

1

INTRODUCTION







Figure 1.2 FILES: figures/neuron.eps (Tue Nov 3 16:23:13 2009). The parts of a nerve cell or neuron. Each neuron consists of a cell body, or soma, that contains a cell nucleus. Branching out from the cell body are a number of fibers called dendrites and a single long fiber called the axon. The axon stretches out for a long distance, much longer than the scale in this diagram indicates. Typically, an axon is 1 cm long (100 times the diameter of the cell body), but can reach up to 1 meter. A neuron makes connections with 10 to 100,000 other neurons at junctions called synapses. Signals are propagated from neuron to neuron by a complicated electrochemical reaction. The signals control brain activity in the short term and also enable long-term changes in the connectivity of neurons. These mechanisms are thought to form the basis for learning in the brain. Most information processing goes on in the cerebral cortex, the outer layer of the brain. The basic organizational unit appears to be a column of tissue about 0.5 mm in diameter, containing about 20,000 neurons and extending the full depth of the cortex about 4 mm in humans).



2

INTELLIGENT AGENTS








Figure 2.9 FILES: figures/simple-reflex-agent.eps (Tue Nov 3 16:23:44 2009). Schematic diagram of a simple reflex agent.

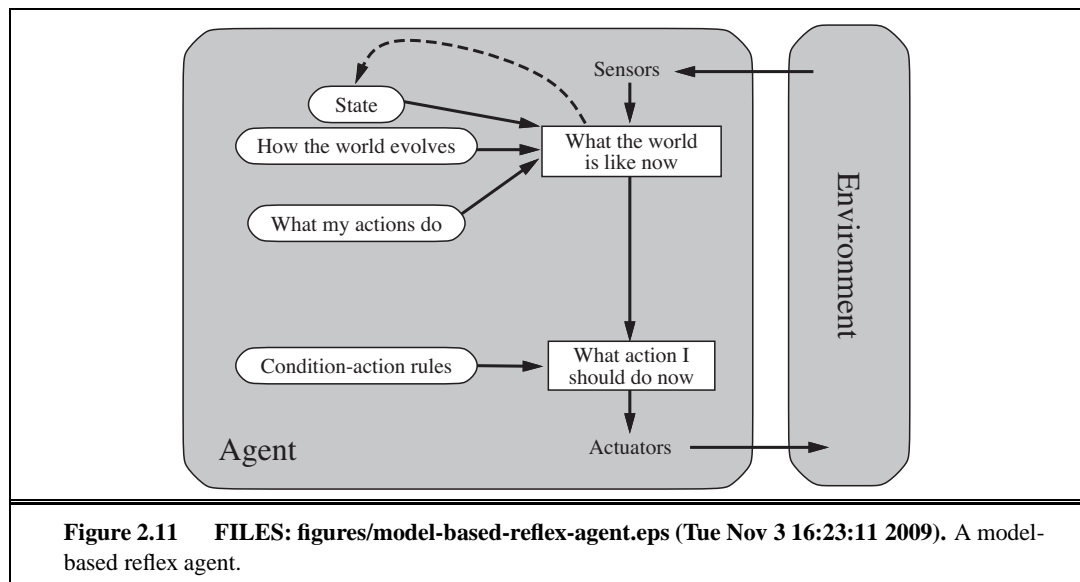
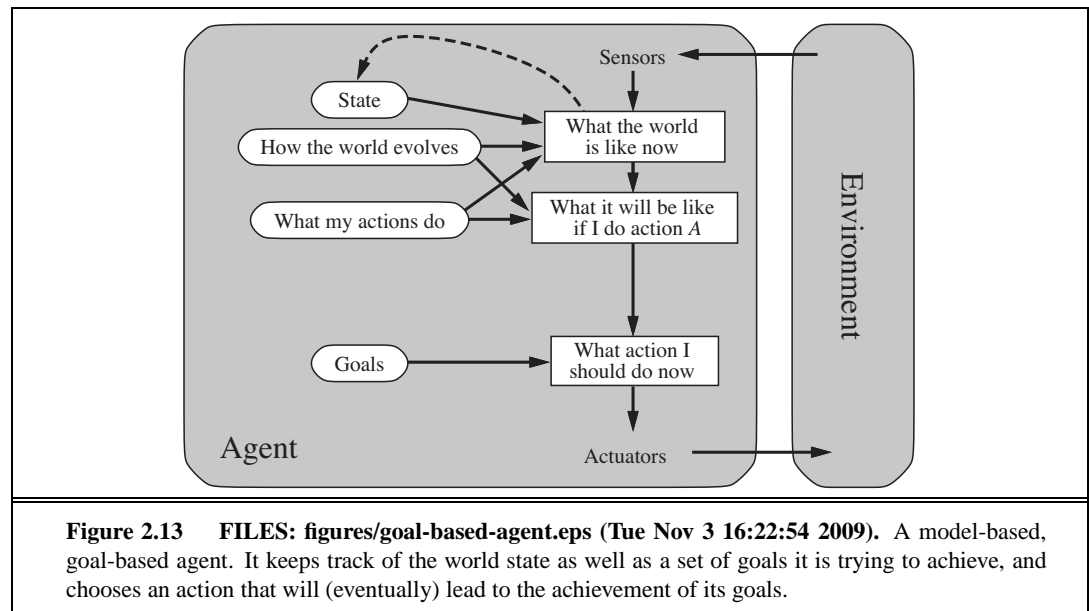
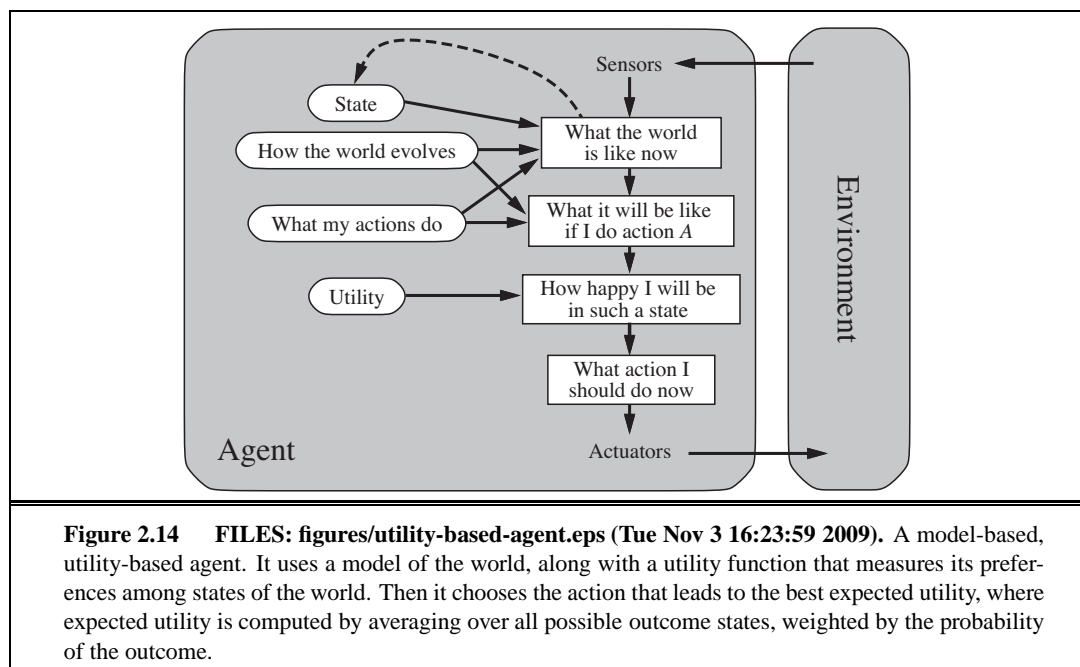


Figure 2.11 FILES: figures/model-based-reflex-agent.eps (Tue Nov 3 16:23:11 2009). A model-based reflex agent.





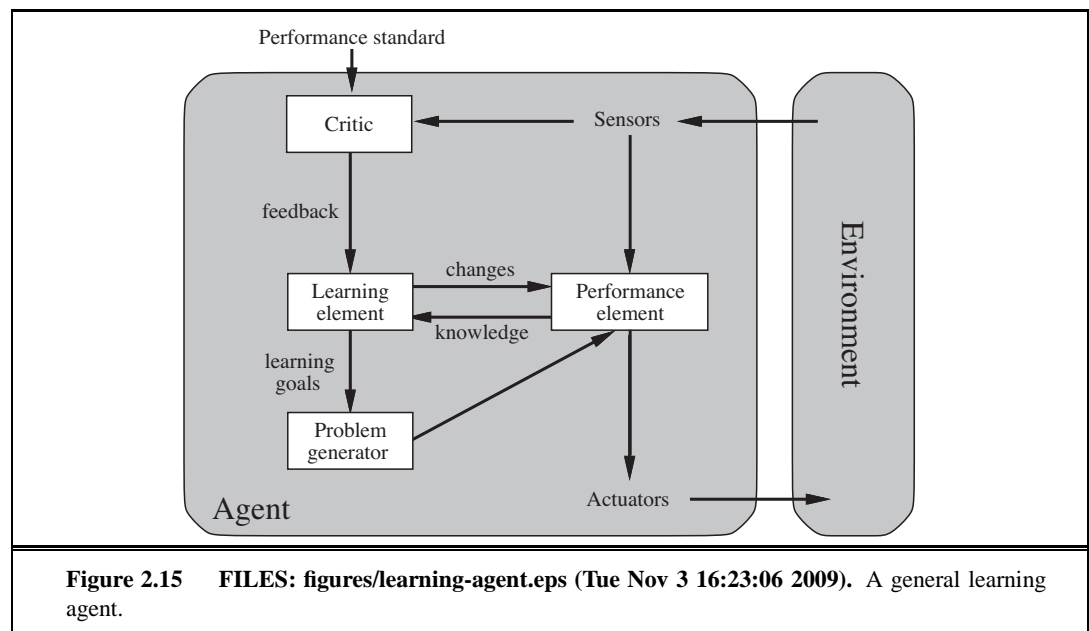
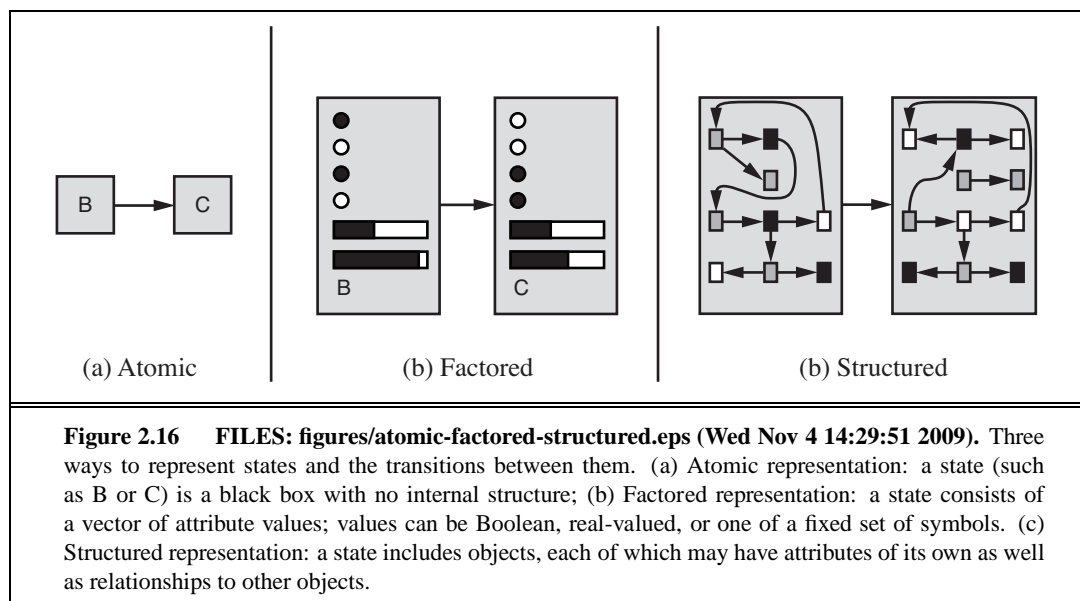




Figure 2.15 FILES: figures/learning-agent.eps (Tue Nov 3 16:23:06 2009). A general learning agent.

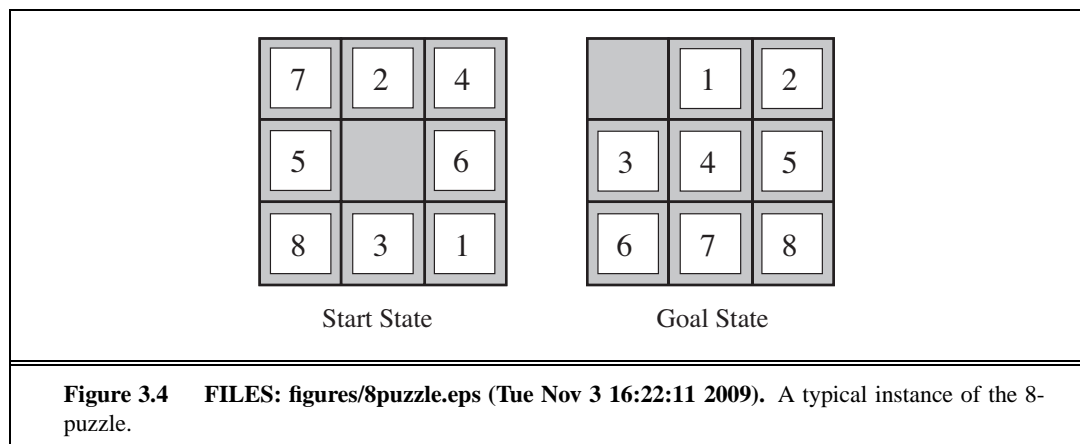


3

SOLVING PROBLEMS BY SEARCHING










Figure 3.5 FILES: figures/8queens.eps (Wed Nov 4 16:21:52 2009). Almost a solution to the 8-queens problem. (Solution is left as an exercise.)





Figure 3.6 FILES: figures/search-map.eps (Tue Nov 3 16:23:38 2009). Partial search trees for finding a route from Arad to Bucharest. Nodes that have been expanded are shaded; nodes that have been generated but not yet expanded are outlined in bold; nodes that have not yet been generated are shown in faint dashed lines.






Figure 3.9 FILES: figures/graph-separation.eps (Tue Nov 3 13:36:17 2009). The separation property of GRAPH-SEARCH, illustrated on a rectangular-grid problem. The frontier (white nodes) always separates the explored region of the state space (black nodes) from the unexplored region (gray nodes). In (a), just the root has been expanded. In (b), one leaf node has been expanded. In (c), the remaining successors of the root have been expanded in clockwise order.



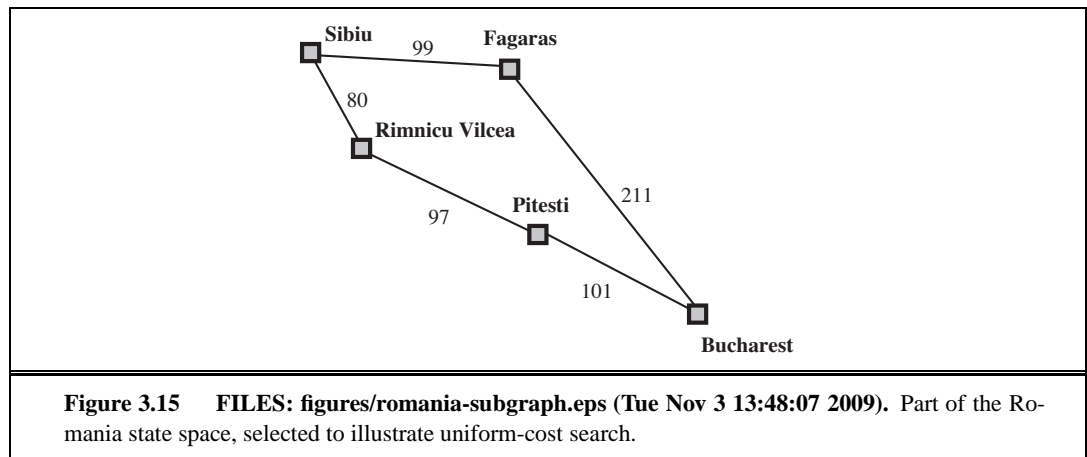


Figure 3.10 FILES: figures/state-vs-node.eps (Tue Nov 3 13:50:06 2009). Nodes are the data structures from which the search tree is constructed. Each has a parent, a state, and various bookkeeping fields. Arrows point from child to parent.








Figure 3.16 FILES: figures/dfs-progress-noblack.eps (Tue Nov 3 13:30:55 2009). Depth-first search on a binary tree. The unexplored region is shown in light gray. Explored nodes with no descendants in the frontier are removed from memory. Nodes at depth 3 have no successors and *M* is the only goal node.




Figure 3.19 FILES: figures/ids-progress.eps (Tue Nov 3 16:23:04 2009). Four iterations of iterative deepening search on a binary tree.


