

# Index

- = (equals), 628
- $\wedge$  (and), 174, 182
- $\leftarrow$  (if), 174, 182, 586
- $\leftrightarrow$  (equivalence), 174
- $\Leftarrow$  (base-level if), 675
- $\perp$  (bottom), 86
- $\epsilon$ -greedy exploration strategy, 558
- $\models$  (entails), 176, 589, 590
- $\neq$  (not equal to), 631
- $\phi$  (denotation of terms), 587, 606
- $\pi$  (denotation of predicate symbols), 587
- $\pi$  (meaning of atoms), 175
- $\mapsto$  (rewritten as), 614
- $\vdash$  (prove), 186
- $\vee$  (or), 679
- % (comment), 586
- & (base-level and), 675
- 0/1 error, 276
  
- abduction, 220, 269, 701
- abductive diagnosis, 222
- abilities, 12
- ABox, 662
- absolute error, 276
  
- absorbing state, 458
- abstractions, 19
- achievement goal, 29, 246
- action, 51, 240
- action constraint, 255
- action features, 255
- action function, 81
- action instance, 257
- action profile, 523, 535
- action replay, 568
- action variable, 252
- activation function, 294, 310
- activation layer, 311
- active learning, 269
- active sensor, 70
- acts, 4
- actuator, 50, 51
- acyclic directed graph, 84
- acyclic knowledge base, 204, 215
- adaptive importance sampling, 410
- additive independence, 433
- additive utility, 433
- adjective, 621
- admissible heuristic, 99, 102, 104
- admissible search algorithm, 102

- aerodynamics, 11
- agent, 4, 11, 50, 51
  - purposive, 13
- agent system, 51
- AI, 3
- AIspace, 737
- algebraic variable, 126
- algorithm
  - asynchronous value iteration, 468
  - back-propagation, 314
  - boosting, 319
  - candidate elimination learner, 327
  - collaborative filtering, 712
  - crossover, 158
  - CSP solver, 141
  - debugging
    - false positive, 201
    - missing answer, 203
  - decision tree
    - learning, 287
  - depth-first branch-and-bound search, 111
  - expectation maximization, 506
  - generalized arc consistency (GAC), 136
  - genetic algorithm, 158
  - Gibbs sampling, 413
  - inductive logic programming, 705
  - iterative deepening, 96
  - k*-means, 501
  - likelihood weighting, 408
  - linear layer (in neural network), 314
  - linear regression, 293
  - local search, 145
  - logistic regression, 296
  - minimax, 531
  - model-based reinforcement learning, 564
  - neural network, 314
  - partial-order planning, 259
  - particle filtering, 411
  - policy hill-climbing, 570
  - policy iteration, 469
  - proof procedure
    - conflict, bottom-up, 210
    - conflict, top-down, 212
    - Datalog, top-down, 604
    - definite clause, bottom-up, 187
    - definite clause, top-down, 191
    - negation as failure, bottom-up, 218
    - negation-as-failure, top-down, 219
  - Q-learning, 557
  - SARSA, 561
  - SARSA with linear function approximation, 567
  - search, 86
  - sigmoid layer (in neural network), 314
  - unification, 601
  - value iteration, 465
  - variable elimination
    - belief network, 378
    - CSP, 143
    - decision network, 451
    - single-stage decision network, 444
    - soft constraints, 163
- Allais Paradox, 434
- alpha-beta ( $\alpha$ - $\beta$ ) pruning, 530
- alternative (ICL), 719
- analysis, 4
- Analytical Engine, 8
- ancestor, 204
- annealing, 152
- annealing schedule, 153
- answer, 185, 190, 207, 595
- answer clause, 189
- answer extraction, 595, 602
- any-conflict algorithm, 151
- anytime algorithm, 26
- aperiodic Markov chain, 385

- application of substitution, 597
- applications of AI, 4, 33, 736
- approximate inference, 370, 402
- approximately correct, 329
- approximately optimal solution, 18
- arc, 82
- arc consistent, 136, 162
- argument, 217, 586
- Aristotelian definition, 660
- array, 746
- Arrow's impossibility theorem, 541
- artificial general intelligence, 735
- artificial intelligence, 3
- ask, 178, 185
- ask-the-user, 194
- askable atom, 194
- assertional knowledge base, 662
- assignment, 126
- assignment space, 132
- assistive technology, 738
- assumable, 206, 220
- assumption-based truth maintenance system, 210, 227
- asymptotic complexity, 93
- asynchronous value iteration, 467
  - algorithm, 468
- ATMS, *see* assumption-based truth maintenance system
- atom, 174, 182
- atomic clause, 182, 586
- atomic proposition, 174, 182
- atomic symbol, 586
- attribute, 649
- auction, 544
- autonomous agent, 736
- autonomous delivery robot, 33, 34
- autonomous vehicles, 736
- average reward, 461
- axiom, 176, 177, 592
- axiom schema, 629
- axiomatizing the domain, 177, 592
- axioms of probability, 347
- axioms of rationality, 426
- back-propagation, 312
  - algorithm, 314
- background knowledge, 13, 702
- backtracking, 90, 134
- backward induction, 529
- backward search, 113
- bag-of-words, 398
- bagging, 318
- base language, 673
- base learner, 318
- base level, 673
- base-level algorithms, 318
- Basic Formal Ontology, 670, 684
- basic income, 741
- batched gradient descent, 293
- Bayes classifier, 491
- Bayes' rule, 353, 487
- Bayesian information criteria (BIC), 498
- Bayesian learning, 512
- Bayesian network, *see* belief network
- Bayesian probability, 344
- beam search, 157
- behavior, 4
- belief, 55
- belief monitoring, 392
- belief network, 360, 361, 507
  - causality, 371
- belief state, 12, 54, 67
- belief state transition function, 55, 58
- best response, 535
- best-first search, 100
- beta distribution, 515
- BFO, *see* Basic Formal Ontology
- bias, 270, 300, 328, 487
- bias-variance trade-off, 301
- bias-free hypothesis space, 328
- bidirectional search, 114
- big data, 300
- bigram model, 398

