimal, 301

- N3, 555
- naive Bayesian classifier, 246, 310, 455
- Nash equilibrium, 436
- natural kind, 309, 313, 559
- natural language processing, 520
- nature, 44, 424
- nearest neighbors, 325
- negation, 158, 185
- negation as failure, 194, 199, 537
- negative examples, 607
- neighbor, 75, 131
- neural network, 315
- no, 166
- no-forgetting agent, 388
- no-forgetting decision network, 388
- node, 75
- noisy, 267
- noisy or, 250

- non-deterministic, 170
- non-deterministic procedure, 170
- non-ground representation, 580
- non-monotonic, 196
- non-parametric distribution, 223
- non-planning, 23
- non-serial dynamic programming, 151
- non-terminal symbol, 521
- nonlinear planning, 364
- normal distribution, 223
- normal-form game, 425
- Notation 3, 555
- NP, 170
- NP-complete, 170
- NP-hard, 170
- number of agents dimension, 25
- number uncertainty, 619

- object, 552, 575
- object aggregate, 575
- object language, 580
- object property, 568
- object-oriented languages, 495
- object-oriented programming (OOP) languages, 558
- objective function, 144
- observation, 11, 17, 174, 225
- occurrent, 573, 575
- occurs check, 518
- Ockham's razor, 287
- off policy, 470
- off-policy learner, 475
- offline, 61
- offline computation, 17
- offline learning, 285
- omniscient agent, 114
- on-policy learner, 475
- one-point crossover, 142
- online, 60, 64
- online computation, 17
- online learning, 285
- ontological, 221
- ontology, 61, 161, 175, 549, 563
- OOP, *see* object-oriented programming
- open-world assumption, 193
- optimal, 74
- optimal algorithm, 105
- optimal policy, 386, 392, 403
- optimal solution, 13, 76
- optimality criterion, 144
- optimization problem, 144
- oracle, 170
- orders of magnitude reasoning, 52
- ordinal, 13
- ordinal preference, 25
- organizations, 6
- outcome, 373, 382, 425
- outgoing arc, 75
- overfitting, 303
- OWL, 564, 568

- PAC learning, 332, 340
- pair, 633
- parametric distribution, 223
- parametrized random variable, 613

- paramodulation, 534
- parents, 235
- partial, 267
- partial evaluation, 590
- partial-order planning, 363
- partially observable, 23
- partially observable game, 428
- partially observable Markov decision process, 401, 411
- particle, 264
- particle filtering, 264, 265
- passive sensor, 64
- past experience, 11
- path, 75
 - consistency, 125
- payoff matrix, 426
- percept, 45
- percept stream, 46
- percept trace, 46
- perception, 58
- perceptron, 306
- perfect information, 430
- perfect information game, 426
- perfect rationality, 26
- philosophy, 9
- physical symbol system, 15
- physical symbol system hypothesis, 15, 319
- pixels, 45
- plan, 356
- planner, 356
- planning, 356–370, 604
 - as a CSP, 360
 - forward, 356
 - partial-order, 363
 - regression, 357
- planning horizon, 22
- planning horizon dimension, 23
- plate model, 338, 613
- point estimate, 288
- policy, 103, 386, 390
- policy iteration, 407
- policy search, 466
- POMDP, *see* partially observable Markov decision process
- population, 141, 264
- positive examples, 607
- possible, 358
- possible world, 113, 115, 382, 391, 615
- posterior probability, 225
- pragmatics, 521
- precision, 293
- precision-recall curve, 293
- precondition, 350, 353, 354, 601
- precondition constraints, 361
- predicate symbol, 494
- preference, 144
- preference bias, 331
- preference dimension, 25
- preference elicitation, 379
- primitive, 204, 354, 566, 600
 - knowledge, 558
- prior count, 296
- prior knowledge, 10