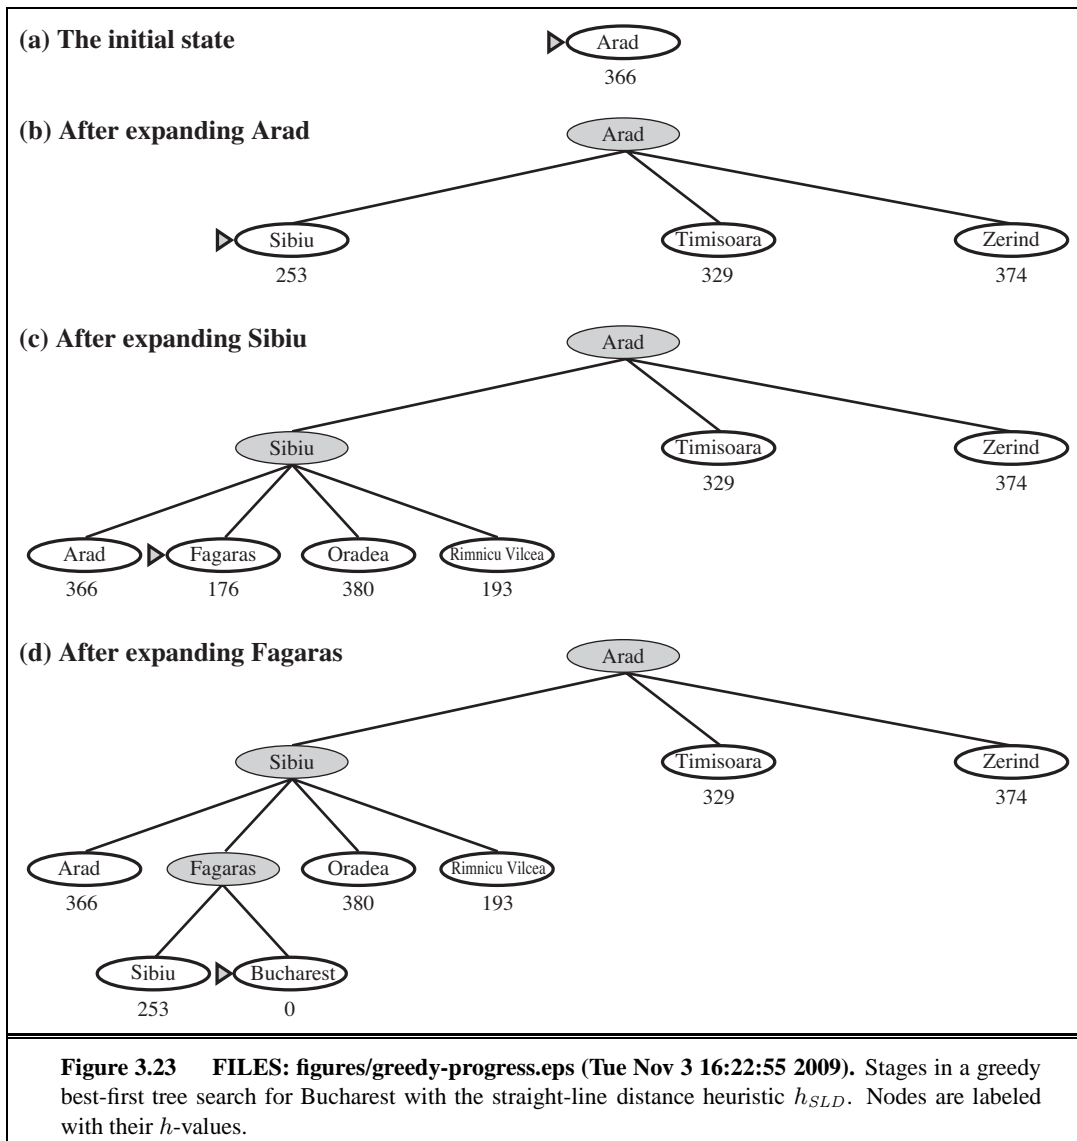
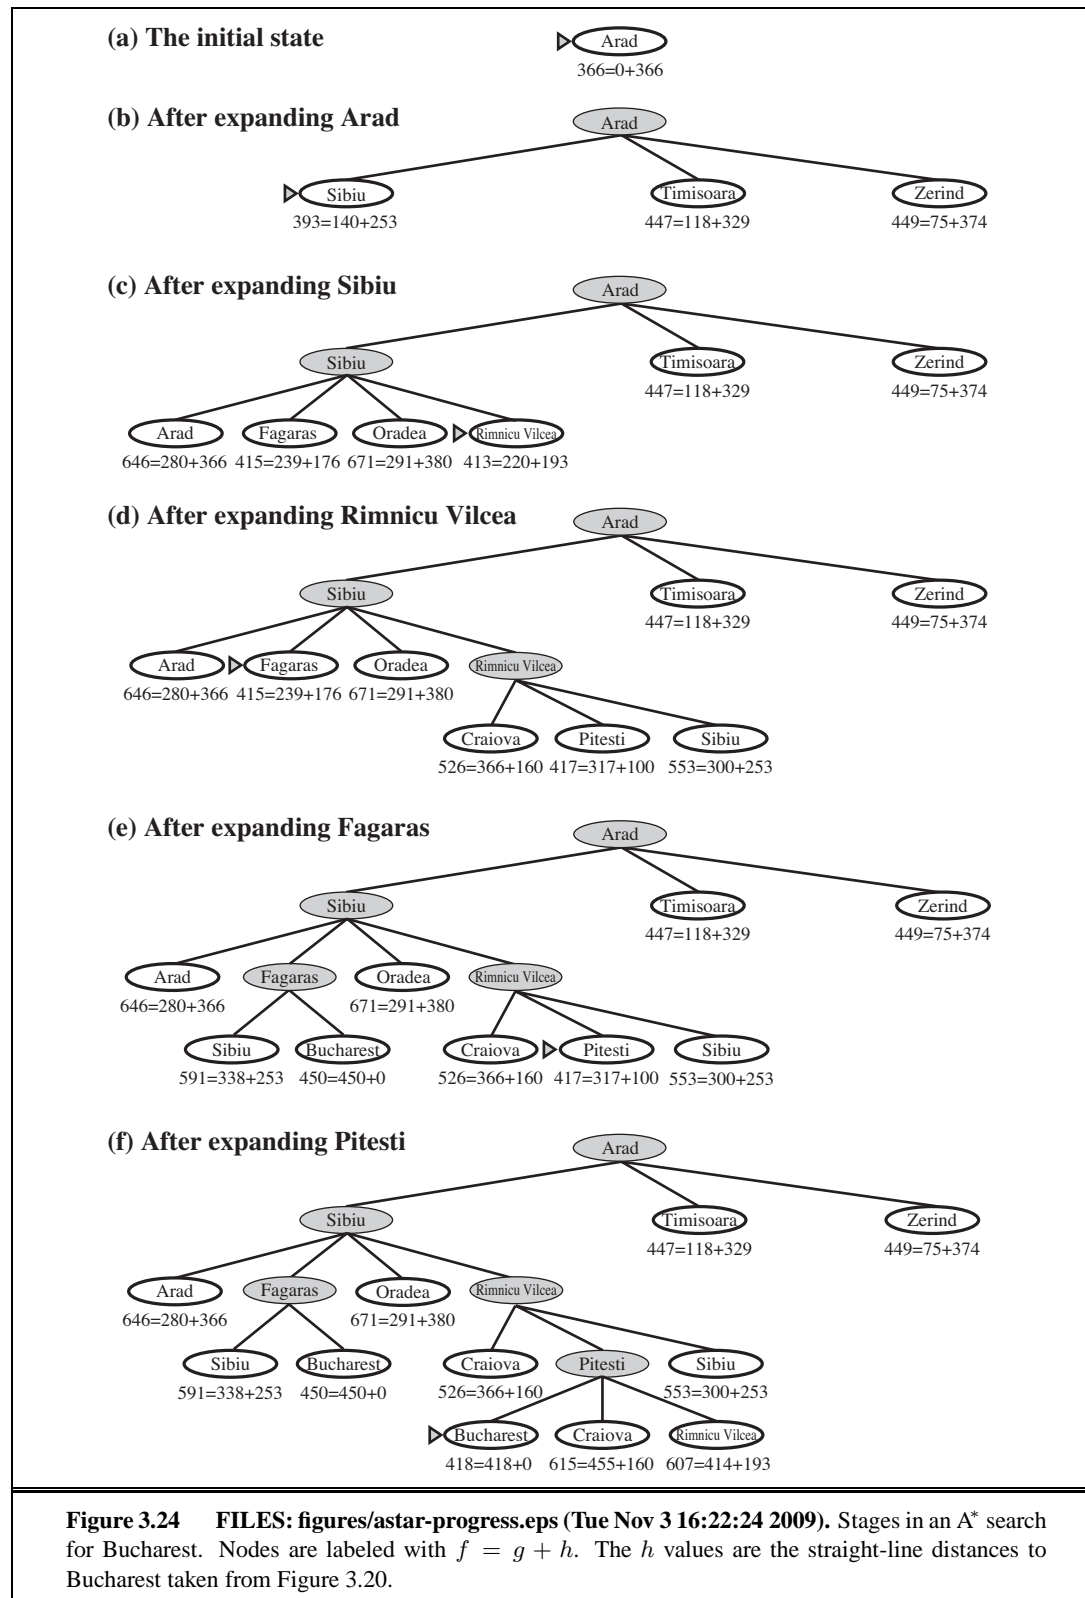



Figure 3.20 FILES: figures/bidirectional.eps (Tue Nov 3 16:22:27 2009). A schematic view of a bidirectional search that is about to succeed when a branch from the start node meets a branch from the goal node.

Arad	366	Mehadia	241
Bucharest	0	Neamt	234
Craiova	160	Oradea	380
Drobeta	242	Pitesti	100
Eforie	161	Rimnicu Vilcea	193
Fagaras	176	Sibiu	253
Giurgiu	77	Timisoara	329
Hirsova	151	Urziceni	80
Iasi	226	Vaslui	199
Lugoj	244	Zerind	374

Figure 3.22 FILES: figures/romania-sld.eps (Tue Nov 3 16:23:37 2009). Values of h_{SLD} —straight-line distances to Bucharest.












Figure 3.27 FILES: figures/rbfs-progress.eps (Tue Nov 3 16:23:27 2009). Stages in an RBFS search for the shortest route to Bucharest. The f -limit value for each recursive call is shown on top of each current node, and every node is labeled with its f -cost. (a) The path via Rimnicu Vilcea is followed until the current best leaf (Pitesti) has a value that is worse than the best alternative path (Fagaras). (b) The recursion unwinds and the best leaf value of the forgotten subtree (417) is backed up to Rimnicu Vilcea; then Fagaras is expanded, revealing a best leaf value of 450. (c) The recursion unwinds and the best leaf value of the forgotten subtree (450) is backed up to Fagaras; then Rimnicu Vilcea is expanded. This time, because the best alternative path (through Timisoara) costs at least 447, the expansion continues to Bucharest.








Figure 3.31 FILES: figures/geometric-scene.eps (Tue Nov 3 16:22:54 2009). A scene with polygonal obstacles. S and G are the start and goal states.




Figure 3.32 FILES: figures/brío.eps (Wed Nov 4 14:35:23 2009). The track pieces in a wooden railway set; each is labeled with the number of copies in the set. Note that curved pieces and “fork” pieces (“switches” or “points”) can be flipped over so they can curve in either direction. Each curve subtends 45 degrees.

4

BEYOND CLASSICAL SEARCH


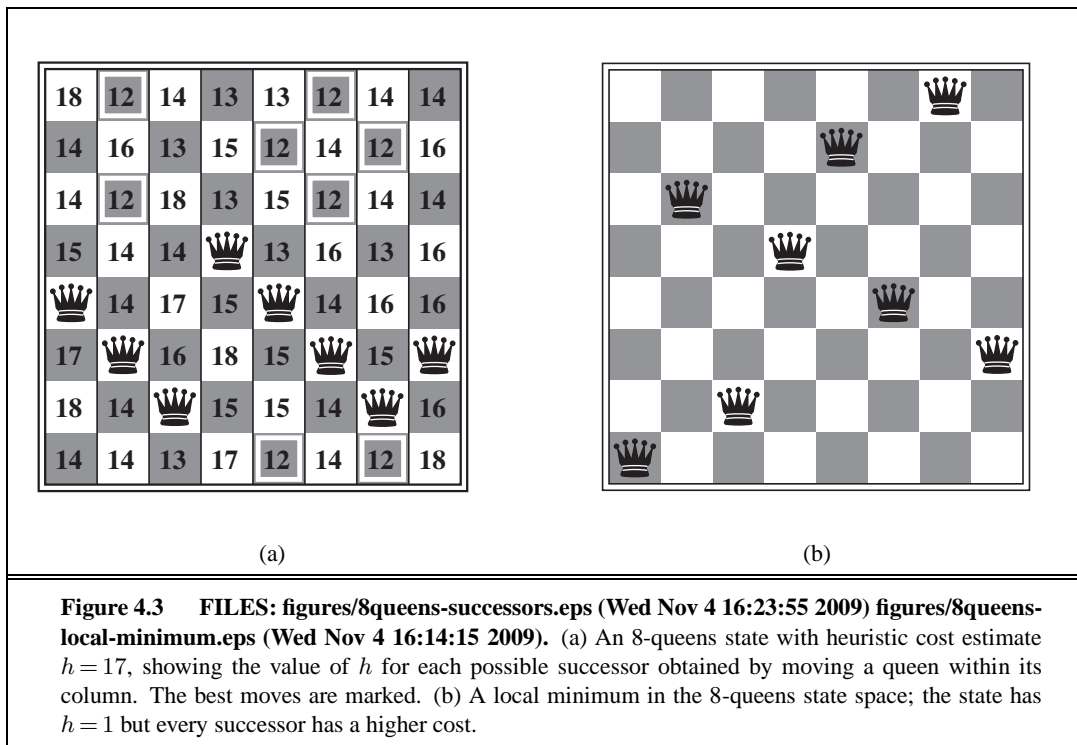


Figure 4.1 FILES: figures/hill-climbing.eps (Tue Nov 3 16:23:03 2009). A one-dimensional state-space landscape in which elevation corresponds to the objective function. The aim is to find the global maximum. Hill-climbing search modifies the current state to try to improve it, as shown by the arrow. The various topographic features are defined in the text.





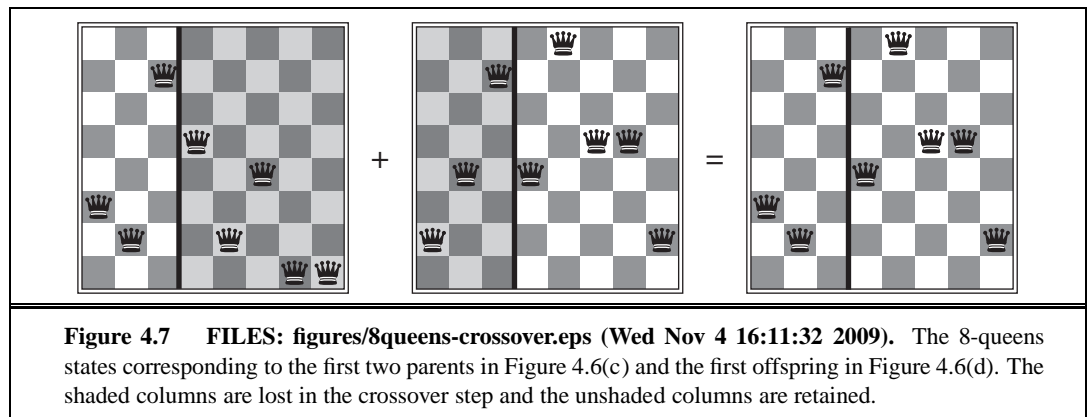
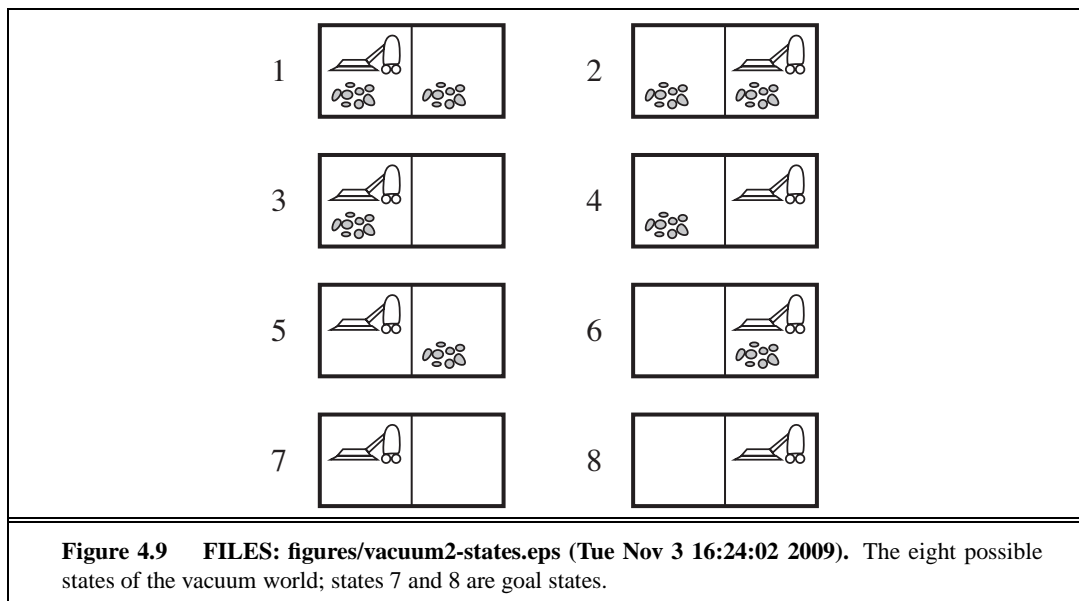


Figure 4.4 FILES: figures/ridge.eps (Tue Nov 3 16:23:29 2009). Illustration of why ridges cause difficulties for hill climbing. The grid of states (dark circles) is superimposed on a ridge rising from left to right, creating a sequence of local maxima that are not directly connected to each other. From each local maximum, all the available actions point downhill.












Figure 4.10 FILES: figures/erratic-vacuum-and-or-plan.eps (Tue Nov 3 13:32:58 2009). The first two levels of the search tree for the erratic vacuum world. State nodes are OR nodes where some action must be chosen. At the AND nodes, shown as circles, every outcome must be handled, as indicated by the arc linking the outgoing branches. The solution found is shown in bold lines.








Figure 4.14 FILES: figures/vacuum2-sets.eps (Tue Nov 3 16:24:01 2009). The reachable portion of the belief-state space for the deterministic, sensorless vacuum world. Each shaded box corresponds to a single belief state. At any given point, the agent is in a particular belief state but does not know which physical state it is in. The initial belief state (complete ignorance) is the top center box. Actions are represented by labeled links. Self-loops are omitted for clarity.








Figure 4.15 FILES: figures/vacuum-prediction-update.eps (Tue Nov 3 13:52:01 2009). Two example of transitions in local-sensing vacuum worlds. (a) In the deterministic world, *Right* is applied in the initial belief state, resulting in a new belief state with two possible physical states; for those states, the possible percepts are $[B, Dirty]$ and $[B, Clean]$, leading to two belief states, each of which is a singleton. (b) In the slippery world, *Right* is applied in the initial belief state, giving a new belief state with four physical states; for those states, the possible percepts are $[A, Dirty]$, $[B, Dirty]$, and $[B, Clean]$, leading to three belief states as shown.