- binary constraint, 130
- binary variable, 126
- binomial distribution, 514
- biology, 7
- bit, 280, 356
- blame attribution problem, 552
- body
  - agent, 50, 51
  - clause, 182, 586
- Boltzmann distribution, 152, 157, 558
- Boolean property, 649
- Boolean variable, 126, 179
- boosting, 318
  - algorithm, 319
- bot, 50
- bottom  $\perp$ , 86
- bottom-up proof procedure, 187
- boundary of a process, 673
- bounded above zero, 102
- bounded rationality, 26
- branch-and-bound search, 110
- branching factor, 84
- breadth-first search, 87
- burn-in, 413
  
- candidate elimination learner, 326
  - algorithm, 327
- canonical form, 433
- canonical representation, 630
- cardinal feature, 274
- cardinal preference, 29
- case, 320
- case analysis, 139
- case-based reasoning, 320
- causal link, 257
- causal model, 225
- causal network, 371
- causal rule, 245, 696
- causal transduction, 54
- causality, 225, 371
- chain rule, 312, 351
- chance node, 442
- characteristic function, 652, 745
  
- children, 84, 159
- choice space, 719
- choose, 88, 189
- Church–Turing thesis, 8
- clarity principle, 127, 363, 371
- Clark normal form, 634
- Clark’s completion, 213, 634
- class, 285, 491, 499, 652, 662
- classification, 268, 285
- classification tree, 285
- clause, 179
  - definite, 182, 586
  - Horn, 205
- closed list, *see* explored set
- closed-world assumption, 213
- clustering, 499
- CNF, *see* conjunctive normal form
- cognitive science, 11
- cold-start problem, 711
- collaborative filtering, 706
  - algorithm, 712
- combinatorial auction, 737
- command, 51, 55, 58, 556
- command function, 55, 58
- command trace, 52
- commonsense reasoning, 6, 17
- competitive agents, 522
- complements, 435
- complete
  - algorithm, 109, 146
  - bottom-up derivation, 189
  - knowledge assumption, 213
  - linear layer, 311
  - proof procedure, 186
- completeness of preferences, 427
- complex preference, 29
- complexity, 93
- compositional measure of belief, 355
- compound proposition, 174
- computational agent, 4
- computational learning theory, 328
- computational limits dimension, 26

- computational linguistics, 612
- computational sustainability, 737
- concept, 669
- conceptualization, 584, 657
- conditional effect, 244, 246
- conditional entropy, 357
- conditional expected value, 356
- conditional odds, 383
- conditional probability, 350
- conditional probability distribution, 351
- conditional probability table, 373
- conditionally independent, 358
- conditioning, 374
- Condorcet paradox, 541
- conflict, 147, 207
- conflicting variable, 151
- conjugate prior, 489
- conjunction, 174
- conjunctive normal form, 209
- consequence set, 187
- consistency-based diagnosis, 207, 209
- consistent heuristic, 107
- consistent hypothesis, 325, 703
- constant, 586
- constrained optimization problem, 160
- constraint, 129
- constraint network, 135
- constraint optimization problem, 160
- constraint satisfaction problem, 131
- context-free grammar, 614
- context-specific independence, 381
- contingent attribute, 655
- continuant, 670, 671
- continuous time, 395
- continuous variable, 126
- controller, 51, 54, 524
  - model-based reinforcement learning, 564
  - policy hill-climbing, 570
  - Q-learning, 557
- SARSA, 561
  - SARSA with linear function approximation, 567
- convolutional neural network, 316
- cooperate, 572
- cooperative agents, 522
- cooperative system, 216
- coordinate, 572
- coordination, 536
- cost, 82, 160
- CPT, *see* conditional probability table
- credit assignment problem, 268, 552
- cross validation, 306, 321
- crossover, 157, 159
  - algorithm, 158
- CSP, *see* constraint satisfaction problem
- CSP solver
  - algorithm, 141
- culture, 7
- cumulative probability distribution, 403
- cumulative reward, 461
- cut, 275
- Cyc, 658
- cycle, 84
- cycle pruning, 106
- cyclic knowledge base, 204
- DAG (directed acyclic graph), 84
- Datalog, 586
- datatype property, 662
- DBN, *see* dynamic belief network
- DCG, *see* definite clause grammar
- dead reckoning, 65
- debugging, 199
  - false negative, 202
  - false positive, 200
    - algorithm, 201
  - incorrect answers, 200
  - infinite loops, 203
  - missing answer, 202