- prior probability, 225
- prisoner's dilemma, 437
- probabilistic independence, 233
- probabilistic inference, 248
- probabilistic relational model (PRM), 612
- probability, 219–275, 295
 - axioms, 224
 - conditional, 225
 - posterior, 225
 - prior, 225
 - semantics, 222
- probability density function, 223
- probability distribution, 222
- probability mass, 262
- probability measure, 221
- probable solution, 14
- probably approximately correct, 259, 332
- probably approximately correct learning, 332
- process, 23, 576
- process aggregate, 576
- processual context, 576
- processual entity, 576
- projection, 636
- Prolog, 9, 494, 496
- proof, 167
 - bottom-up, 167
 - top-down, 169
- proof procedure, 167
- proof tree, 179
- prop*, 552
- property, 552, 568
- property inheritance, 561
- proposal distribution, 260, 261
- proposition, 21, 158
- propositional calculus, 158
- propositional definite clause resolution, 169
- propositional definite clauses, 163
- propositional satisfiability, 126
- prospect theory, 380
- proved, 167
- pseudocount, 296
- psychology, 9
- punishment, 399
- pure strategy, 435
- pure symbol, 127
- purposive agent, 44

- Q^* , 404
- Q-learning, 469
- Q-value, 404, 469
- Q^π , 403
- qualitative derivatives, 52
- qualitative reasoning, 52
- quality, 575
- quantitative reasoning, 52
- query, 166, 494, 496
- querying the user, 175

- random initialization, 131
- random restart, 131
- random sampling, 132
- random variable, 221
- random walk, 132
- random worlds, 233, 234

- range, 552, 633
- rational, 376
- rational agent, 373
- RDF, 555, 559, 564
- RDF Schema, 559
- RDF-S, 564
- rdf:type*, 559
- rdfs:subClassOf*, 559
- reachable, 599
- realizable entity, 575
- reasoning, 17
- recall, 293
- recognition, 201
- record linkage, 619
- recursively enumerable, 517
- reflection, 579
- regression, 284, 304
- regression planning, 357, 358
- regression tree, 314
- reify, 552
- reinforcement learning, 285, 463
- rejection sampling, 259
- relation, 22, 492, 633, 635
- relational algebra, 635
- relational database, 635
- relational descriptions, 22
- relational probability model, 612
- relational uncertainty, 619
- representation, 12
- representation scheme, 12
- representation scheme dimension, 20, 22
- resampling, 264
- resolution, 169, 171
 - SLD, 169, 510
- resolvent, 171
- resource, 555, 564
- Resource Description Framework, 555, 564
- restriction bias, 331
- return, 469
- revelation principle, 447
- reward, 399
- rewrite rule, 521, 534
- risk averse, 377
- robot, 10, 30, 44
- ROC curve, 293
- role, 63, 575
- root, 75
- rule, 163, 494, 496
- rule of inference, 167
- run, 427
- run-time distribution, 138

- safety goal, 18, 355
- sample average, 258
- sample complexity, 333
- SARSA(λ), 479
- satisfiable, 201
- satisficing, 14
- satisficing solution, 14
- satisfies, 116
 - policy, 391
- satisfy, 115
- scenario, 201
- scheme, 115, 635