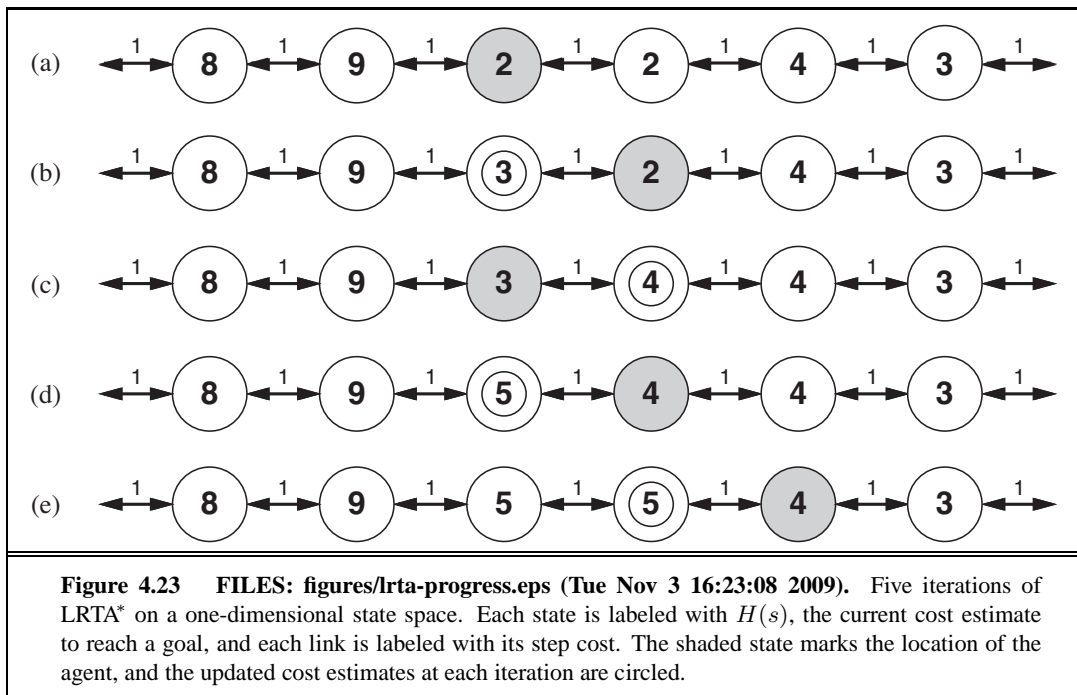


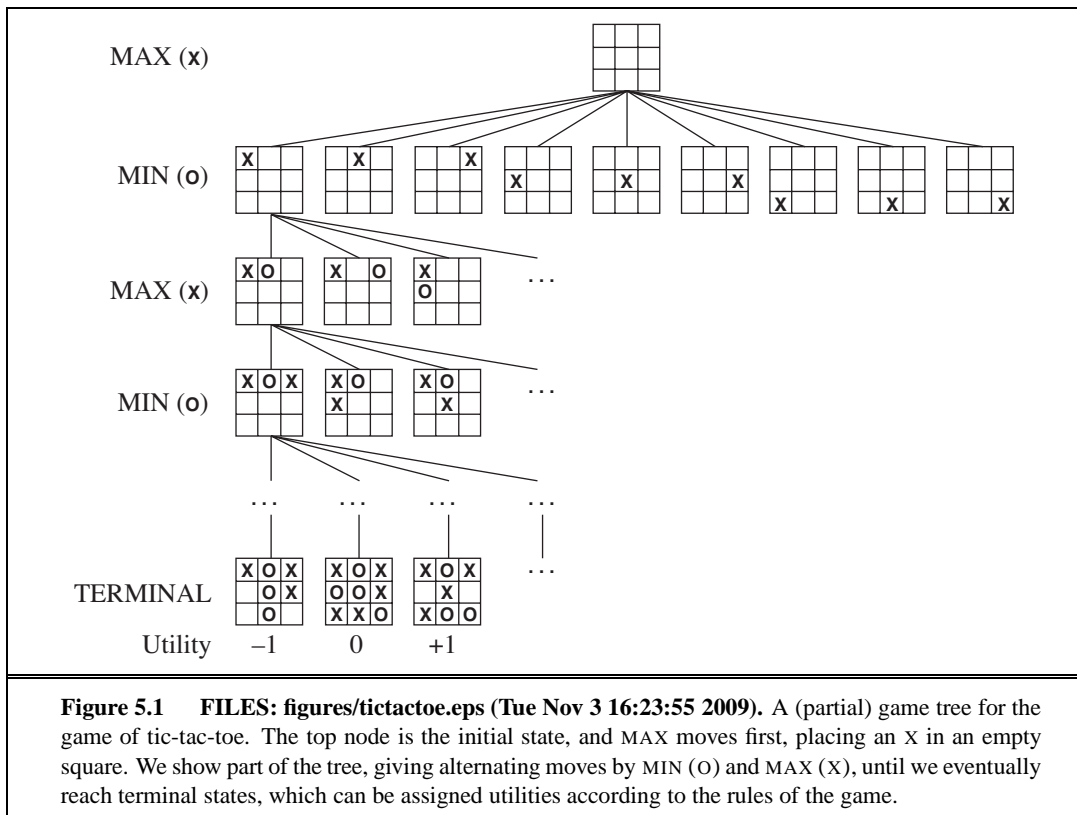
Figure 4.20 FILES: figures/adversary-spaces.eps (Tue Nov 3 16:22:18 2009) figures/adversary-blocks.eps (Sun Oct 25 01:08:26 2009). (a) Two state spaces that might lead an online search agent into a dead end. Any given agent will fail in at least one of these spaces. (b) A two-dimensional environment that can cause an online search agent to follow an arbitrarily inefficient route to the goal. Whichever choice the agent makes, the adversary blocks that route with another long, thin wall, so that the path followed is much longer than the best possible path.





5

ADVERSARIAL SEARCH







Figure 5.2 FILES: figures/minimax.eps (Tue Nov 3 16:23:11 2009). A two-ply game tree. The \triangle nodes are “MAX nodes,” in which it is MAX’s turn to move, and the ∇ nodes are “MIN nodes.” The terminal nodes show the utility values for MAX; the other nodes are labeled with their minimax values. MAX’s best move at the root is a_1 , because it leads to the state with the highest minimax value, and MIN’s best reply is b_1 , because it leads to the state with the lowest minimax value.






Figure 5.5 FILES: figures/alpha-beta-progress.eps (Tue Nov 3 16:22:20 2009). Stages in the calculation of the optimal decision for the game tree in Figure 5.2. At each point, we show the range of possible values for each node. (a) The first leaf below B has the value 3. Hence, B , which is a MIN node, has a value of at most 3. (b) The second leaf below B has a value of 12; MIN would avoid this move, so the value of B is still at most 3. (c) The third leaf below B has a value of 8; we have seen all B 's successor states, so the value of B is exactly 3. Now, we can infer that the value of the root is at least 3, because MAX has a choice worth 3 at the root. (d) The first leaf below C has the value 2. Hence, C , which is a MIN node, has a value of at most 2. But we know that B is worth 3, so MAX would never choose C . Therefore, there is no point in looking at the other successor states of C . This is an example of alpha-beta pruning. (e) The first leaf below D has the value 14, so D is worth at most 14. This is still higher than MAX's best alternative (i.e., 3), so we need to keep exploring D 's successor states. Notice also that we now have bounds on all of the successors of the root, so the root's value is also at most 14. (f) The second successor of D is worth 5, so again we need to keep exploring. The third successor is worth 2, so now D is worth exactly 2. MAX's decision at the root is to move to B , giving a value of 3.

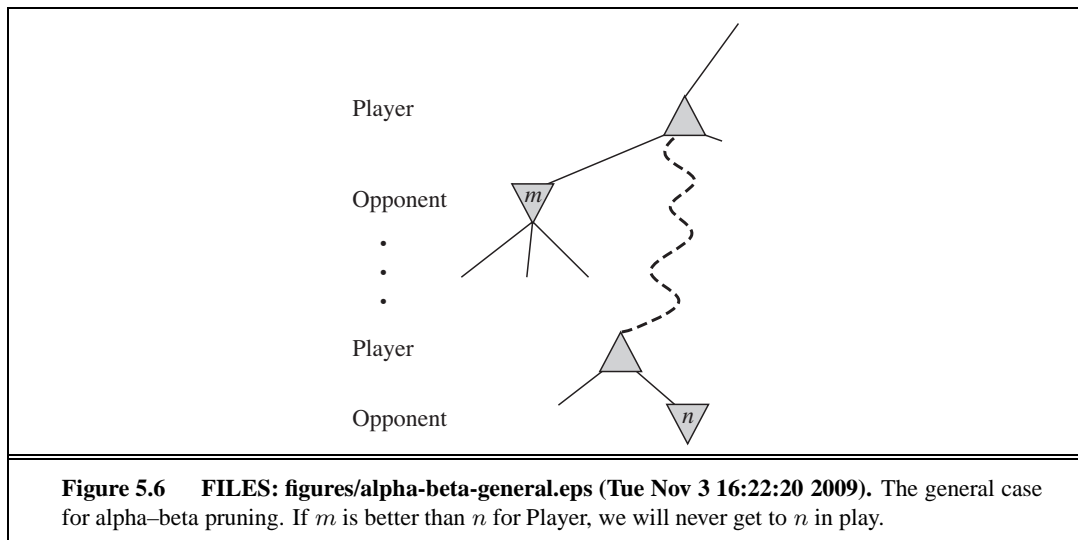

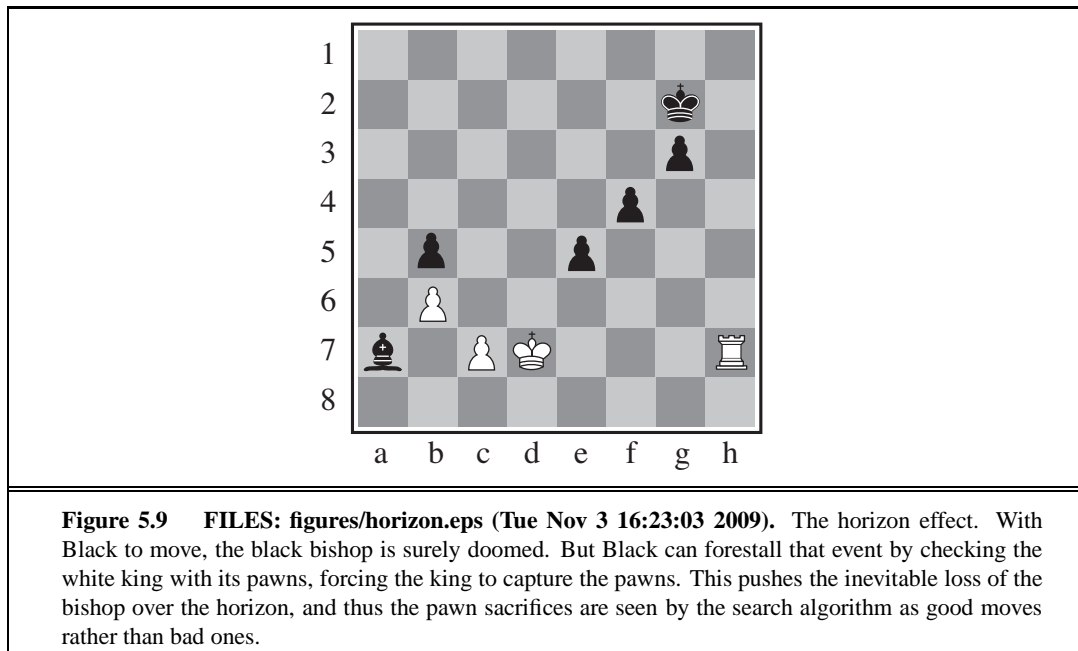


Figure 5.6 FILES: figures/alpha-beta-general.eps (Tue Nov 3 16:22:20 2009). The general case for alpha-beta pruning. If m is better than n for Player, we will never get to n in play.






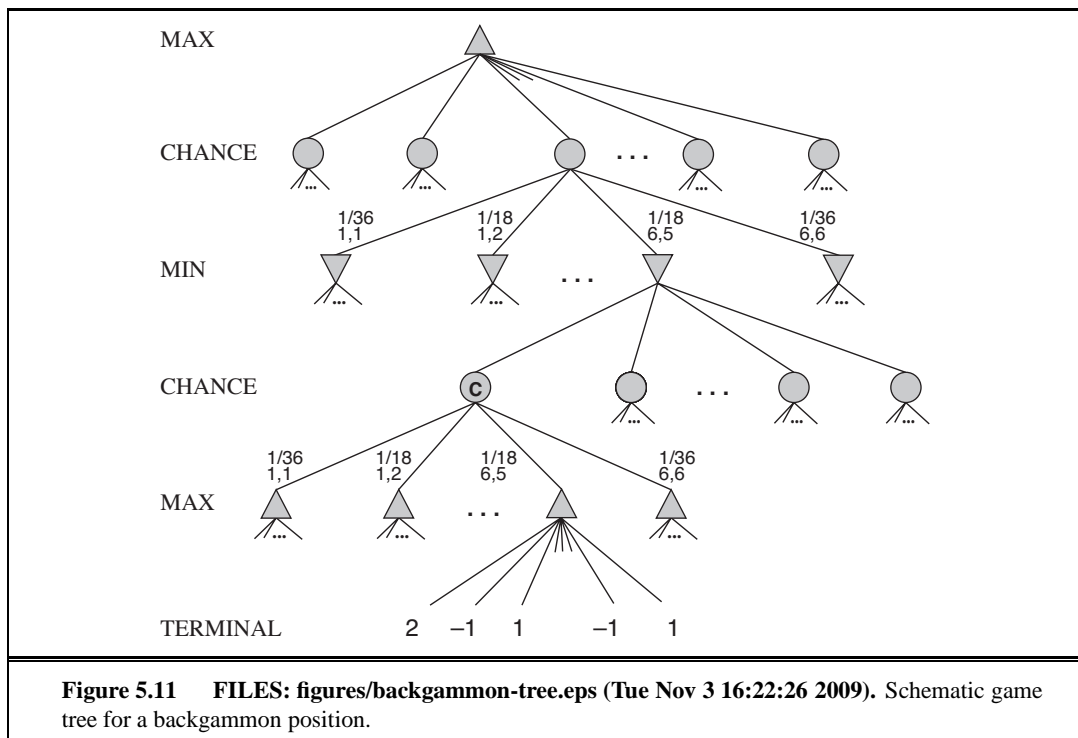


Figure 5.10 FILES: figures/backgammon-position.eps (Tue Nov 3 16:22:26 2009). A typical backgammon position. The goal of the game is to move all one's pieces off the board. White moves clockwise toward 25, and Black moves counterclockwise toward 0. A piece can move to any position unless multiple opponent pieces are there; if there is one opponent, it is captured and must start over. In the position shown, White has rolled 6-5 and must choose among four legal moves: (5-10,5-11), (5-11,19-24), (5-10,10-16), and (5-11,11-16), where the notation (5-11,11-16) means move one piece from position 5 to 11, and then move a piece from 11 to 16.






Figure 5.12 FILES: figures/chance-evaluation.eps (Tue Nov 3 16:22:32 2009). An order-preserving transformation on leaf values changes the best move.

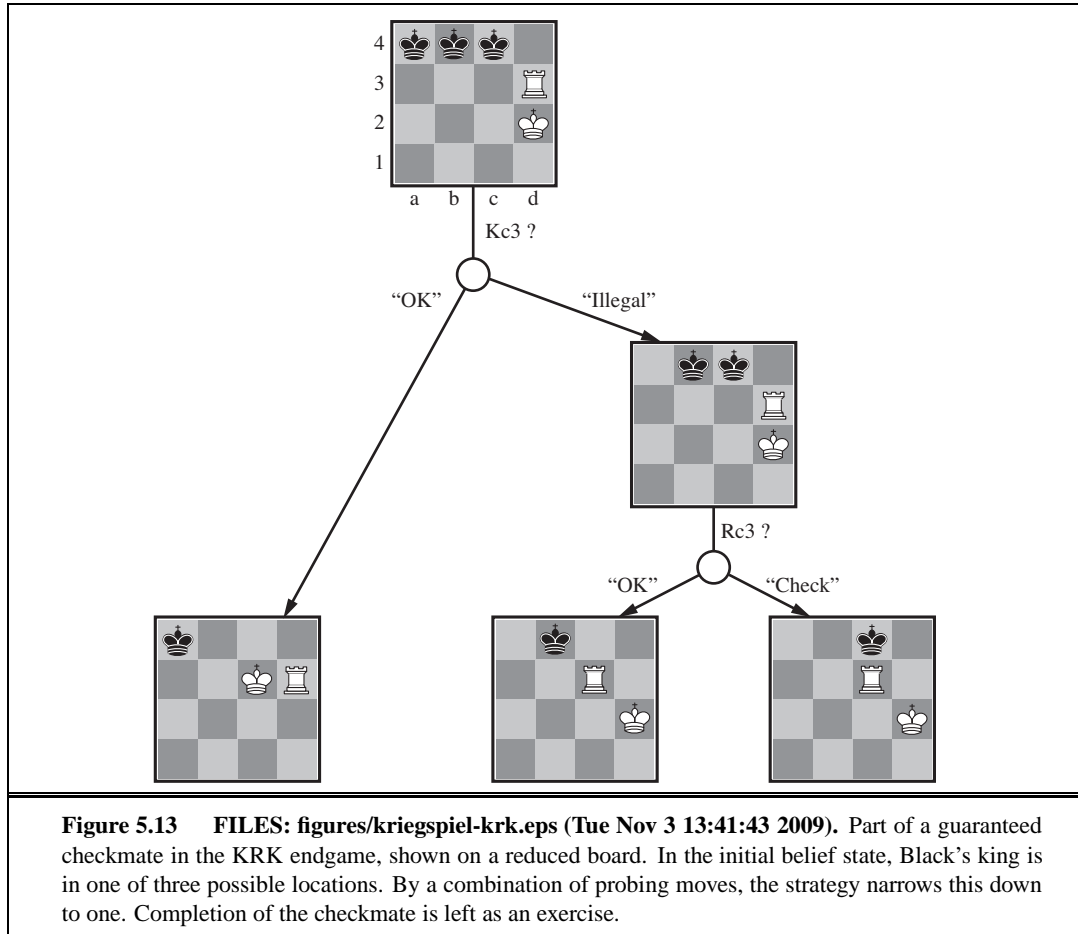
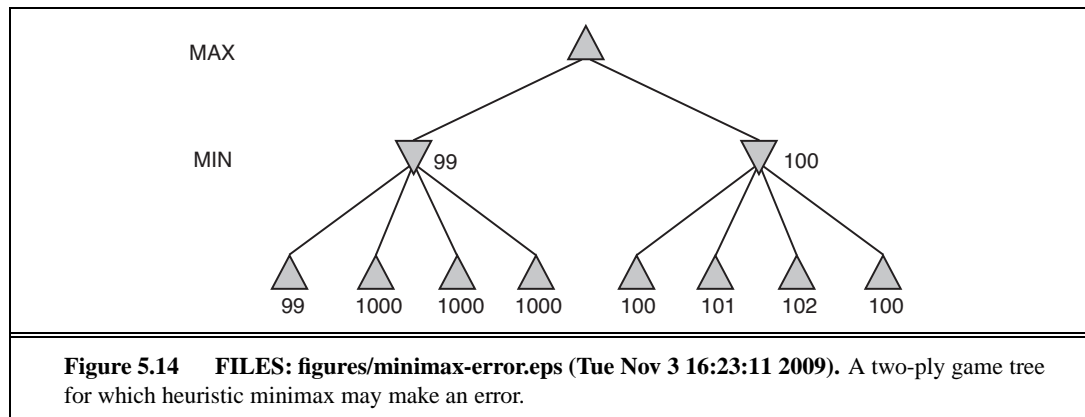
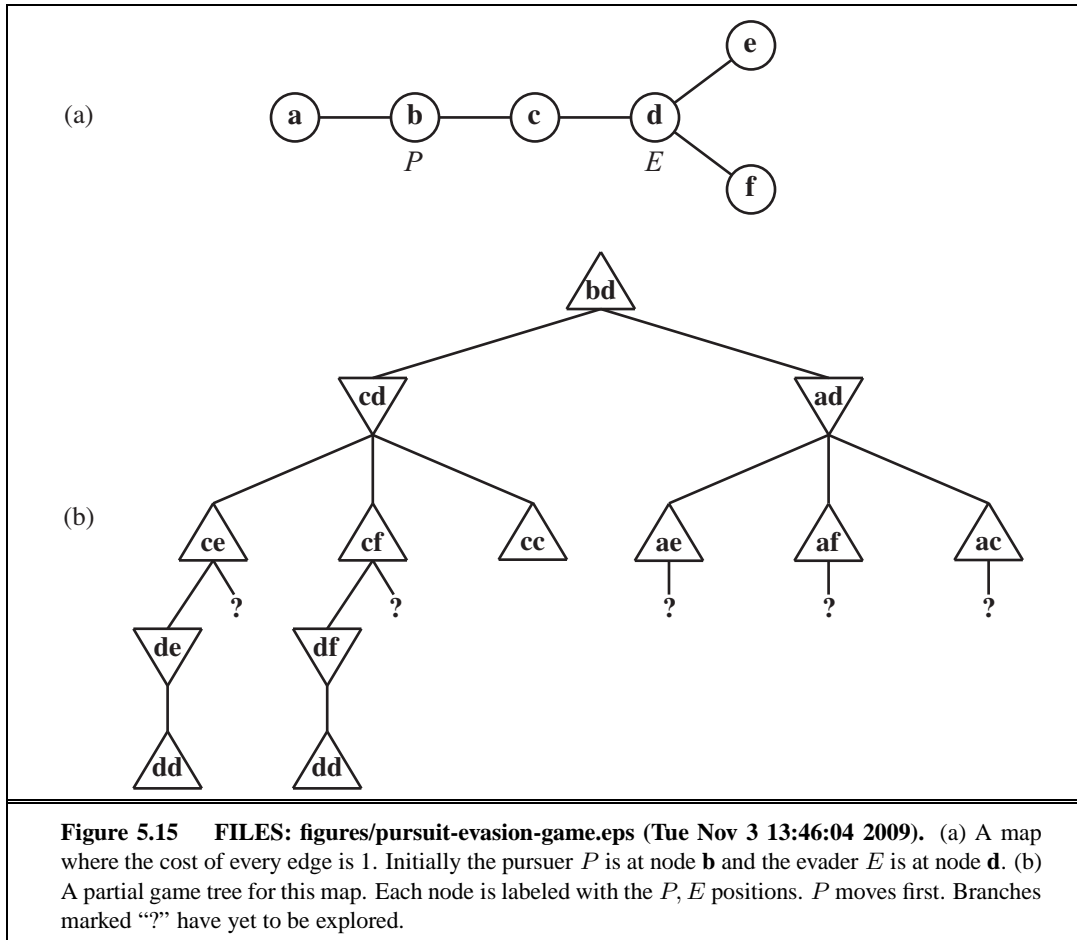


Figure 5.13 FILES: figures/kriegspiel-krk.eps (Tue Nov 3 13:41:43 2009). Part of a guaranteed checkmate in the KRK endgame, shown on a reduced board. In the initial belief state, Black's king is in one of three possible locations. By a combination of probing moves, the strategy narrows this down to one. Completion of the checkmate is left as an exercise.