- algorithm, 203
- decision
  - sequential, 444, 445
  - single, 441
- decision function, 449
- decision network, 445
  - multiagent, 527
  - single-stage, 442
- decision node, 442
- decision tree, 285, 381, 438, 508, 524
  - learning, 285, 357, 496
    - algorithm, 287
- decision variable, 438
- decision-theoretic planning, 470
- deduction, 186, 220, 269
- deep learning, 32, 308, 316
- deep reinforcement learning, 568
- Deep Space One, 232
- default, 196, 216
- definite clause, 182, 586
- definite clause grammar (DCG), 614
- definite clause resolution, 189
- delay, 632, 682
- delivery robot, 34
- DENDRAL, 9
- denote, 587
- dense time, 52
- dependent continuant, 670–672
- depth bound, 95
- depth-bounded meta-interpreter, 680
- depth-bounded search, 95
  - algorithm, 96
- depth-first branch-and-bound search, 110
  - algorithm, 111
- depth-first search, 90
- derivation, 187, 190, 603
- derived, 186
  - atom, 225
  - feature, 243
  - knowledge, 652
  - relation, 694
- description logic, 662
- descriptive theory, 435
- design, 4, 222
- design space, 21
- design time computation, 13, 67
- desire, 55
- deterministic, 29
- diagnosis, 222
  - abductive, 222
  - consistency-based, 207
  - decision theoretic, 441
- diagnostic assistant, 33, 36
- dictator, 543
- dictionary, 410
- difference list, 614
- differentia, 660
- dimension
  - computational limits, 25
  - effect uncertainty, 28
  - interaction, 30
  - learning, 27
  - modularity, 21
  - number of agents, 30
  - planning horizon, 23
  - preference, 29
  - representation, 23
  - sensing uncertainty, 28
- directed acyclic graph, 84
- directed graph, 82
- Dirichlet distribution, 515
- discount factor, 462
- discounted reward, 462
- discrete time, 52
- discrete variable, 126, 345
- discretization, 348, 395
- disjunction, 174, 205, 679
- disjunctive normal form, 209
- disposition, 671, 672
- DNF, *see* disjunctive normal form
- do command, 61, 556, 692
- document, 395
- domain, 126, 271, 345, 587, 649, 745

- domain consistent, 135
- domain expert, 69
- domain ontology, 668
- domain splitting, 139
- dominant strategy, 538, 542
- don't-care non-determinism, 88
- don't-know non-determinism, 88
- dot product, 392
- DPLL, 181
- dynamic belief network, 393
- dynamic decision network, 470
- dynamic programming, 115
- dynamic relation, 694
- dynamics, 28, 385, 458
  
- economically efficient mechanism, 542
- effect, 240, 243
- effect constraint, 252
- effect uncertainty dimension, 28
- effectively computable function, 8
- effector, 50
- elimination ordering, 376
- EM, *see* expectation maximization
- embodied agent, 50
- empirical frequency, 284
- empirical systems, 4
- empty body, 182, 586
- engineering goal, 4
- ensemble learning, 318
- entails, 176
- entity, 24, 582, 670
- entropy, 279, 280, 357
- environment, 12
- epistemology, 10, 344
- equal, 628
- equilibrium distribution, 387
- equivalence, 174
- ergodic Markov chain, 385
- error, 274, 313
  - 0/1, 276
  - absolute, 276
  - data, 271
  - false negative, 202, 281
  - false positive, 200, 281
  - sum squares, 276
  - worst case, 277
- error layer, 312
- error of hypothesis, 329
- Euclidean distance, 108, 321
- evaluation function, 146, 532
- event calculus, 699
- evidence, 350
- evidential model, 225
- evolutionary algorithm, 553
- exact inference, 370
- exclusive-or, 296
- existence uncertainty, 723
- existentially quantified variable, 591
- expanding a path, 86
- expectation, 505
- expectation maximization, 504, 509
  - algorithm, 506
- expected monetary value, 430
- expected utility, 450
- expected value, 355
  - of utility of a policy, 450
- experience, 555
- expert opinion, 490
- expert system, 9, 12, 67
- explained away, 222, 366
- explanation, 220
- explanation-based learning, 683
- exploit, 557
- explore, 558
- explore probability, 558
- explore–exploit dilemma, 552
- explored set, 106
- expression, 586
- extension, 130, 173
- extensional set, 652
- extensive form, 525
- external knowledge source, 70
- extrapolation, 271
  