- scientific goal, 4
- scope, 115, 145, 634, 635
- search, 71–110
 - A^* , 89
 - best-first, 88
 - bidirectional, 101
 - breadth-first, 84
 - cycle-checking, 93
 - depth first, 80
 - dynamic programming, 103
 - gradient descent, 149
 - heuristic, 87
 - island driven, 101
 - iterative deepening, 95
 - local search, 130
 - lowest-cost-first, 86
- search and score, 462
- search bias, 331
- second-order logic, 517
- second-price auction, 448
- select, 170
- selection, 635
- selector function, 615
- semantic interoperability, 65, 550
- semantic network, 554
- semantic web, 555, 564
- semantics, 159, 520
 - propositional calculus, 159
 - variables, 496
- sensing uncertainty dimension, 23
- sensor, 44, 45, 64
- sensor fusion, 270
- separable control problem, 58
- sequential decision problem, 386
- sequential prisoner's dilemma, 438
- set, 633
- set difference, 636
- short-term memory, 61
- sigmoid, 307
- simulated agent, 59
- simulated annealing, 136
- single agent, 25
- single decision, 384
- single-stage decision network, 384
- site, 575
- situation, 598
- situation calculus, 598
- SLD derivation, 171, 510
- SLD resolution, 169, 510
- slot, 556
- Smalltalk, 558
- smoothing, 267, 271
- SNLP, 370
- social preference function, 446
- society, 6
- soft clustering, 452
- soft constraint, 111, 145
- soft-max, 473
- software agent, 10
- software engineer, 63
- software engineering, 13
- solution, 73, 75
- sound, 167
- spatial region, 573, 575
- spatio-temporal region, 576
- specialization operator, 609
- specific boundary, 329
- squashed linear function, 306
- stable, 453
- stack, 80
- stage, 23
- start node, 75
- start states, 74
- starvation, 170, 518
- state, 20, 72
- state constraints, 361
- state space, 72
- state-space graph, 351, 356
- state-space problem, 74
- static, 599
- stationary, 266, 399
- stationary policy, 403
- step function, 306
- step size, 150, 151
- stimuli, 45
- stochastic, 24
- stochastic beam search, 142
- stochastic local search, 134
- stochastic simulation, 256
- stochastic strategy, 435
- stopping state, 399
- strategic form of a game, 425
- strategically, 424
- strategy, 426, 427, 435
- strategy profile, 427, 435
- strictly dominated, 147
- strictly dominates, 440
- strictly preferred, 373
- STRIPS assumption, 354
- STRIPS representation, 354
- structure learning, 461
- subClassOf, 559
- subgame-perfect equilibrium, 439
- subgoal, 171
- subject, 552
- subjective probability, 220
- substitutes, 381
- substitution, 506
- sum-of-squares error, 290, 453
- supervised learning, 284, 288
- support set, 435, 439
- support vector machine (SVM), 314
- SVM, *see* support vector machine
- symbol, 15, 114
- symbol level, 17
- symbol system, 15
- syntax, 520
 - Datalog, 494
 - propositional definite clauses, 163
- synthesis, 4
- systematicity, 370
- tabu list, 135
- target features, 288
- TD error, 467
- tell, 161
- temperature, 473

- temporal difference error, 467
- temporal region, 575
- term, 494, 514
- terminal symbol, 521
- terminological knowledge base, 568
- test examples, 288
- theorem, 167
- there exists, 500
- thing, 573
- thinking, 4
- threat, 438
- time, 46, 598
- tit-for-tat, 438
- top-down proof procedure, 169
- total assignment, 115
- total reward, 402
- trading agent, 30, 37
- training examples, 284, 288, 328
- transduction, 48
- transitivity of preferences, 374
- tree, 75
- tree augmented naive Bayesian (TAN) network, 314
- treewidth, 130, 256
- triple, 552, 555, 606, 633
- triple representation, 552
- true, 159, 498, 633
- true-positive rate, 293
- truth maintenance system, 207
- truth value, 159
- truthful, 446
- try, 131
- tuple, 115, 633, 635
- Turing test, 5
- Turtle, 555
- tutoring system, 30
- two-step belief network, 272
- type, 559
- type*, 553, 559

- UML, 558
- unary constraint, 116
- unconditionally independent, 235
- unification, 511
- unifier, 507
- Uniform Resource Identifier, 555, 564, 568
- uninformed search strategies, 80
- union, 636
- unique names assumption, 535
- unit resolution, 127
- units, 315
- universally quantified, 498
- universally quantified variable, 498
- unsatisfiable, 186
- unstable, 319
- unsupervised learning, 285, 451

- URI, *see* Uniform Resource Identifier
- useful, 358
- user, 64, 114, 175
- utility, 14, 376
- utility node, 385

- V^π , 403
- V^* , 404
- validation set, 324
- value, 401
- value of control, 398
- value of information, 397
- vanilla meta-interpreter, 582, 583
- variable, 113
 - algebraic, 113
 - Boolean, 113
 - continuous, 113
 - decision, 382
 - discrete, 113
 - existentially quantified, 500
 - logical, 494
 - random, 221
 - universally quantified, 498
- variable assignment, 498
- variable elimination, 386
 - belief networks, 248
 - CSPs, 127
 - decision networks, 392
 - optimizing soft constraints, 147
- variational inference, 248
- VCG mechanism, 447
- VE, *see* variable elimination
- verb, 552
- version space, 329
- virtual body, 50

- walk, 131
- weakest precondition, 358
- weakly dominated, 147
- weakly preferred, 373
- Web Ontology Language, 564
- what, 17
- why question, 177
- whynot question, 177
- win or learn fast (WoLF), 445
- word, 494
- world, 10
- worst-case error, 291
- wrapper, 65

- XML, 564

- yes, 166

- zero-sum game, 424, 430