Figure 5.16 FILES: figures/line-game4.eps (Tue Nov 3 16:23:06 2009). The starting position of a simple game. Player *A* moves first. The two players take turns moving, and each player must move his token to an open adjacent space in either direction. If the opponent occupies an adjacent space, then a player may jump over the opponent to the next open space if any. (For example, if *A* is on 3 and *B* is on 2, then *A* may move back to 1.) The game ends when one player reaches the opposite end of the board. If player *A* reaches space 4 first, then the value of the game to *A* is +1; if player *B* reaches space 1 first, then the value of the game to *A* is -1 .





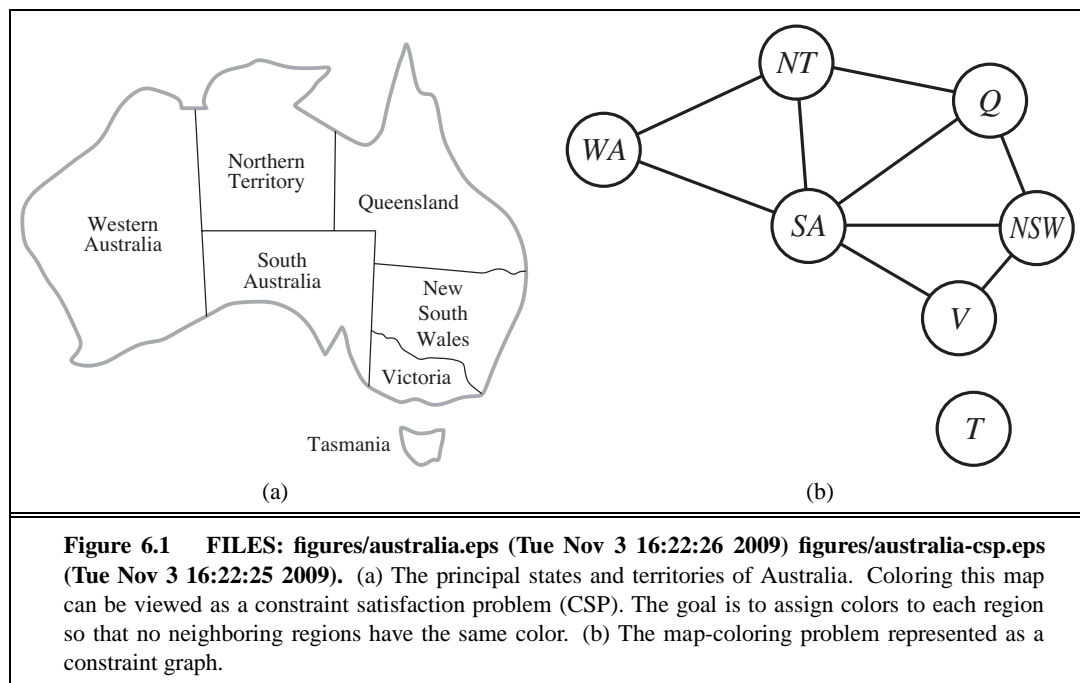


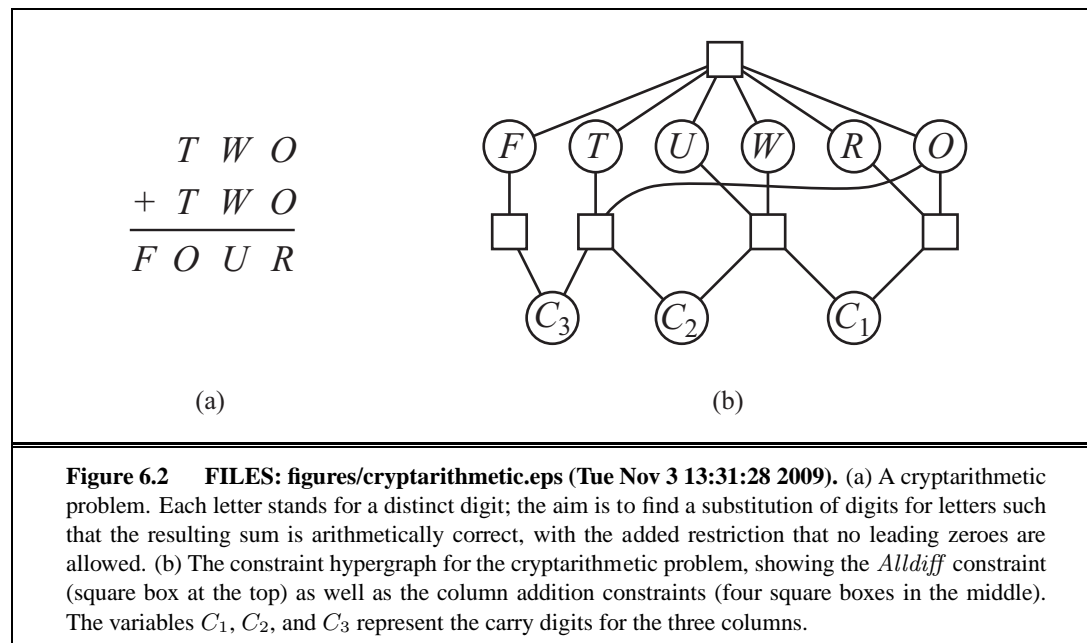
Figure 5.18 FILES: figures/pruning.eps (Tue Nov 3 16:23:22 2009). The complete game tree for a trivial game with chance nodes.

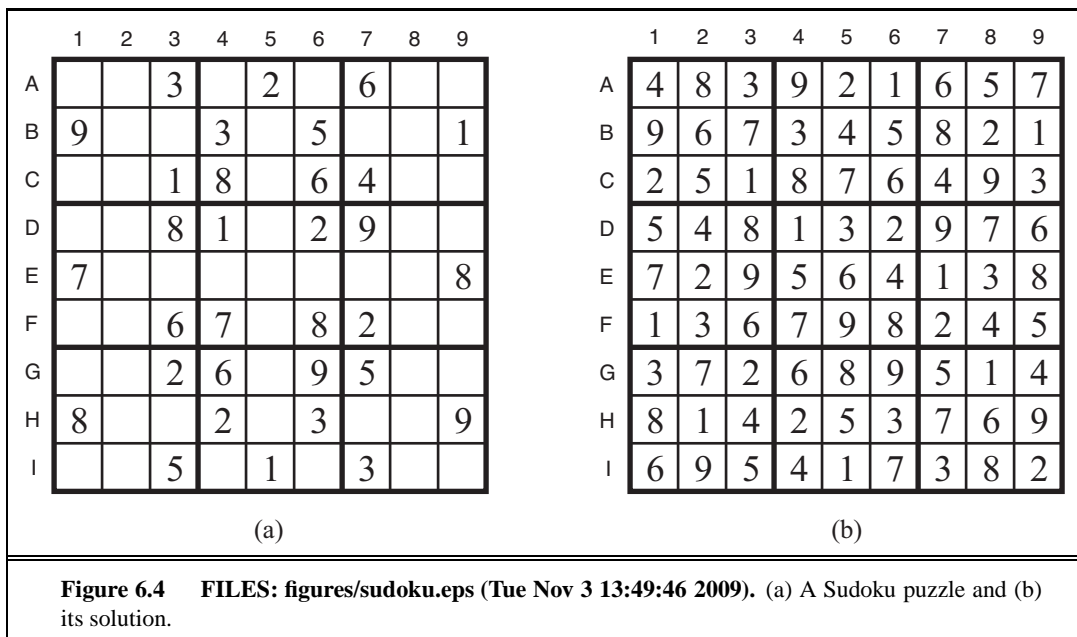



6

CONSTRAINT SATISFACTION PROBLEMS



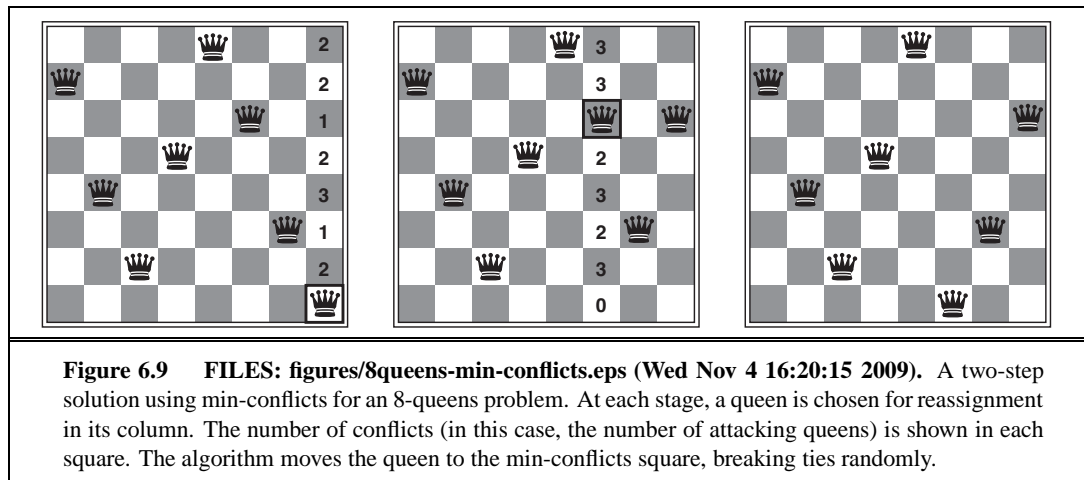









	WA	NT	Q	NSW	V	SA	T
Initial domains	R G B	R G B	R G B	R G B	R G B	R G B	R G B
After $WA=red$	Ⓡ	G B	R G B	R G B	R G B	G B	R G B
After $Q=green$	Ⓡ	B	Ⓢ	R B	R G B	B	R G B
After $V=blue$	Ⓡ	B	Ⓢ	R	Ⓟ		R G B

Figure 6.7 FILES: figures/australia-fc.eps (Tue Nov 3 16:22:25 2009). The progress of a map-coloring search with forward checking. $WA = red$ is assigned first; then forward checking deletes red from the domains of the neighboring variables NT and SA . After $Q = green$ is assigned, $green$ is deleted from the domains of NT , SA , and NSW . After $V = blue$ is assigned, $blue$ is deleted from the domains of NSW and SA , leaving SA with no legal values.



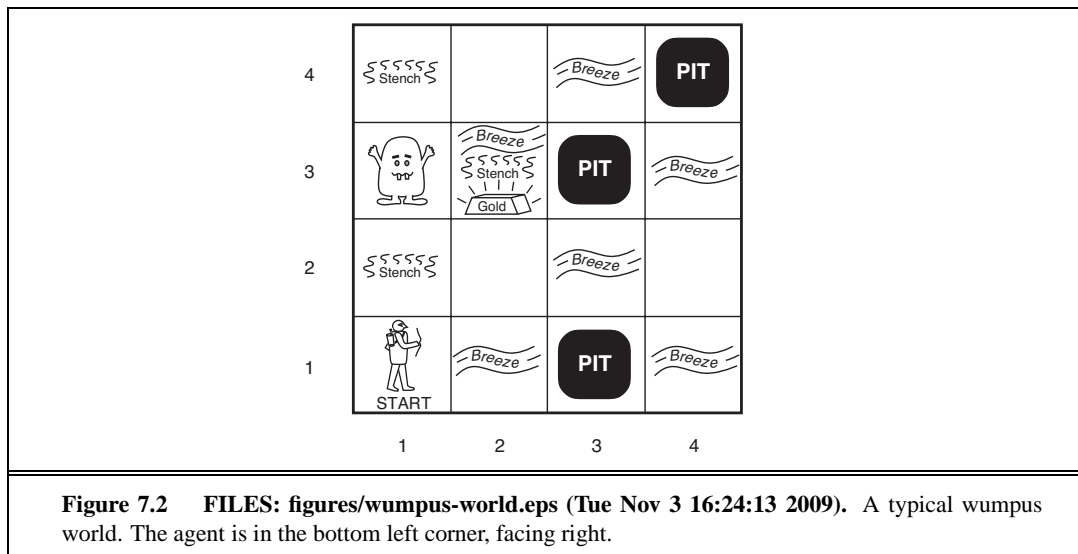





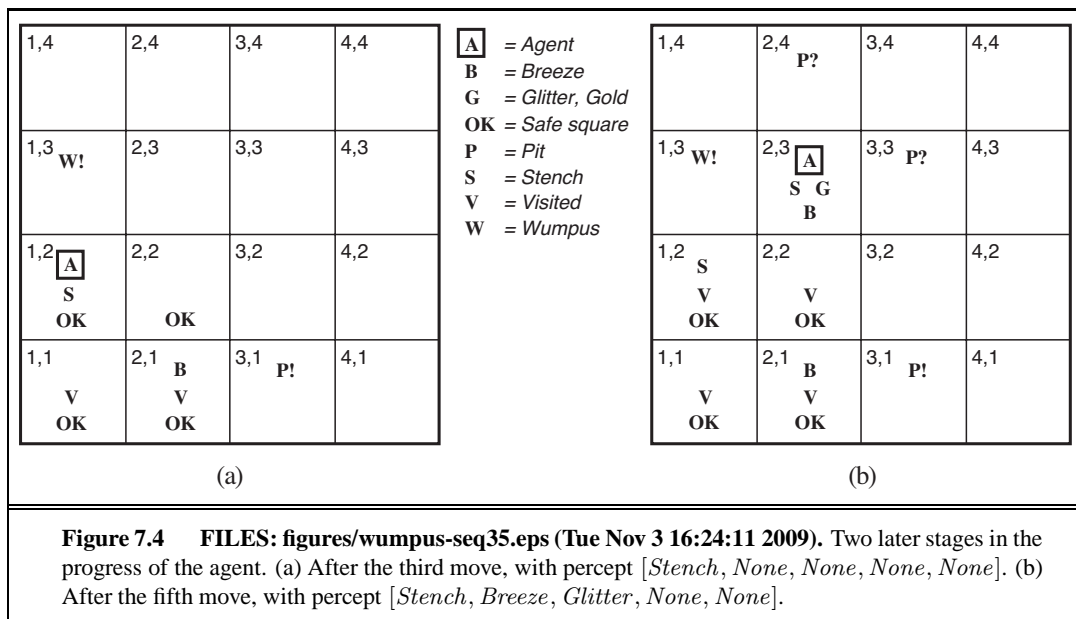



7

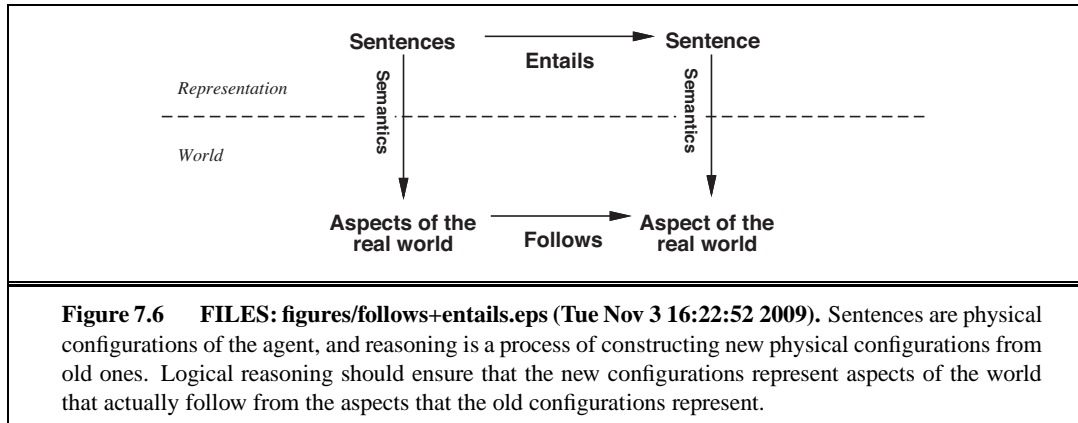
LOGICAL AGENTS


















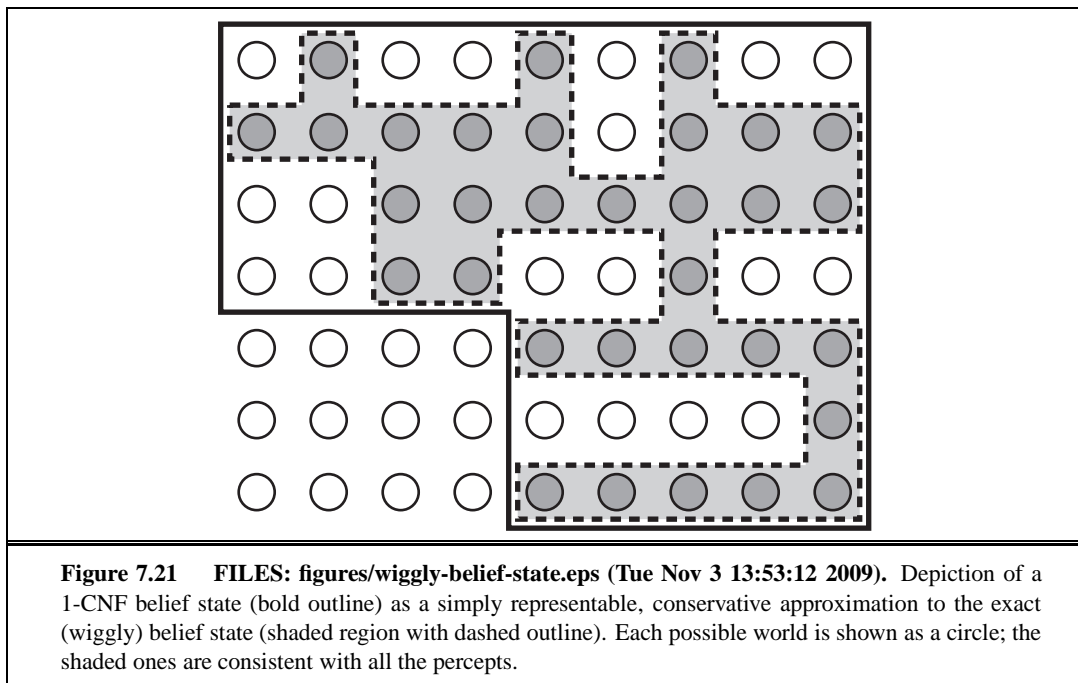



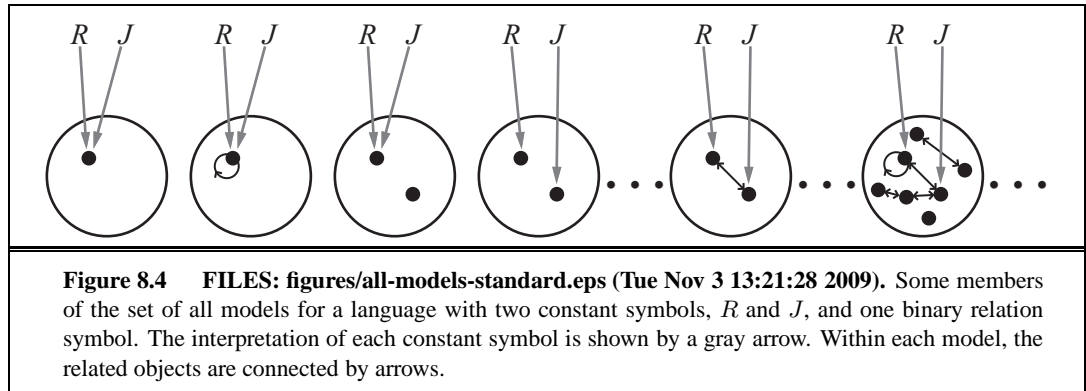
Figure 7.19 FILES: (a) Graph showing the probability that a random 3-CNF sentence with $n = 50$ symbols is satisfiable, as a function of the clause/symbol ratio m/n . (b) Graph of the median run time (measured in number of recursive calls to DPLL, a good proxy) on random 3-CNF sentences. The most difficult problems have a clause/symbol ratio of about 4.3.