- $f(p)$ , 100
- fact, 182, 586

- factor, 374, 381, 746
- factored finite state machine, 56
- factored representation, 56
- factorization, 361
- failure, 88, 217
  - natural, 95
- fairness, 15, 88, 610
- false, 175, 589
- false-negative error, 18, 202, 281
- false-positive error, 18, 200, 281
- false-positive rate, 281
- fault, 207
- feature, 24, 125, 271
- feature engineering, 565, 568
- feature selection, 306
- feature-based representation of actions, 244
- feed-forward neural network, 310
- FIFO, 87
- filtering, 65, 388, 392
- finite failure, 220
- finite horizon, 23
- finite state controller, 56
- finite state machine, 56
- first-order predicate calculus, 609
- fixed point, 188
- flat structure, 22
- floundering goal, 637
- fluent, 694
- flying machines, 11
- foaf (friend-of-a-friend), 661
- fold, 308
- for all ( $\forall$ ), 591
- forward chaining, 187
- forward planner, 247
- forward sampling, 404
- forward search, 113
- found a solution, 86
- frame constraint, 252
- frame problem, 724
- frame rule, 245, 696
- framing effect, 434
- fringe, 85
- frontier, 85
- fully observable Markov decision process (MDP), 461
- fully observable world, 28
- function, 671, 672, 745
  - symbol, 606
- functional gradient boosting, 319
- fuzzy terms, 59
- gambling, 344
- game tree, 525
- general boundary, 325
- general game playing, 545
- general-to-specific search, 703
- generalization, 270
- generalized additive independence, 435
- generalized answer clause, 602
- generalized arc consistency (GAC), 136
  - algorithm, 136
- generate and test, 132, 162
- generic search algorithm, 85
- genetic algorithm, 157
  - algorithm, 158
- genus, 660
- Gibbard–Satterthwaite theorem, 543
- Gibbs distribution, 152, 157, 558
- Gibbs sampling, 413
  - algorithm, 413
- global optimum, 147, 165
- goal, 12, 29, 55, 81, 82, 246
- goal constraint, 252
- goal node, 82
- goal state, 81
- Google, 386, 402
- gradient descent, 165, 292
- grammar, 614
  - context-free, 614
  - definite clause, 614
- granularity, 394
- graph, 82



- greedy, 23
- greedy ascent, 146
- greedy best-first search, 100
- greedy descent, 146
- greedy look-ahead, 704
- greedy optimal split, 289
- ground expression, 587
- ground instance, 597, 599
- ground representation, 674
- ground truth, 300
- guaranteed bounds, 370
  
- $h(n)$ , *see* heuristic function
- hard clustering, 499
- hard constraint, 125, 129
- head, 182, 586
- help system, 396, 493
- Herbrand interpretation, 600
- heuristic, 88
- heuristic depth-first search, 99
- heuristic function, 99
- heuristic knowledge, 79
- heuristic search, 98
- hidden Markov model, 461
- hidden Markov model (HMM), 387
- hidden property, 706
- hidden units, 310
- hidden variable, 363, 503, 509
- hierarchical control, 56
- hierarchical structure, 22
- hill climbing, 146
- history, 12, 54, 555
- HMM, *see* hidden Markov model
- Hoeffding's inequality, 404
- horizon, 23, 56, 252
- Horn clause, 205
- how question, 196, 681
- hybrid system, 58
- hyperplane, 296
- hypothesis space, 324
  
- i.i.d., *see* independent and identically distributed
- identity uncertainty, 723
- immaterial entity, 670, 671
- imperfect information, 456
- imperfect-information game, 526, 532
- implication, 174
- importance sampling, 407, 408
- incoming arc, 82
- inconsistent, 206
- incorrect answer, 200
- incremental gradient descent, 292
- indefinite horizon, 23, 458
- independent and identically distributed (i.i.d.), 513
- independent choice logic (ICL), 719
- independent continuant, 670, 671
- independent variables, 358
- indicator variable, 180, 275
- indifferent, 427
- individual, 24, 157, 582, 587, 662
- individual–property–value triple, 648
- induction, 220, 269
- inductive logic programming, 701
  - algorithm, 705
- inference, 186
- infinite horizon, 23, 458
- influence diagram, 445
- information content, 280, 357
- information gain, 280, 289, 357
- information seeking actions, 444
- information set, 526
- information theory, 280, 356, 498
- inheritance, 653
- init* (initial situation), 692
- initial part of a path, 82
- initial-state constraint, 252
- input features, 271
- input layer, 310
- insects, 15
- instance, 595, 597
  - ground, 599
- instance space, 324
- insurance, 431

- integrity constraint, 205
- intelligent action, 4
- intelligent tutoring system, 39
- intended interpretation, 177, 583, 592
- intension, 130, 173
- intensional set, 652
- intention, 55
- interaction dimension, 30
- interoperate, 645
- interpolation, 271
- interpretation, 175, 587
- intersection, 748
- intervention, 225, 371
- inverse graph, 113
- inverse reinforcement learning, 739
- involve a variable, 129
- is (Prolog predicate), 620, 680
- is\_a, 649
- island, 114
- island-driven search, 114
- iterative best improvement, 146
- iterative deepening, 95
  - $A^*$  ( $IDA^*$ ), 103
  - algorithm, 96
- iterative soft-thresholding, 306
  
- Java, 585, 656
- join, 748
- joint probability distribution, 347, 361
  
- $k$ -fold cross validation, 306
- $k$ -means, 499
  - algorithm, 501
- $kd$ -tree, 322
- kernel function, 316, 317
- knowledge, 16, 66
- knowledge base, 13, 16, 67, 176, 182, 586
- knowledge base designer, 177
- knowledge engineer, 69
- knowledge graph, 650
- knowledge is given, 27
- knowledge is learned, 27
  
- knowledge level, 21, 196
- knowledge-based system, 66
- knowledge-level debugging, 199
  