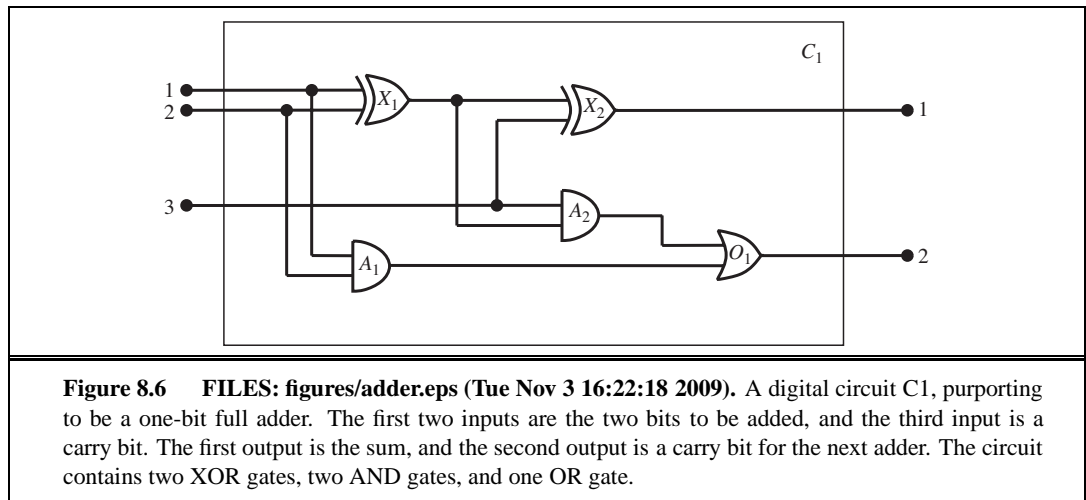
8

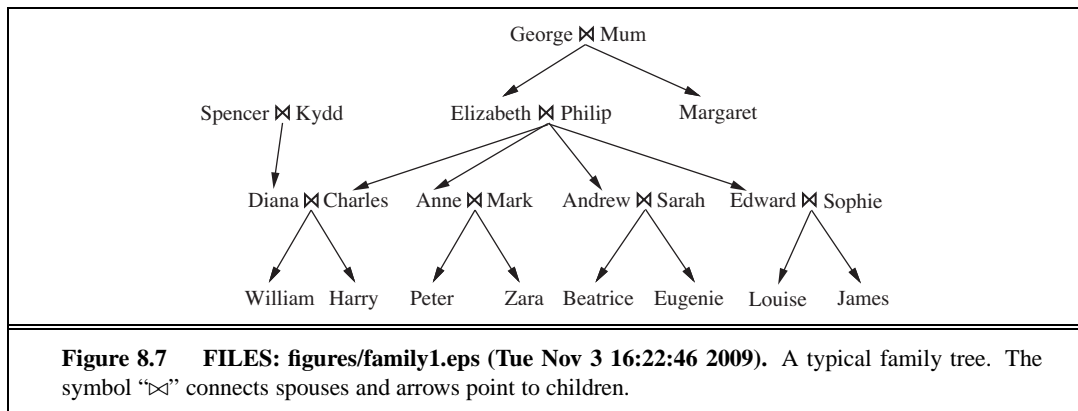
FIRST-ORDER LOGIC














9

INFERENCE IN FIRST-ORDER LOGIC





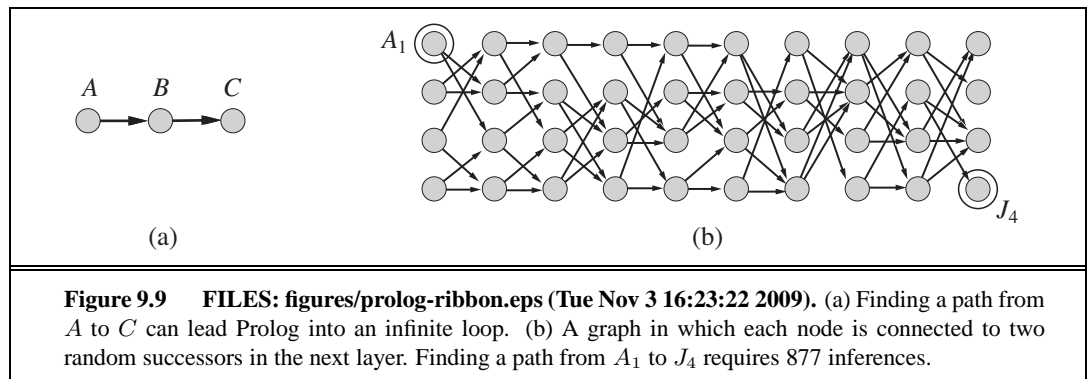



Figure 9.2 FILES: figures/subsumption-lattices.eps (Tue Nov 3 16:23:50 2009). (a) The subsumption lattice whose lowest node is $Employs(IBM, Richard)$. (b) The subsumption lattice for the sentence $Employs(John, John)$.













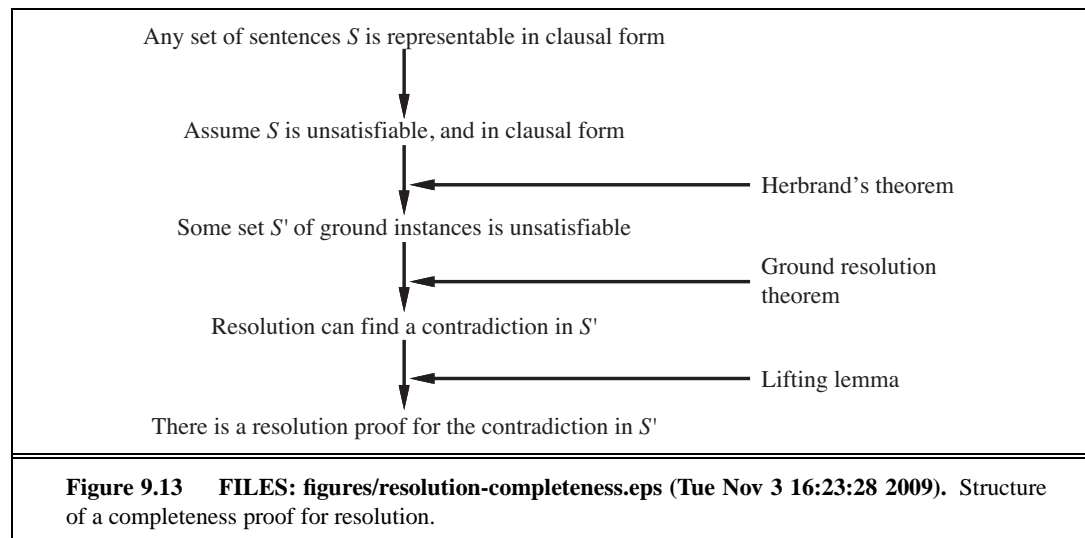




Figure 9.11 FILES: figures/crime-resolution.eps (Tue Nov 3 16:22:35 2009). A resolution proof that West is a criminal. At each step, the literals that unify are in bold.





10 CLASSICAL PLANNING






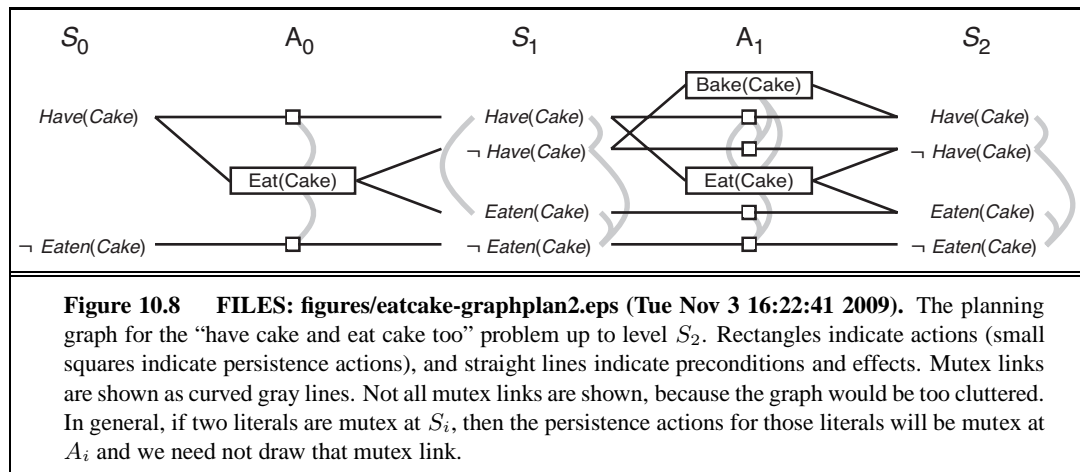


Figure 10.6 FILES: figures/ignore-del.eps (Tue Nov 3 16:23:04 2009). Two state spaces from planning problems with the ignore-delete-lists heuristic. The height above the bottom plane is the heuristic score of a state; states on the bottom plane are goals. There are no local minima, so search for the goal is straightforward. From ? (?).






Figure 10.10 FILES: figures/tire-graphplan2.eps (Tue Nov 3 16:23:55 2009). The planning graph for the spare tire problem after expansion to level S_2 . Mutex links are shown as gray lines. Not all links are shown, because the graph would be too cluttered if we showed them all. The solution is indicated by bold lines and outlines.



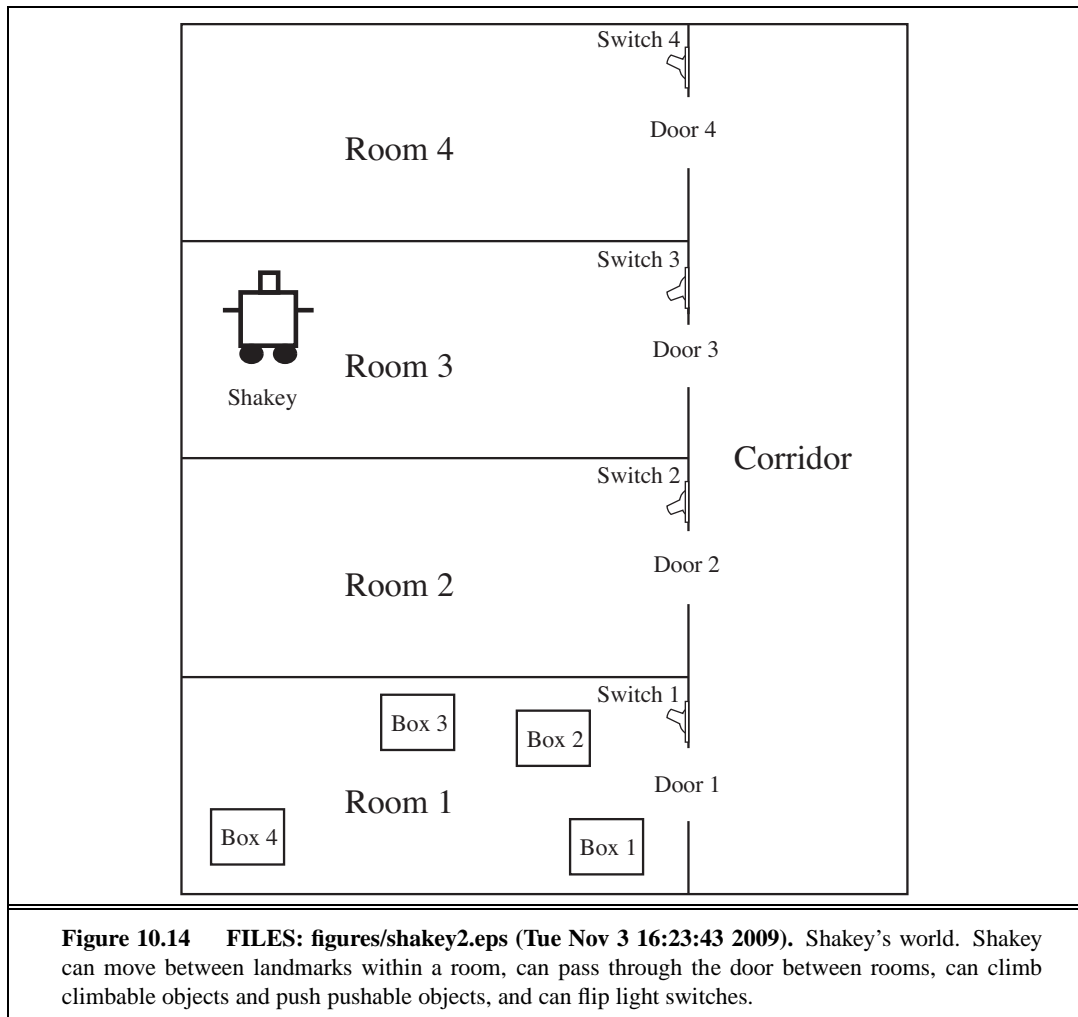



Figure 10.12 FILES: figures/situations.eps (Tue Nov 3 16:23:45 2009). Situations as the results of actions in the wumpus world.





11 PLANNING AND ACTING IN THE REAL WORLD




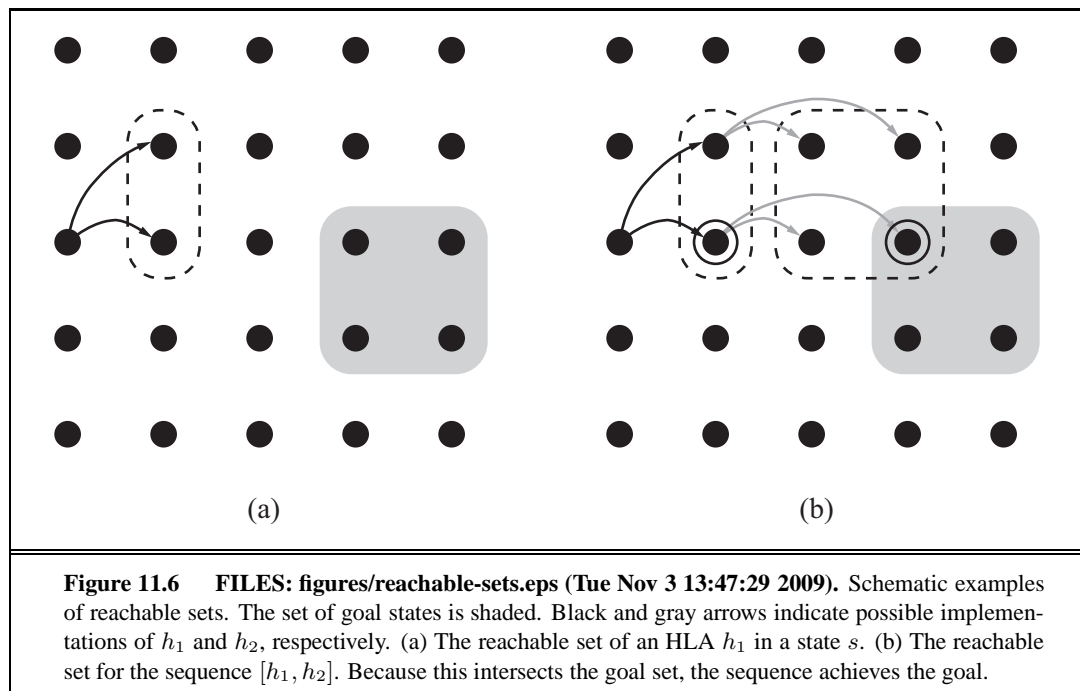



Figure 11.3 FILES: figures/jobshop-resources.eps (Tue Nov 3 16:23:05 2009). A solution to the job-shop scheduling problem from Figure ??, taking into account resource constraints. The left-hand margin lists the three reusable resources, and actions are shown aligned horizontally with the resources they use. There are two possible schedules, depending on which assembly uses the engine hoist first; we've shown the shortest-duration solution, which takes 115 minutes.









Figure 11.9 FILES: figures/plan-repair.eps (Tue Nov 3 16:23:19 2009). Before execution, the planner comes up with a plan, here called *whole plan*, to get from S to G . The agent executes steps of the plan until it expects to be in state E , but observes it is actually in O . The agent then replans for the minimal repair plus continuation to reach G .



12 KNOWLEDGE REPRESENTATION


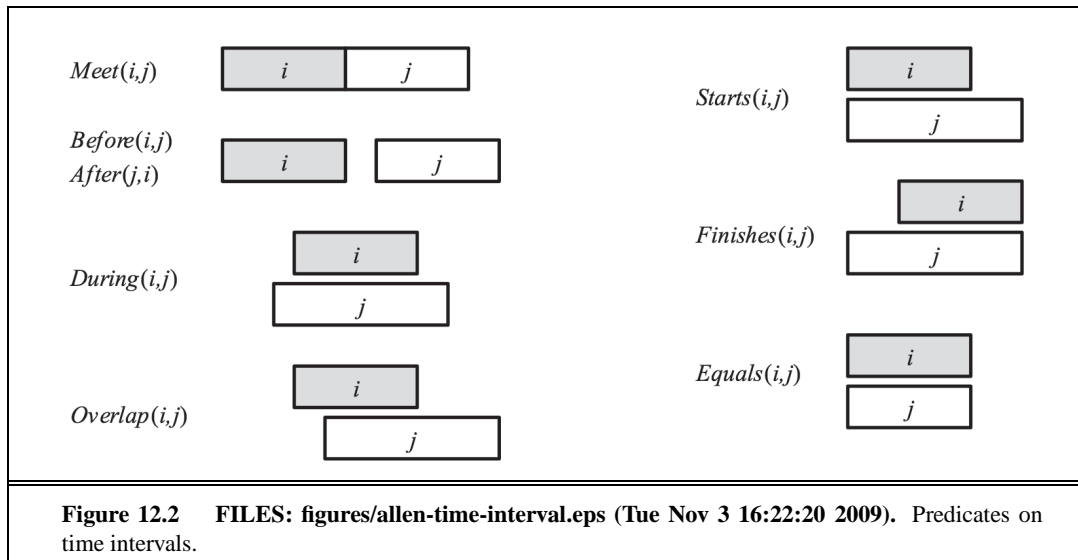


Figure 12.1 FILES: figures/everything.eps (Tue Nov 3 16:22:41 2009). The upper ontology of the world, showing the topics to be covered later in the chapter. Each link indicates that the lower concept is a specialization of the upper one. Specializations are not necessarily disjoint; a human is both an animal and an agent, for example. We will see in Section ?? why physical objects come under generalized events.