- $L_0$  error, 277
- $L_1$  error, 277
- $L_1$  regularizer, 305
- $L_2$  error, 277
- $L_2$  regularizer, 305
- $L_\infty$  error, 277
- landmark, 59
- language, 614
  - natural, 395, 612
- language bias, 328
- Laplace smoothing, 304, 489, 517
- latent tree model, 496
- latent variable, *see* hidden variable
- law of large numbers, 404
- Laws of Robotics, 738
- layer (neural network), 310
- leaf, 84
- learning, 67, 267–339, 487–520, 549–577, 701–711
  - Bayesian, 512
  - belief network, 507–512
  - bias, 270, 300, 328, 487
  - case-based, 320
  - collaborative filtering, 706
  - decision tree, 285, 357, 496
  - deep, 308
  - expectation maximization (EM), 503, 509
  - inductive logic programming, 701
  - $k$ -means, 499
  - minimum description length, 498
  - missing data, 509
  - multiagent, 569
  - naive Bayes classifier, 491
  - neural network, 308
  - PAC, 328, 517
  - probabilistic classifier, 491
  - probabilities, 487, 508
  - reinforcement, 549

- relational, 701
- structure, 510
- supervised, 268
- to coordinate, 569
- to act, 549
- unsupervised, 268, 499
- version space, 325
- learning dimension, 27
- learning rate, 292
- least fixed point, 189
- leave-one-out cross validation, 308, 321
- level of abstraction, 20
- lifelong learning, 7
- LIFO, 90
- likelihood, 353
- likelihood of data, 278, 487
- likelihood ratio, 383
- likelihood weighting, 407
  - algorithm, 408
- linear classifier, 294
- linear function, 291
- linear layer (in neural network)
  - class, 314
- linear programming, 160
- linear regression, 291, 567
  - algorithm, 293
- linear rule for differentiation, 312
- linearly separable, 296
- Linnaean taxonomy, 660
- list, 608
- literal, 179, 214
- liveness, 15
- local optimum, 147, 165
- local search, 144, 145
  - algorithm, 145
- localization, 389
- log loss, 278, 295
- log-likelihood, 278
- log-linear model, 382
- logic, 585
- logic program, 227, 607
- logical variable, 586
- logical connectives, 174
- logical consequence, 176, 589, 590
- logical form, 625
- logical formula, 174
- logically follows, 176
- logistic function, 294
- logistic regression, 295, 383, 492, 508
  - algorithm, 296
- long-term memory, 67
- loop pruning, 106
- lottery, 427
- lowest-cost-first search, 97
- machine learning, *see* learning
- maintenance goal, 29, 246
- MAP model, 487
- mapping, 745
- margin, 317
- Markov assumption, 384
- Markov blanket, 360, 413
- Markov chain, 384
- Markov chain Monte Carlo, 412
- Markov decision process, 458–475, 550
- material entity, 670, 671
- matrix factorization, 710
- matrix multiplication, 392
- maximization, 505
- maximum a posteriori probability, 487
- maximum entropy, 359
- maximum likelihood model, 278, 488
- maximum-likelihood estimate, 284
- MCMC, *see* Markov chain Monte Carlo
- MDL, *see* minimum description length
- MDP, *see* Markov decision process
- mean, 283
- measure, 346, 348
- mechanism, 522, 542
- mechanism design, 542
- median, 283
- memory, 54

- meta-interpreter, 673, 676
  - build proof tree, 681
  - built-in predicates, 679
  - delayed goals, 683
  - depth-bounded, 680
  - disjunction, 679
  - vanilla, 676
- meta-level, 673
- metalanguage, 673
- MGU, *see* most general unifier
- min-factor elimination ordering, 144
- minimal conflict, 207
- minimal diagnosis, 209
- minimal explanation, 221
- minimal model, 189, 600
- minimax, 529, 569
  - algorithm, 531
- minimum deficiency elimination ordering, 144
- minimum description length (MDL), 498
  - MDL principle, 498
- missing at random, 510
- missing data, 271, 492, 509
- mode, 283
- model, 19, 65, 129, 176, 589
- model averaging, 513
- model complexity, 299
- model-based reinforcement learning
  - algorithm, 564
- modular structure, 22
- modularity, 21, 22
- modus ponens, 187
- money pump, 428
- monitoring, 388, 392
- monotone restriction, 107
- monotonic logic, 216
- moral graph, 380
- more general hypothesis, 325
- more specific hypothesis, 325
- most general unifier, 598, 601
- most improving step, 150
- multiagent decision network, 527
- multiagent reasoning, 30, 521–547
- multiple-path pruning, 106
- mutex constraint, 255
- MYCIN, 9
- myopic, 23
- myopically optimal split, 289
- n*-gram, 399
- naive Bayes classifier, 396, 491, 504
- Nash equilibrium, 535
- natural kind, 491, 495, 653
- natural language processing, 395, 612
- nature, 13, 30, 522
- nearest neighbors, 321
- negation, 174, 205
- negation as failure, 214, 219, 633, 634
- negative examples, 702
- negative income tax, 741
- neighbor, 82
- Netflix Prize, 708
- neural network, 308–316, 509
  - algorithm, 314
- neuron, 9, 308
- no* answer, 185
- no-forgetting agent, 448
- no-forgetting decision network, 449
- node, 82
- noise, 271, 301
- noisy observation, 387
- noisy-or, 382, 401
- non-deterministic choice, 88
- non-deterministic procedure, 88
- non-ground representation, 674
- non-monotonic logic, 216
- non-parametric distribution, 348
- non-planning agent, 23
- non-serial dynamic programming, 167
- non-terminal symbol, 614
- nonlinear planning, 257
- norm, 277
- normal form game, 523
- normative theory, 435

- noun, 621
- NP*, 88
- NP*-complete, 88
- NP*-hard, 88
- #NP* (sharp-*NP*), 370
- number of agents dimension, 30
- number uncertainty, 723
  
- object, 24, 582, 623, 649, 671, 672
- object language, 673
- object property, 662
- object-oriented languages, 585, 656
- objective function, 160
- observation, 12, 14, 194, 350, 371
- occurrent, 670–672
- occurs check, 610, 611
- Ockham's razor, 272
- odds, 383
- off-policy learner, 560
- offline computation, 13, 30, 67
- offline learning, 269
- offline reasoning, *see* offline computation
- omniscient agent, 127
- on-policy learner, 561
- one-point crossover, 159
- online computation, 14, 31, 67, 69
- online learning, 269
- online reasoning, *see* online computation
- ontology, 67, 178, 194, 344, 646, 657
- open-world assumption, 213
- optimal algorithm, 118
- optimal policy, 116, 443, 451, 463
- optimal solution, 17, 18, 81, 83, 97
- optimality criterion, 160
- optimism in the face of uncertainty, 558
- optimization problem, 160
- oracle, 88
- orders of magnitude reasoning, 59
- ordinal feature, 275
- ordinal preference, 29
  
- organizations, 6, 741
- outcome, 426, 440, 523
- outgoing arc, 82
- outlier, 278
- overconfidence, 301
- overfitting, 290, 298
- OWL, 658, 662
- OWL functional-style syntax, 664
  
- PAC learning, 328, 517
- Pagerank, 386
- pair, 745
- parameterized random variable, 714
- parametric distribution, 348
- paramodulation, 630
- parents, 361
- partial evaluation, 683
- partial observation, 387
- partial restart, 156
- partial-order planning, 257
  - algorithm, 259
- partially observable game, 526, 532
- partially observable Markov decision
  - process, 461, 474
- partially observable world, 28
- particle, 410
- particle filtering, 410
  - algorithm, 411
- partition function, 488, 513
- passive sensor, 70
- past experience, 12
- path, 82
- path consistency, 139
- pattern database, 105, 117
- payoff matrix, 524
- percept, 51, 556
- percept function, 58
- percept stream, 52
- percept trace, 52
- perceptron, 9, 294
- perfect information, 528
- perfect rationality, 26
- perfect-information game, 525, 569

- periodic Markov chain, 385
- personalized recommendations, 707
- philosophy, 10
- physical symbol system, 19
- physical symbol system hypothesis, 19, 315
- pixel, 51
- plan, 246
- planner, 246
- planning, 246–265, 698
  - CSP, 252
  - forward, 247
  - partial-order, 257
  - regression, 249
- planning horizon, 23, 56
- planning horizon dimension, 23
- plate, 715
- plate model, 715
- point estimate, 274
- policy, 116, 443, 449, 463
  - stochastic, 570
- policy hill-climbing, 571
  - algorithm, 570
- policy iteration, 468
  - algorithm, 469
- policy search, 553
- POMDP, *see* partially observable Markov decision process
- population, 157, 410
- positive examples, 702
- possible action, 249
- possible world, 125, 126, 346, 438, 450, 719
- posterior distribution, 362, 370
- posterior probability, 350, 353
- pragmatics, 613, 628
- precision, 281
- precision-recall space, 281
- precondition, 240, 243, 695
- precondition constraint, 252
- predicate symbol, 586
- prediction error, 276
- predictor, 281
- preference, 12, 160
- preference bias, 328
- preference dimension, 29
- preference elicitation, 433
- preposition, 623
- primitive
  - atom, 225
  - feature, 243
  - knowledge, 652
  - relation, 659, 694
- primitive proposition, 345
- prior count, 302, 489
- prior knowledge, 12
- prior odds, 383
- prior probability, 350, 353, 487
- prisoner's dilemma, 536
- privacy, 708
- probabilistic bounds, 370
- probabilistic independence, 358
- probabilistic inference, 370
- probabilistic relational model, 714
- probabilistically depends on, 361
- probability, 284, 343–423
  - axioms, 347
  - Bayes' rule, 353, 487
  - conditional, 350
  - density function, 348
  - distribution, 346
  - measure, 346, 348
  - posterior, 350
  - prior, 350
  - semantics, 346
- probable solution, 18
- probably approximately correct, 329, 404
- probably approximately correct learning, 328
- process, 23, 671, 673
- process boundary, 671
- projection, 748
- Prolog, 9, 584