Figure 12.3 FILES: figures/president-usa.eps (Tue Nov 3 16:23:22 2009). A schematic view of the object *President(USA)* for the first 15 years of its existence.






Figure 12.5 FILES: figures/semantic-network.eps (Tue Nov 3 16:23:41 2009). A semantic network with four objects (John, Mary, 1, and 2) and four categories. Relations are denoted by labeled links.



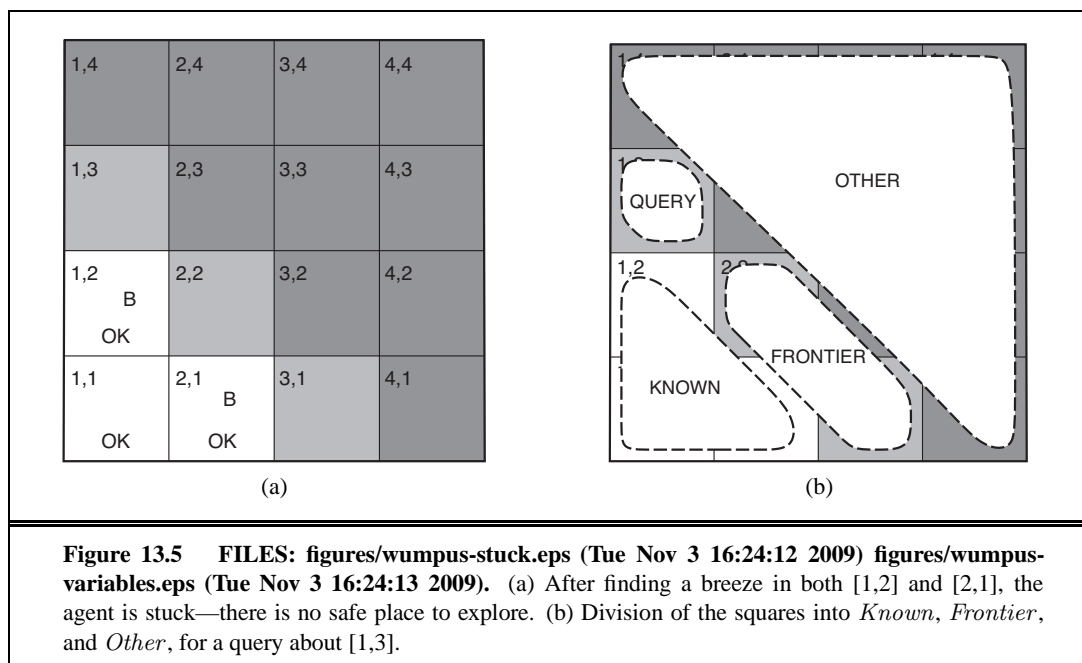



Figure 12.6 FILES: figures/flying-network.eps (Tue Nov 3 16:22:52 2009). A fragment of a semantic network showing the representation of the logical assertion $Fly(Shankar, NewYork, NewDelhi, Yesterday)$.

13 QUANTIFYING UNCERTAINTY

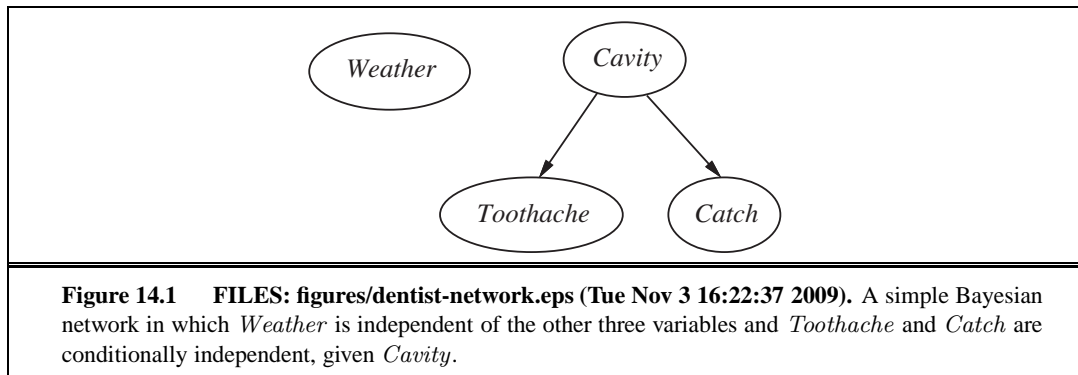








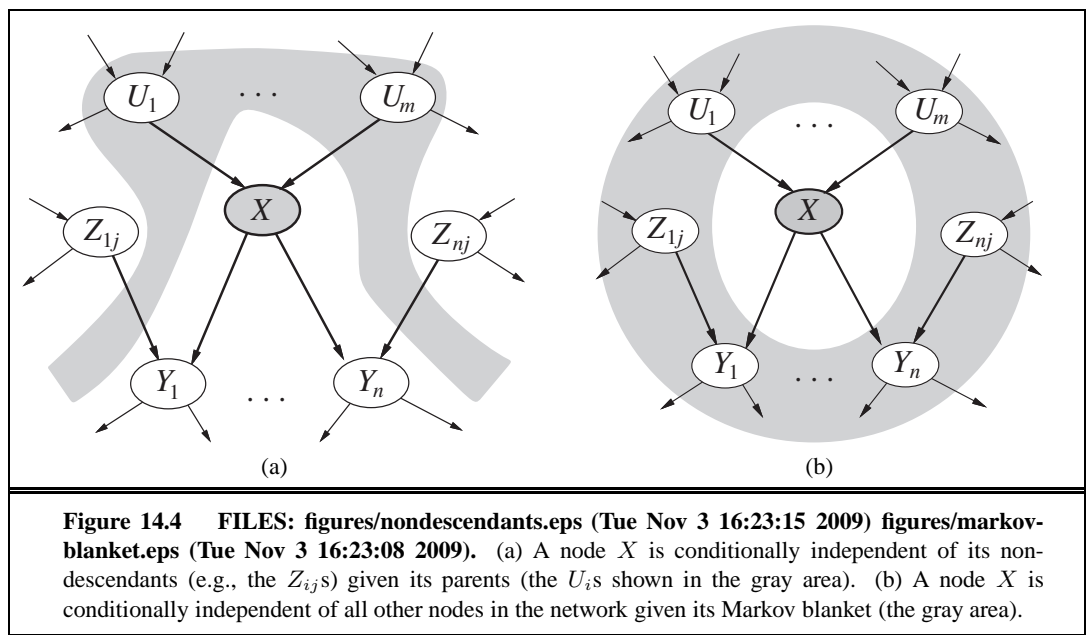
14


PROBABILISTIC REASONING




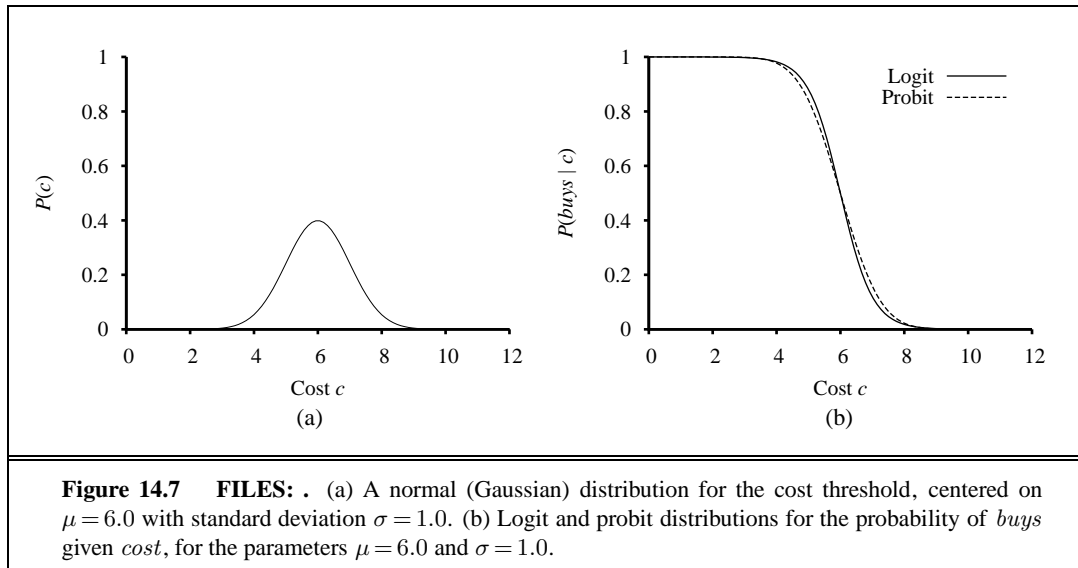















Figure 14.8 FILES: figures/enumeration-tree.eps (Tue Nov 3 16:22:41 2009). The structure of the expression shown in Equation (??). The evaluation proceeds top down, multiplying values along each path and summing at the “+” nodes. Notice the repetition of the paths for j and m .




Figure 14.12 FILES: figures/rain-clustering1.eps (Tue Nov 3 16:23:27 2009) figures/rain-clustering2.eps (Tue Nov 3 16:23:27 2009). (a) A multiply connected network with conditional probability tables. (b) A clustered equivalent of the multiply connected network.

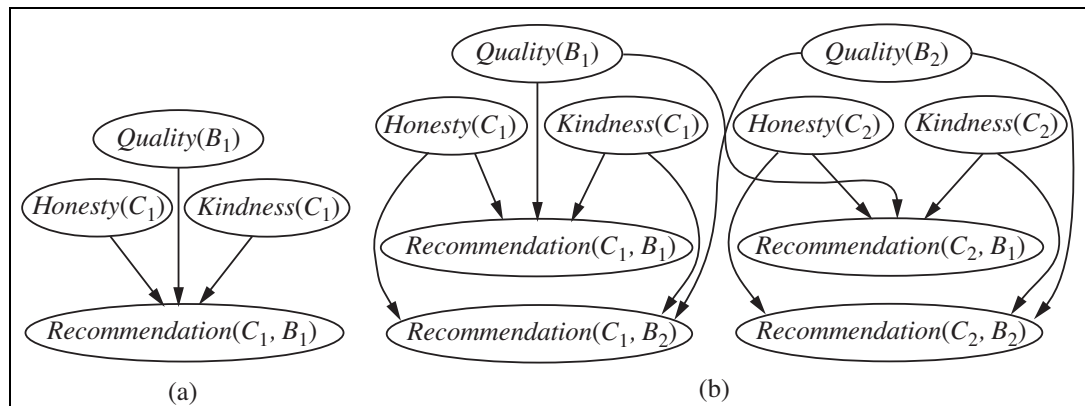


Figure 14.17 FILES: figures/new-14-16-1.eps (Tue Nov 3 16:23:14 2009) figures/new-14-16-2.eps (Tue Nov 3 16:23:14 2009) figures/new-14-16-1.eps (Tue Nov 3 16:23:14 2009). (a) Bayes net for a single customer C_1 recommending a single book B_1 . $Honest(C_1)$ is Boolean, while the other variables have integer values from 1 to 5. (b) Bayes net with two customers and two books.

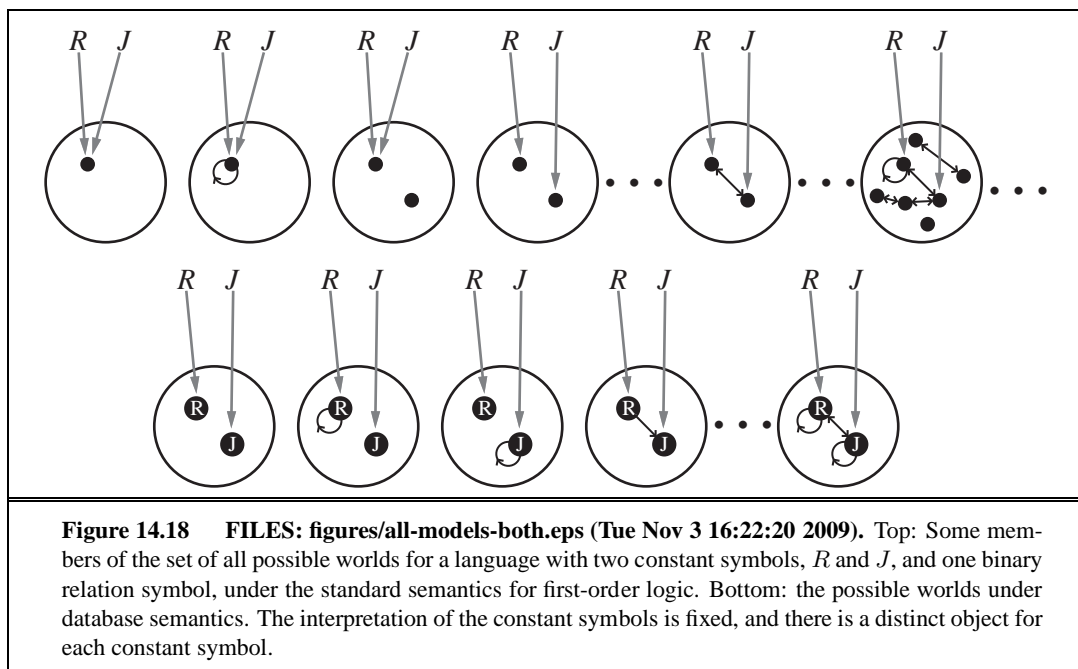

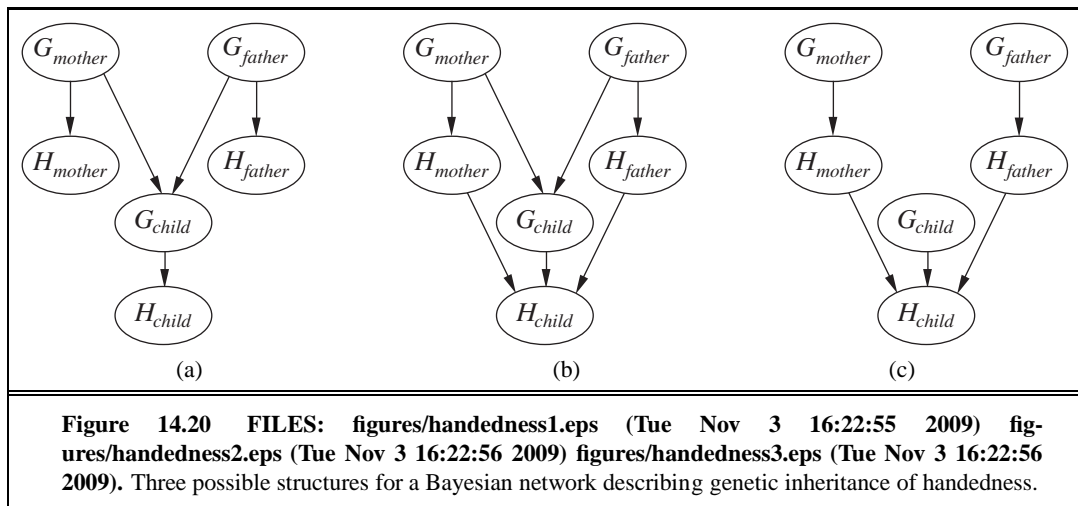


Figure 14.18 FILES: figures/all-models-both.eps (Tue Nov 3 16:22:20 2009). Top: Some members of the set of all possible worlds for a language with two constant symbols, R and J , and one binary relation symbol, under the standard semantics for first-order logic. Bottom: the possible worlds under database semantics. The interpretation of the constant symbols is fixed, and there is a distinct object for each constant symbol.






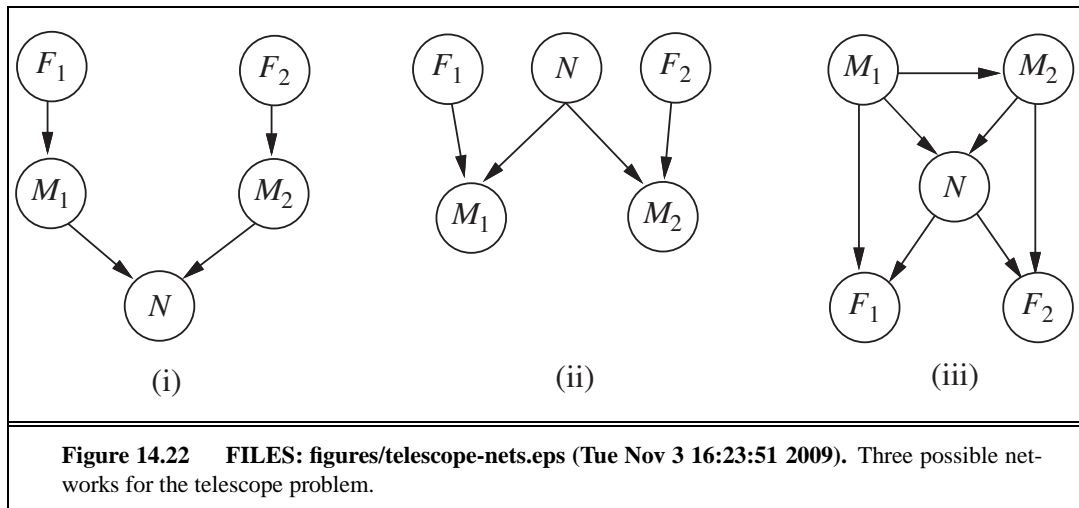


Figure 14.21 FILES: figures/car-starts.eps (Tue Nov 3 16:22:32 2009). A Bayesian network describing some features of a car's electrical system and engine. Each variable is Boolean, and the *true* value indicates that the corresponding aspect of the vehicle is in working order.



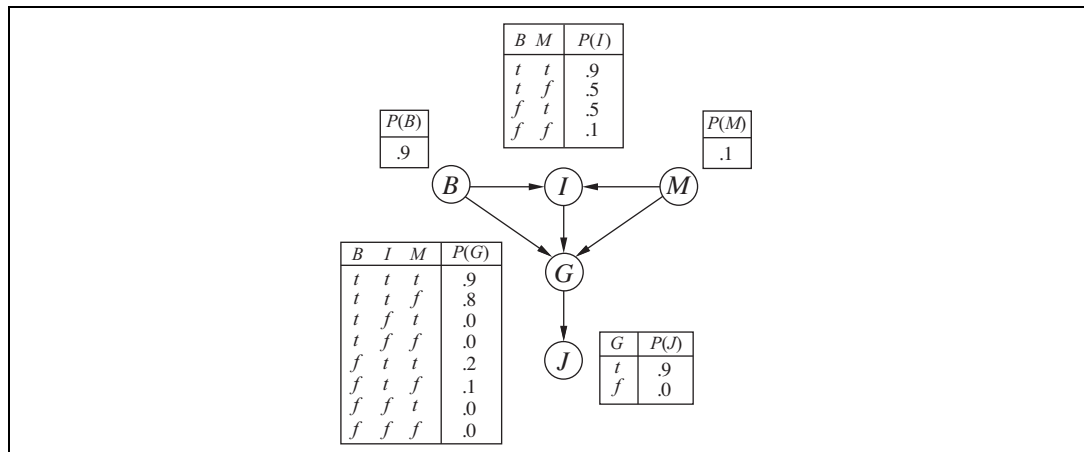

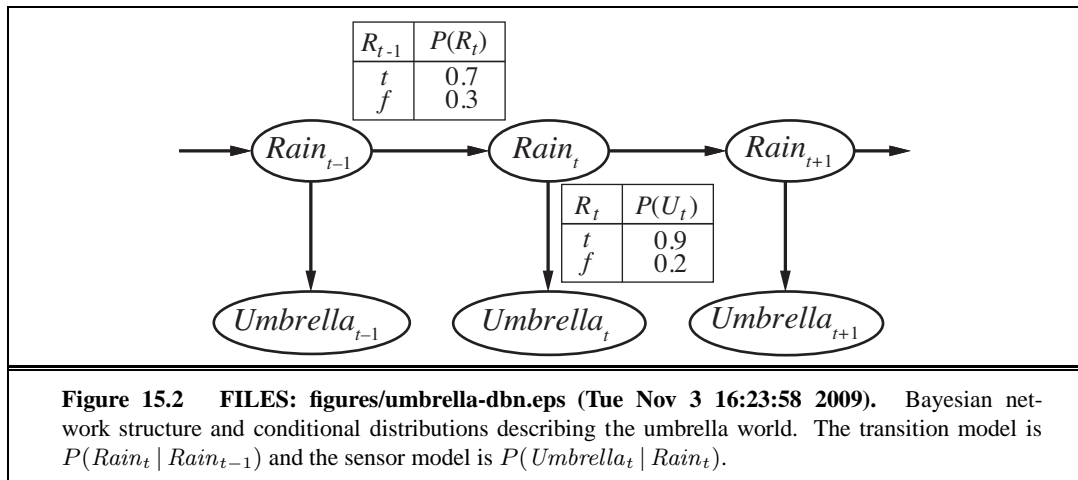



Figure 14.23 FILES: figures/politics.eps (Tue Nov 3 16:23:20 2009). A simple Bayes net with Boolean variables $B = BrokeElectionLaw$, $I = Indicted$, $M = PoliticallyMotivatedProsecutor$, $G = FoundGuilty$, $J = Jailed$.