- proof, 186
- proof procedure, 186
  - bottom-up, 187
  - conflict, bottom-up
    - algorithm, 210
  - conflict, top-down
    - algorithm, 212
  - Datalog, top-down
    - algorithm, 604
  - definite clause, bottom-up
    - algorithm, 187
  - definite clause, top-down
    - algorithm, 191
  - negation as failure, bottom-up
    - algorithm, 218
  - negation-as-failure, top-down
    - algorithm, 219
  - top-down, 189
- proof tree, 199, 681
- prop*, 648
- property, 24, 582, 648, 649, 662, 672
- property inheritance, 653
- proposal distribution, 409
- proposition, 24, 174, 582
- propositional calculus, 174
- propositional definite clause, 182
- propositional definite clause resolution, 189
- propositional satisfiability, 179
- prospect theory, 434, 435
- Protégé, 668
- proved, 186
- pruning belief network, 379
- pseudocount, 302, 307, 489, 493, 517
- psychology, 10
- punishment, 458
- pure literal, 181
- pure strategy, 534
- purposive agent, 13, 50
- Python, 656
- $Q^\pi$ , 463
- $Q^*$ , 464, 556
- Q-learning, 556
  - algorithm, 557
- Q-value, 463
- qualitative derivatives, 59
- qualitative reasoning, 59
- quality, 671, 672
- quantitative reasoning, 59
- query, 185, 586
- query variable, 362
- querying the user, 194
- queue, 87
- random forest, 317
- random initialization, 145
- random restart, 145, 148, 156, 503
- random sampling, 146
- random variable, 345
- random walk, 146, 148
- random worlds, 359
- range, 649, 745
- rational, 430
- rational agent, 426
- RDF, 658
- RDF-S, 658
- reachable state, 693
- reactive system, 65
- recall, 281
- receiver operating characteristic, 281
- recognition, 222
- recommender system, 706
- record linkage, 723
- rectified linear unit, 311
- recurrent neural network, 316
- reference point, 436
- reflection, 673
- regression, 268, 291
- regression planning, 249
- regression to the mean, 299, 303
- regression tree, 317
- regularization, 304, 710
- regularization parameter, 304
- regularizer, 304
- regulatory capture, 740

- reify, 648, 692
- reinforcement learning, 268, 549–577
- rejection sampling, 405
- relation, 24, 129, 582, 745, 747
- relational algebra, 747
- relational database, 747
- relational learning, 701
- relational probability model, 714
- relational representations, 25
- relational uncertainty, 723
- remember, 55, 58
- representation, 16
- representation bias, 300
- representation dimension, 23, 25
- representation language, 16
- resampling, 411
- resolution, 189
  - SLD, 189, 602
- resolvent, 190
- resource, 658
- Resource Description Framework, 658
- resource goal, 246
- restriction bias, 328
- retry, 86
- return, 461, 556
- revelation principle, 543
- reward, 458
- reward function, 459
- rewrite rule, 614, 630
- ridge regression, 305
- risk averse, 431
- RL, *see* reinforcement learning
- RoboCup, 737
- robot, 12, 34, 50
- ROC space, 281
- role, 671, 672
- root, 84
- root-mean-square (RMS) error, 277
- rule, 182, 586
- rule of inference, 187
- run, 525
- run-time distribution, 154
- running average, 303, 554
- safety, 15, 736, 739
- safety goal, 246
- sample average, 404
- sample complexity, 330, 517
- SARSA, 561
  - algorithm, 561
- SARSA with linear function approximation, 565
  - algorithm, 567
- satisfiable, 220
- satisficing solution, 18, 79
- satisfy
  - constraint, 129
  - policy, 450
- scenario, 220
- scheme, 747
- scientific goal, 4
- scope, 129, 160, 381, 746, 747
- search, 77–123
  - $A^*$ , 100
  - algorithm, 86
  - bidirectional, 114
  - breadth-first, 87
  - cycle-pruning, 105
  - depth-first, 90
  - domain splitting, 139
  - dynamic programming, 115
  - for CSPs, 133
  - gradient descent, 165
  - greedy best-first, 100
  - heuristic, 98
  - island driven, 114
  - iterative deepening, 95
  - local search, 144
  - lowest-cost-first, 97
  - minimax, 530
  - policy, 553
  - stochastic local, 148
- search and score, 511
- search bias, 300, 328
- search strategy, 87



- second-order logic, 609
- second-price auction, 544
- select, 88, 189
- selection, 748
- selector function, 719
- semantic interoperability, 70, 646
- semantic network, 650
- semantic web, 658
- semantics, 175, 584
  - Datalog, 587
  - natural language, 613
  - propositional calculus, 175
  - variables, 589
- semi-decidable logic, 609
- sensing uncertainty dimension, 28
- sensor, 50, 51, 69
- sensor fusion, 392
- sentence, 395, 614
- separable control problem, 66
- sequential decision problem, 445
- sequential Monte Carlo, 410
- sequential prisoner's dilemma, 537
- set, 745
- set difference, 748
- set-of-words, 396
- short-term memory, 67
- sigmoid function, 294, 311, 383
- sigmoid layer (in neural network)
  - class, 314
- simulated annealing, 152
- simultaneous action games, 526
- single agent, 30
- single decision, 441
- single-stage decision network, 442
- singularity, 740
- situation, 692
- situation calculus, 692
- SLD derivation, 190, 603
- SLD resolution, 189, 602
- Smalltalk, 656
- smart house, 34, 43
- smoothing, 388, 393
- SNLP, 265
- social preference function, 541
- society, 7
- soft clustering, 499
- soft constraint, 125, 160
- soft-max, 558
- software agent, 12
- software engineering, 17
- solution, 79, 82
- sound, 186
- spatio-temporal region, 671, 673
- specialization operator, 704
- specific boundary, 325
- specific-to-general search, 704
- squashed linear function, 294
- stable assignment, 501
- stack, 90
- stage, 23, 384
- start node, 82
- start state, 81
- starvation, 88, 610
- state, 23, 79, 384
- state constraint, 253
- state features, 255
- state space, 79
- state variable, 252
- state-space graph, 242, 247
- state-space problem, 81
- static relation, 694
- stationary distribution, 385, 412
- stationary model, 385, 393, 458
- stationary policy, 463
- statistical relational AI, 711–722
- step function, 294
- step size, 165, 166
- stimuli, 12, 51
- stochastic beam search, 157
- stochastic dynamics, 29
- stochastic gradient descent, 293, 312
- stochastic local search, 148
- stochastic policy, 570, 571
- stochastic simulation, 402

- stochastic strategy, 534
- stopping state, 458
- strategic agent, 522
- strategic form game, 523
- strategy, 524, 525, 534
- strategy profile, 525, 535
- strictly dominated, 162, 539
- strictly preferred, 427
- STRIPS assumption, 243
- STRIPS representation, 243
- structure learning, 510
- subgame-perfect equilibrium, 538
- subgoal, 189, 190, 249
- subject, 623, 649
- subjective probability, 344
- substitutes, 435
- substitution, 597
- successor, 145
- sufficient statistics, 505
- sum-of-squares error, 276, 292, 295, 500
- supervised learning, 268, 271–339, 508
- support set, 535, 538
- support vector machine (SVM), 317
- sustainable system, 737
- SVM, *see* support vector machine
- symbol, 19, 127
- symbol level, 21
- symbol system, 19
- symptoms, 37
- syntax
  - Datalog, 586
  - natural language, 613
  - propositional definite clauses, 182
- synthesis, 4
- systematicity, 265
- systems 1 and 2, 57, 71
- tabu search, 149
- tabu tenure, 149
- target features, 271
- TBox, 662
- TD error, 554
- tell, 58, 177
- temperature, 558
- temporal difference error, 554
- temporal region, 671, 672
- term, 586, 606
- terminal symbol, 614
- terminological knowledge base, 662
- test examples, 270, 271
- theorem, 186
- there exists ( $\exists$ ), 591
- thing, 24, 582, 670
- thought, 4
- threat, 538
- time, 52, 692
- time granularity, 394
- time-homogenous model, 385
- tit-for-tat, 537
- TMS, *see* truth maintenance system
- top- $n$ , 706
- top-down proof procedure, 189
- top-level ontology, 670
- topic model, 401
- total assignment, 126
- total reward, 461
- tractable, 17
- trading agent, 34, 41
- training examples, 268, 270, 271, 324
- transduction, 52
- transient goal, 246
- transitivity of preferences, 427
- tree, 84
- tree augmented naive Bayes (TAN)
  - network, 496
- treewidth, 144, 380
- triangle inequality, 107
- trigram, 399
- triple, 648, 701, 745
- triple representation, 648
- trolley problems, 439
- true, 175, 589, 745
- true-positive rate, 281

- trust, 738, 739
- truth maintenance system, 227
- truthful, 542
- try (local search), 145
- tuple, 745, 747
- Turing test, 5
- tutoring system, 33
- two-stage choice, 151
- two-step belief network, 393
- type*, 649
- type I error, 281
- type II error, 281
- UML, 656
- unary constraint, 130
- unary relations, 582
- unbiased learning algorithm, 328
- unconditionally independent, 360
- unfolded network, 393, 471
- unification, 601
  - algorithm, 601
- unifier, 598
- Uniform Resource Identifier, 658, 661
- unigram, 398
- uninformed search strategy, 87, 98
- union, 748
- unique names assumption (UNA), 631
- unit (neural network), 309
- unit resolution, 181
- universal basic income, 741
- universally quantified variable, 589, 591
- unnormalized probabilities, 373
- unsatisfiable, 206
- unsupervised learning, 268, 499, 503
- URI, 658, 661
- useful action, 249
- user, 69, 127, 194
- utility, 18, 431
- utility node, 442
- $V^\pi$ , 463
- $V^*$ , 464
- validation set, 306
- value, 461, 462
  - decision network, 443
- value iteration, 464
  - algorithm, 465
- value of control, 457
- value of information, 455
- vanilla meta-interpreter, 676
- variable, 125, 126
  - action, 252
  - algebraic, 126
  - Boolean, 126, 179
  - continuous, 126
  - decision, 438
  - discrete, 126
  - existentially quantified, 591
  - logical, 586
  - random, 345
  - state, 252
  - universally quantified, 589, 591
- variable assignment, 589
- variable elimination, 141, 372
  - belief network, 372
    - algorithm, 378
  - CSP, 141
    - algorithm, 143
    - decision network, 451
      - algorithm, 451
    - monitoring, 392
    - single-stage decision network, 443
      - algorithm, 444
    - soft constraints, 163
      - algorithm, 163
- variance, 301
- variational inference, 372
- VCG mechanism, 543
- VE, *see* variable elimination
- verb, 623, 649
- version space, 325
- violated, 129
- virtual body, 56

792

Index

walk, 145  
Watson, 9, 639  
weak learner, 318  
weakly dominated, 162  
weakly preferred, 427  
Web Ontology Language, 658  
web services, 41  
why question, 196  
whynot question, 196, 202  
Wikidata, 684  
win or learn fast (WoLF), 574  
Winograd schema, 5  
word, 395, 586, 614  
world, 12  
    possible, 125, 346, 438, 719  
World Wide Mind, 741  
worst-case error, 277  
wrapper, 70  
  
XML, 658  
  
YAGO, 684  
*yes* answer, 185  
  
Zeno's paradox, 98  
zero-sum game, 522, 529