15


PROBABILISTIC REASONING OVER TIME













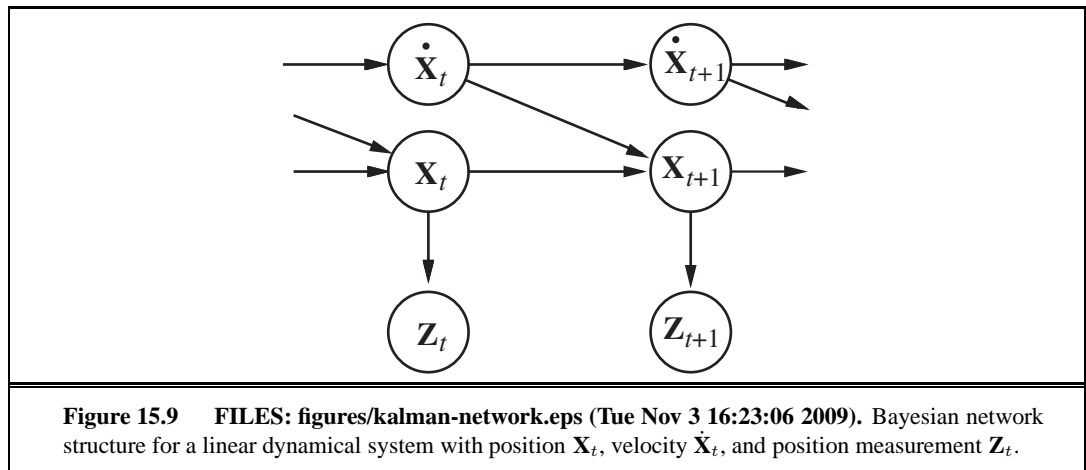
(a) Posterior distribution over robot location after $E_1 = NSW$



(b) Posterior distribution over robot location after $E_1 = NSW, E_2 = NS$

Figure 15.7 FILES: figures/localization-figures-b.eps (Tue Nov 3 16:23:07 2009). Posterior distribution over robot location: (a) one observation $E_1 = NSW$; (b) after a second observation $E_2 = NS$. The size of each disk corresponds to the probability that the robot is at that location. The sensor error rate is $\epsilon = 0.2$.










Figure 15.10 FILES: . Stages in the Kalman filter update cycle for a random walk with a prior given by $\mu_0 = 0.0$ and $\sigma_0 = 1.0$, transition noise given by $\sigma_x = 2.0$, sensor noise given by $\sigma_z = 1.0$, and a first observation $z_1 = 2.5$ (marked on the x -axis). Notice how the prediction $P(x_1)$ is flattened out, relative to $P(x_0)$, by the transition noise. Notice also that the mean of the posterior distribution $P(x_1 | z_1)$ is slightly to the left of the observation z_1 because the mean is a weighted average of the prediction and the observation.





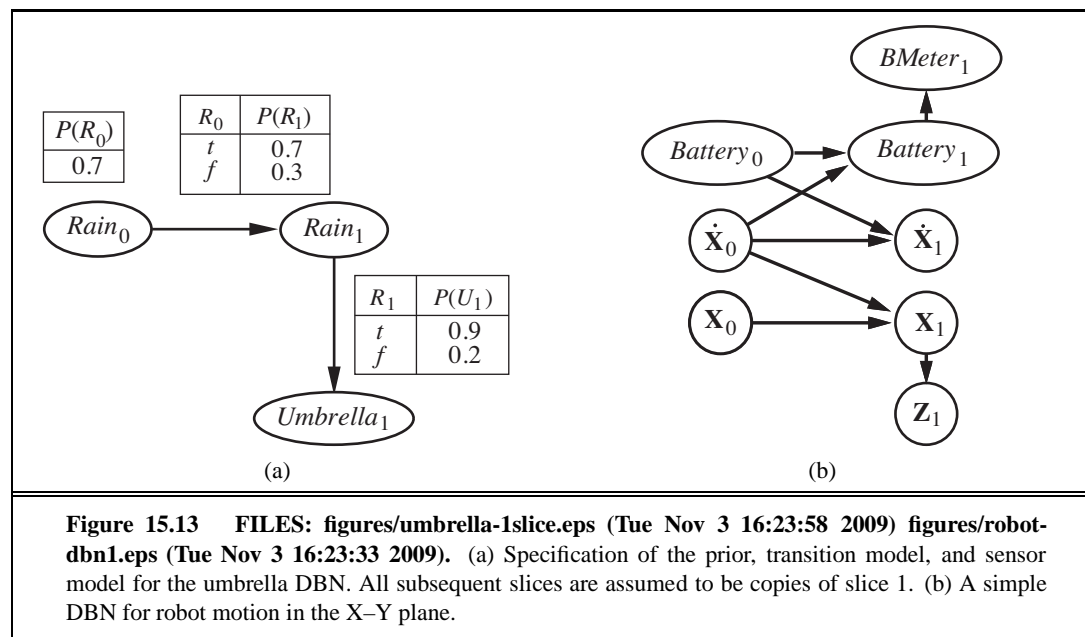


Figure 15.13 FILES: figures/umbrella-1slice.eps (Tue Nov 3 16:23:58 2009) figures/robot-dbn1.eps (Tue Nov 3 16:23:33 2009). (a) Specification of the prior, transition model, and sensor model for the umbrella DBN. All subsequent slices are assumed to be copies of slice 1. (b) A simple DBN for robot motion in the X-Y plane.




Figure 15.14 FILES: . (a) Upper curve: trajectory of the expected value of Battery_t for an observation sequence consisting of all 5s except for 0s at $t=21$ and $t=22$, using a simple Gaussian error model. Lower curve: trajectory when the observation remains at 0 from $t=21$ onwards. (b) The same experiment run with the transient failure model. Notice that the transient failure is handled well, but the persistent failure results in excessive pessimism about the battery charge.


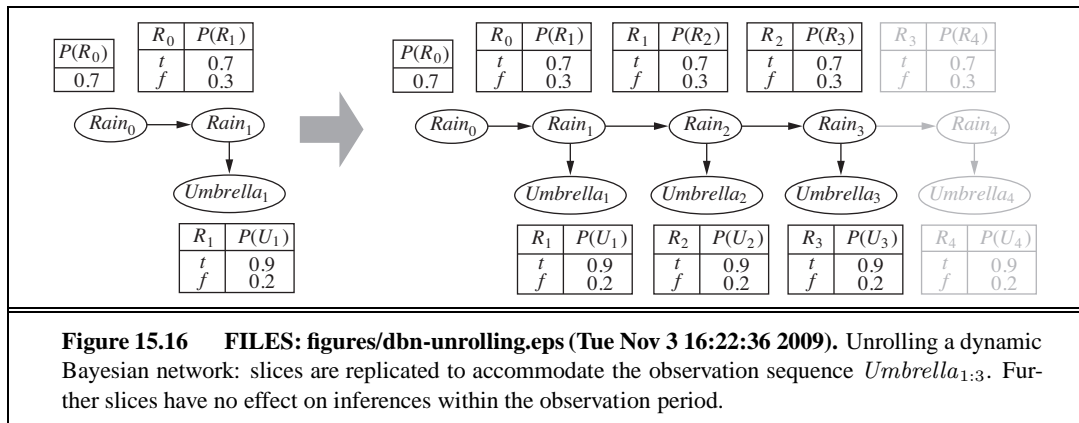


Figure 15.15 FILES: figures/battery-persistence.eps (Tue Nov 3 16:22:26 2009). (a) A DBN fragment showing the sensor status variable required for modeling persistent failure of the battery sensor. (b) Upper curves: trajectories of the expected value of $Battery_t$ for the “transient failure” and “permanent failure” observations sequences. Lower curves: probability trajectories for $BMBroken$ given the two observation sequences.



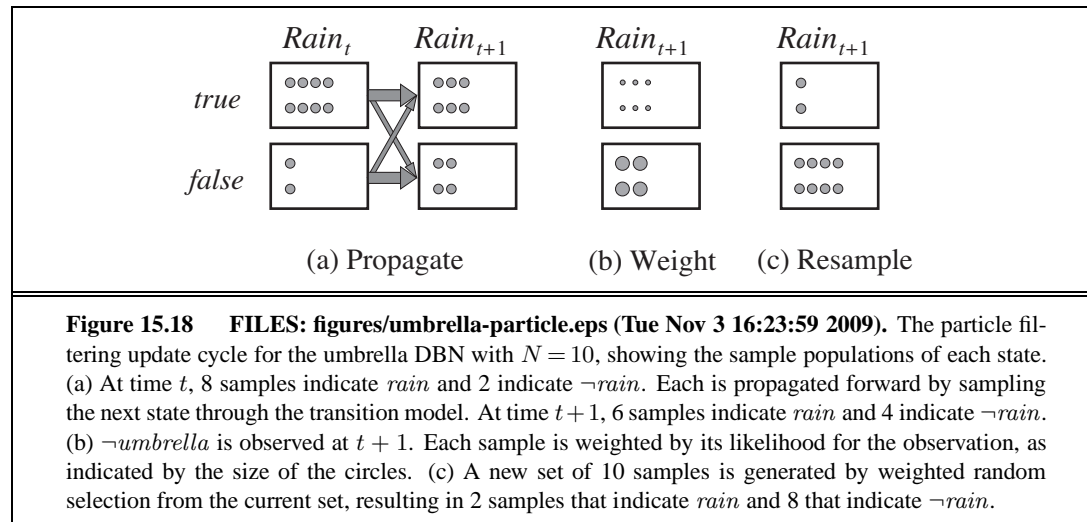

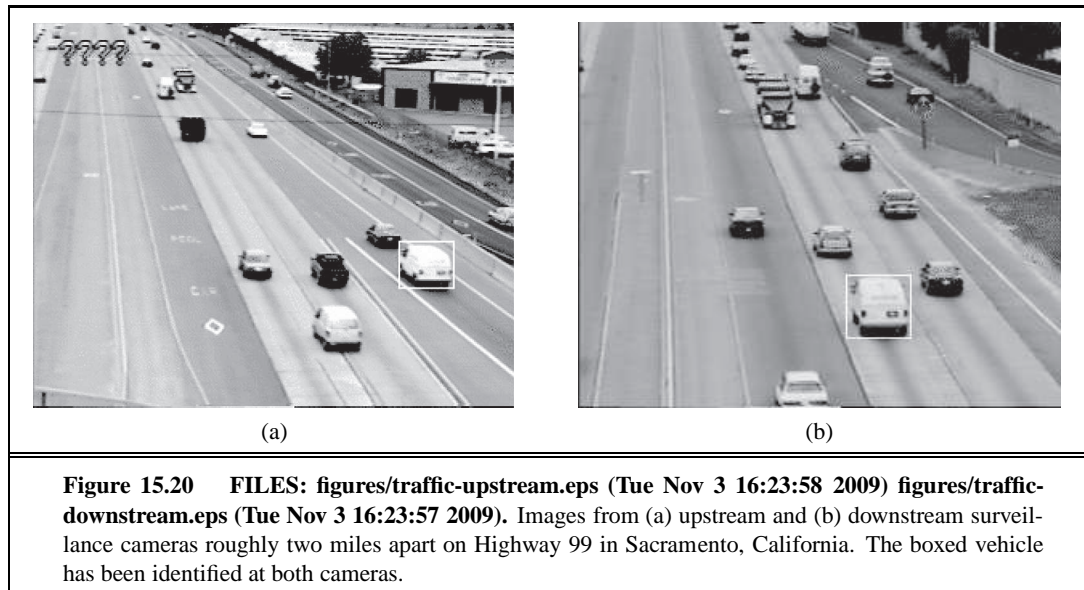
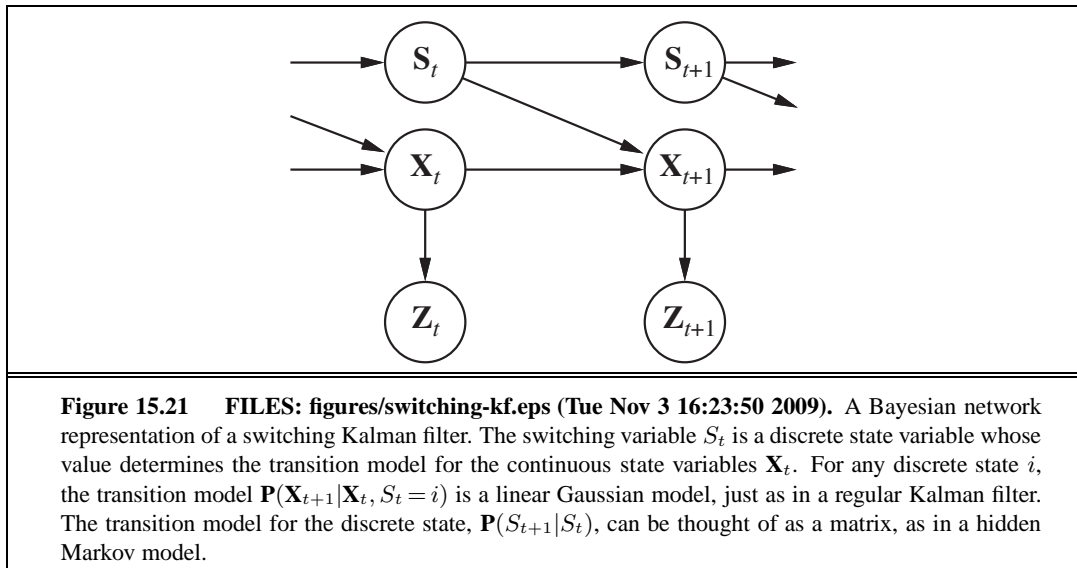


Figure 15.18 FILES: figures/umbrella-particle.eps (Tue Nov 3 16:23:59 2009). The particle filtering update cycle for the umbrella DBN with $N = 10$, showing the sample populations of each state. (a) At time t , 8 samples indicate *rain* and 2 indicate \neg *rain*. Each is propagated forward by sampling the next state through the transition model. At time $t + 1$, 6 samples indicate *rain* and 4 indicate \neg *rain*. (b) \neg *umbrella* is observed at $t + 1$. Each sample is weighted by its likelihood for the observation, as indicated by the size of the circles. (c) A new set of 10 samples is generated by weighted random selection from the current set, resulting in 2 samples that indicate *rain* and 8 that indicate \neg *rain*.







16 MAKING SIMPLE DECISIONS




Figure 16.1 FILES: figures/cash-machine-and-decomposability.eps (Tue Nov 3 13:30:24 2009).
 (a) A cycle of exchanges showing that the nontransitive preferences $A \succ B \succ C \succ A$ result in irrational behavior. (b) The decomposability axiom.


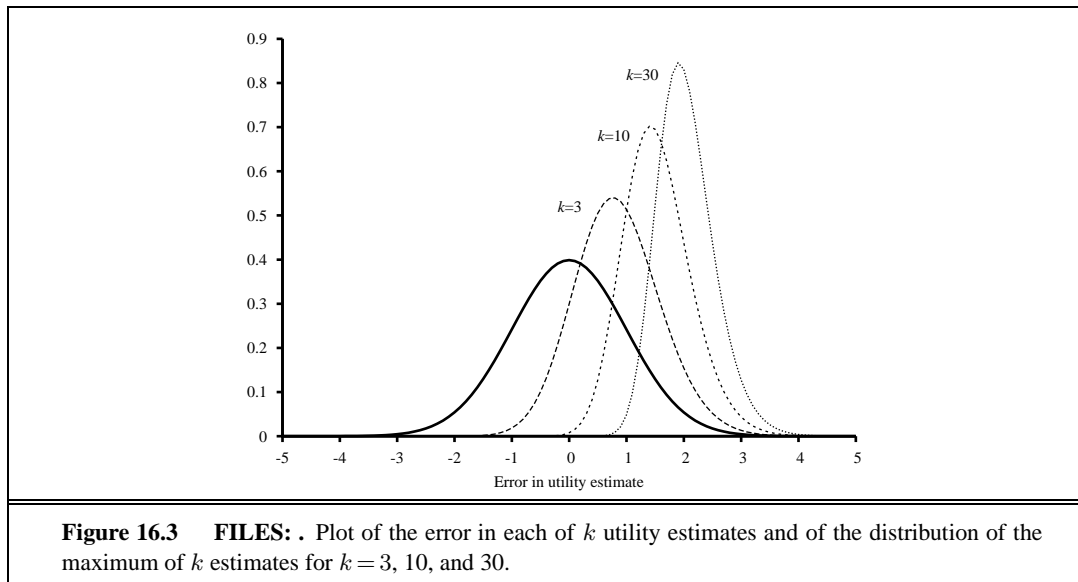


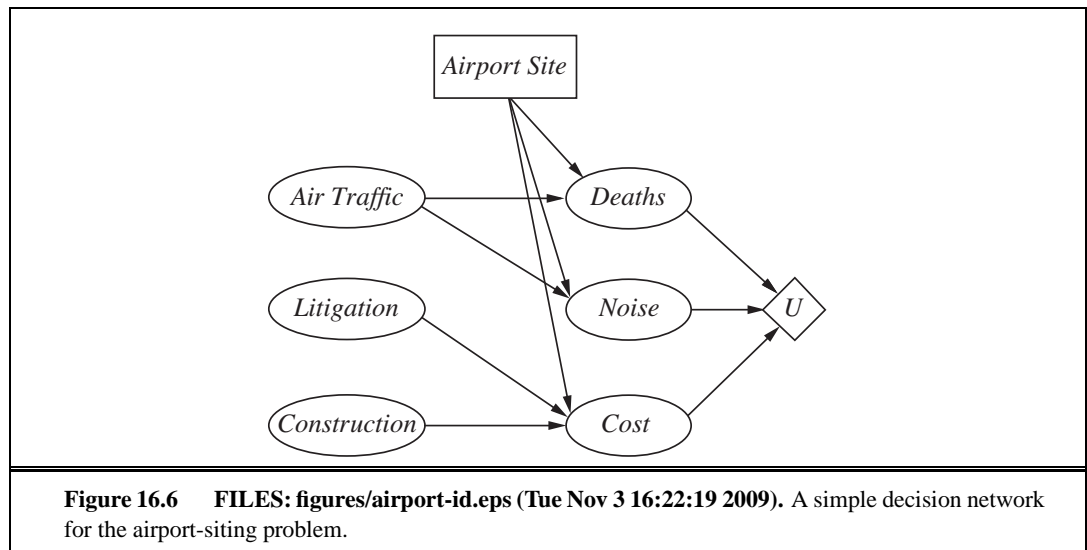



Figure 16.2 FILES: figures/utility-curve.eps (Tue Nov 3 16:24:00 2009). The utility of money. (a) Empirical data for Mr. Beard over a limited range. (b) A typical curve for the full range.














Figure 16.8 FILES: figures/3cases.eps (Tue Nov 3 16:22:10 2009). Three generic cases for the value of information. In (a), a_1 will almost certainly remain superior to a_2 , so the information is not needed. In (b), the choice is unclear and the information is crucial. In (c), the choice is unclear, but because it makes little difference, the information is less valuable. (Note: The fact that U_2 has a high peak in (c) means that its expected value is known with higher certainty than U_1 .)

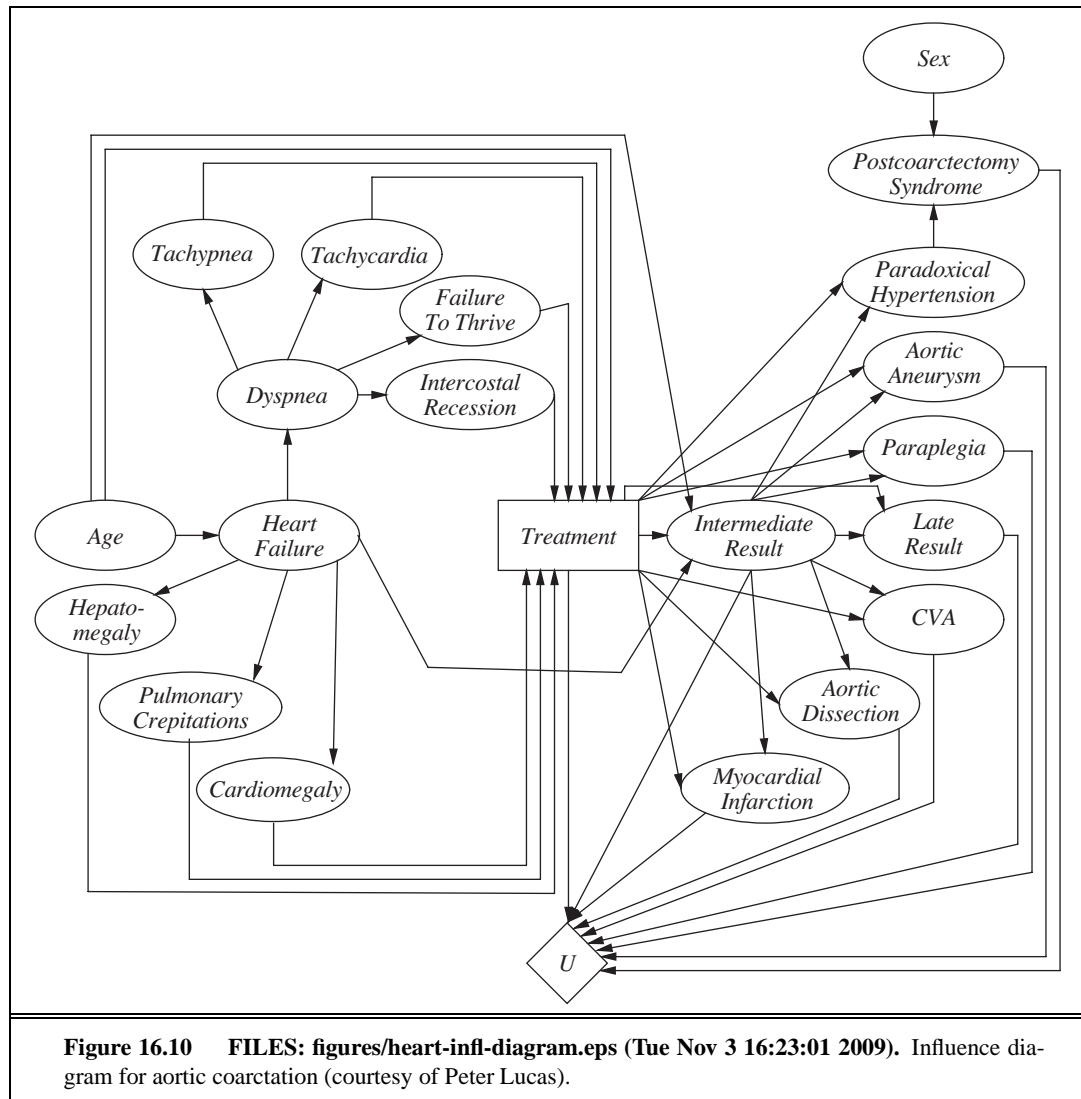


Figure 16.10 FILES: figures/heart-infl-diagram.eps (Tue Nov 3 16:23:01 2009). Influence diagram for aortic coarctation (courtesy of Peter Lucas).

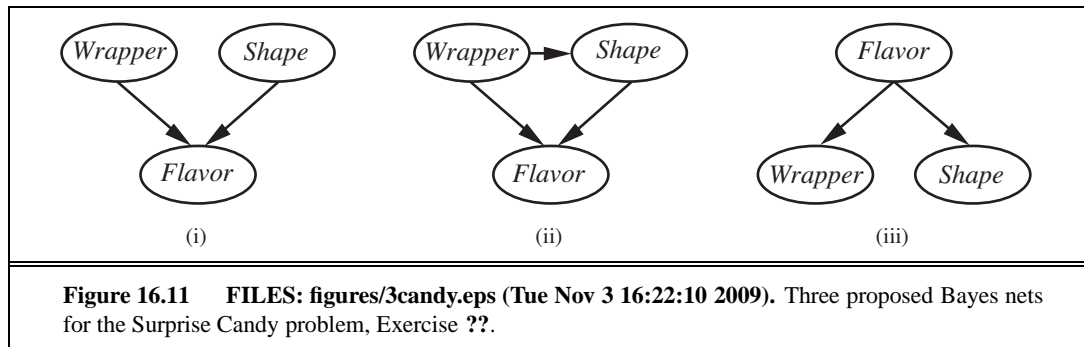





Figure 16.11 FILES: figures/3candy.eps (Tue Nov 3 16:22:10 2009). Three proposed Bayes nets for the Surprise Candy problem, Exercise ??.



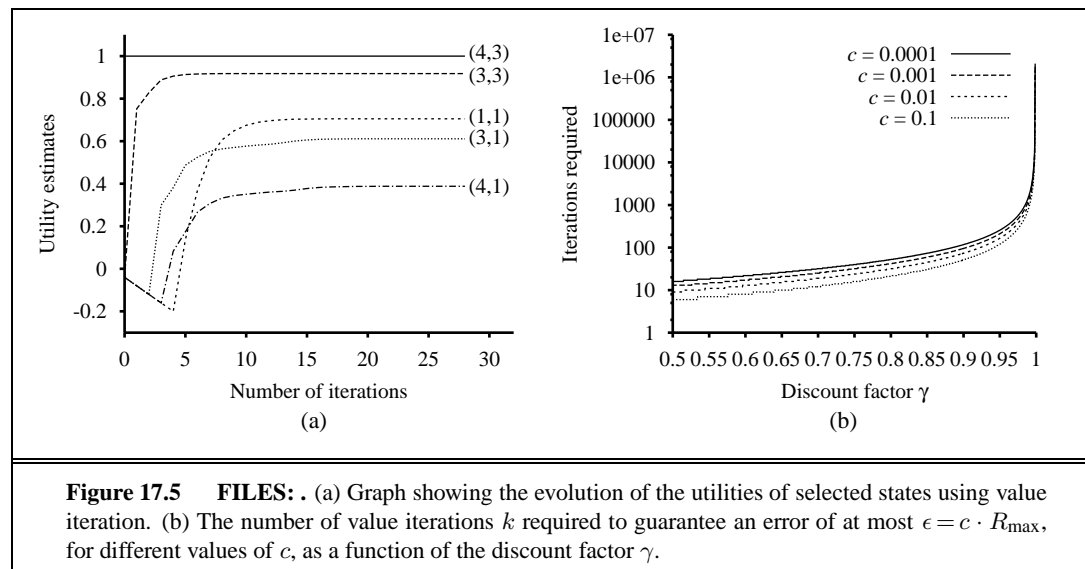
17 MAKING COMPLEX DECISIONS







3	0.812	0.868	0.918	+1
2	0.762		0.660	-1
1	0.705	0.655	0.611	0.388
	1	2	3	4

Figure 17.3 FILES: figures/sequential-decision-values.eps (Tue Nov 3 16:23:42 2009). The utilities of the states in the 4×3 world, calculated with $\gamma = 1$ and $R(s) = -0.04$ for nonterminal states.










Figure 17.10 FILES: figures/generic-ddn.eps (Tue Nov 3 16:22:53 2009). The generic structure of a dynamic decision network. Variables with known values are shaded. The current time is t and the agent must decide what to do—that is, choose a value for A_t . The network has been unrolled into the future for three steps and represents future rewards, as well as the utility of the state at the look-ahead horizon.




Figure 17.11 FILES: figures/pomdp-tree.eps (Tue Nov 3 16:23:20 2009). Part of the look-ahead solution of the DDN in Figure 17.10. Each decision will be taken in the belief state indicated.

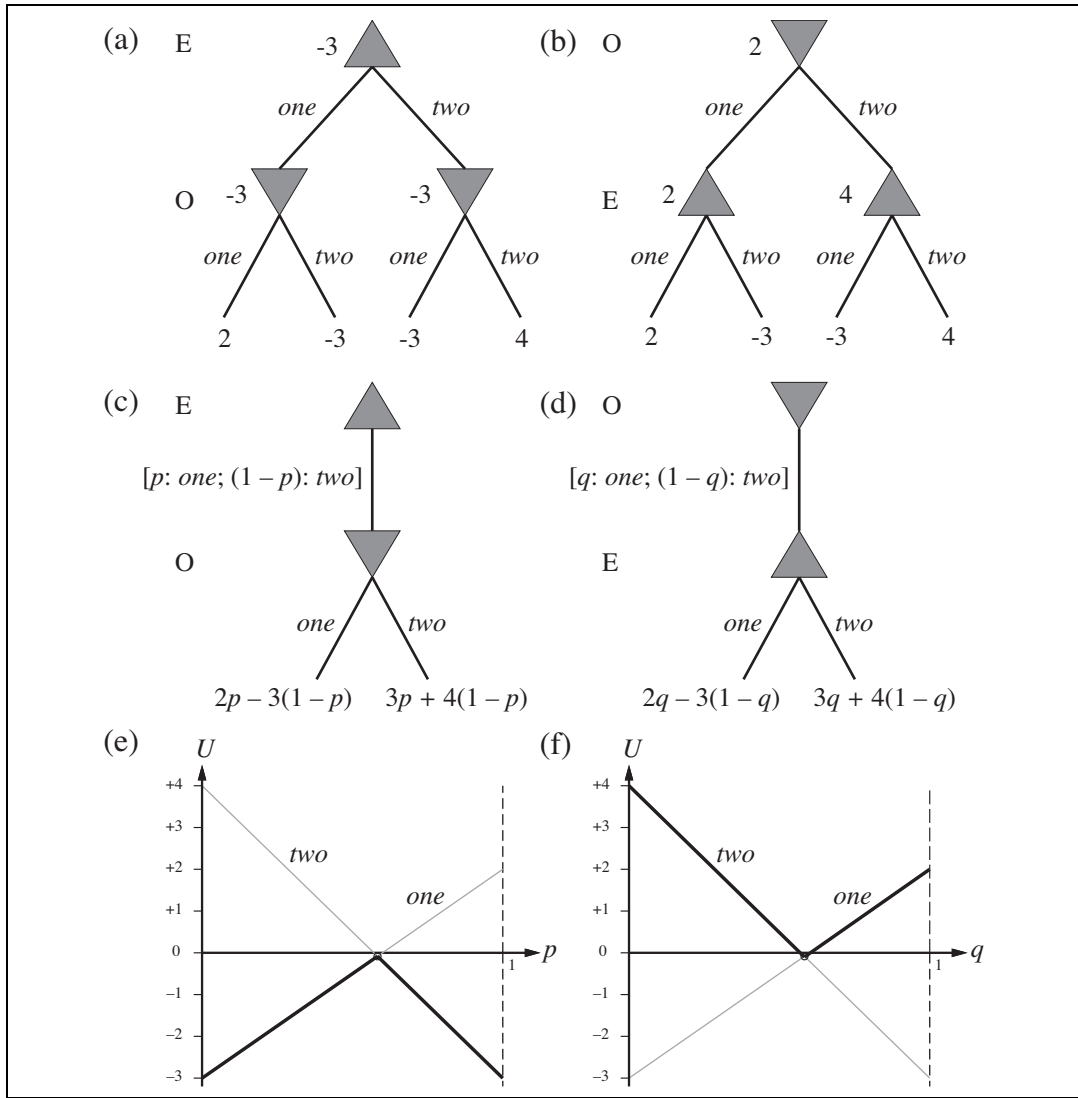





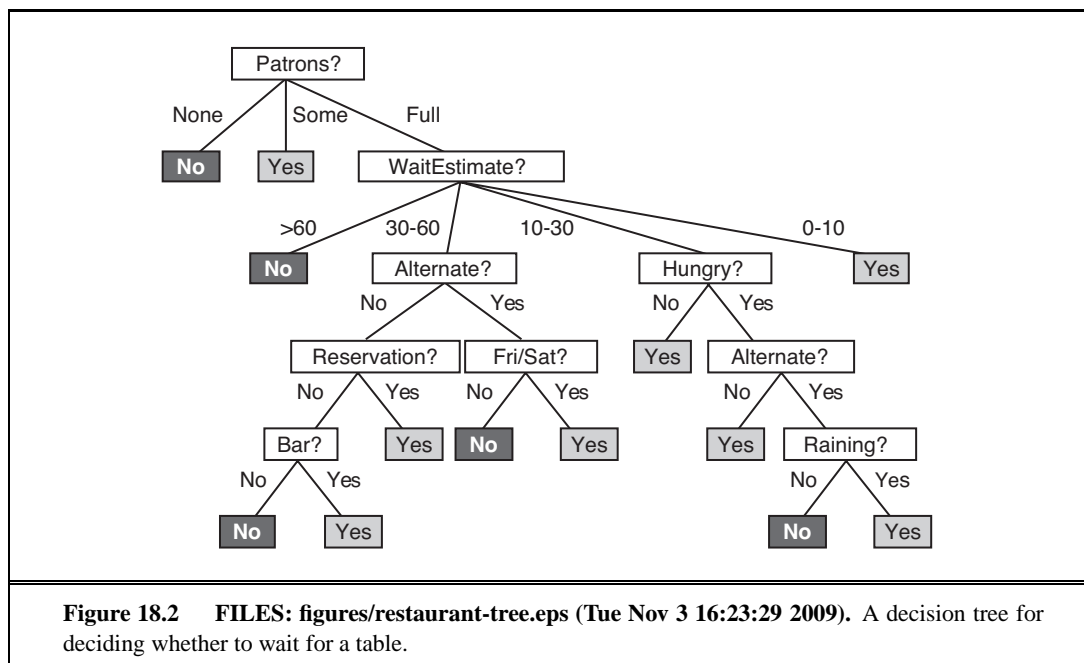
Figure 17.12 FILES: figures/morra-trees.eps (Tue Nov 3 16:23:11 2009). (a) and (b): Minimax game trees for two-finger Morra if the players take turns playing pure strategies. (c) and (d): Parameterized game trees where the first player plays a mixed strategy. The payoffs depend on the probability parameter (p or q) in the mixed strategy. (e) and (f): For any particular value of the probability parameter, the second player will choose the “better” of the two actions, so the value of the first player’s mixed strategy is given by the heavy lines. The first player will choose the probability parameter for the mixed strategy at the intersection point.








18 LEARNING FROM EXAMPLES





























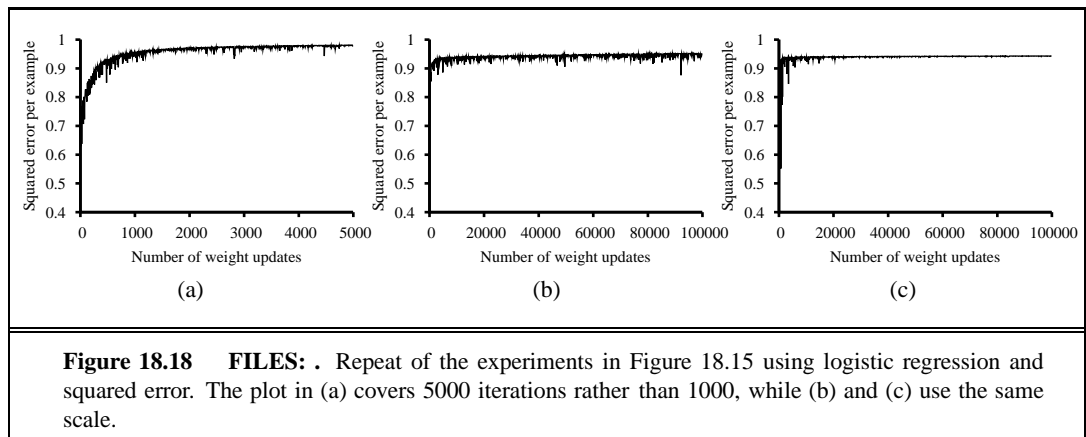



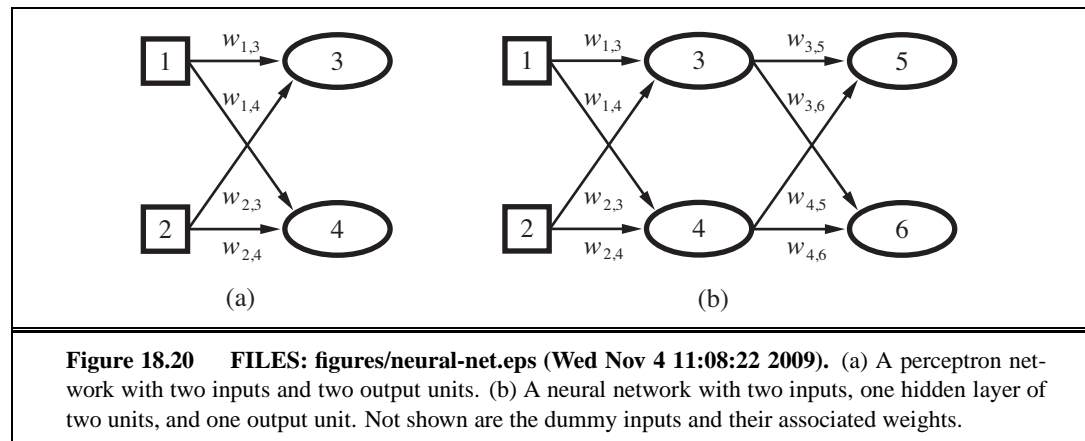
Figure 18.15 FILES: . (a) Plot of two seismic data parameters, body wave magnitude x_1 and surface wave magnitude x_2 , for earthquakes (white circles) and nuclear explosions (black circles) occurring between 1982 and 1990 in Asia and the Middle East (?). Also shown is a decision boundary between the classes. (b) The same domain with more data points. The earthquakes and explosions are no longer linearly separable.















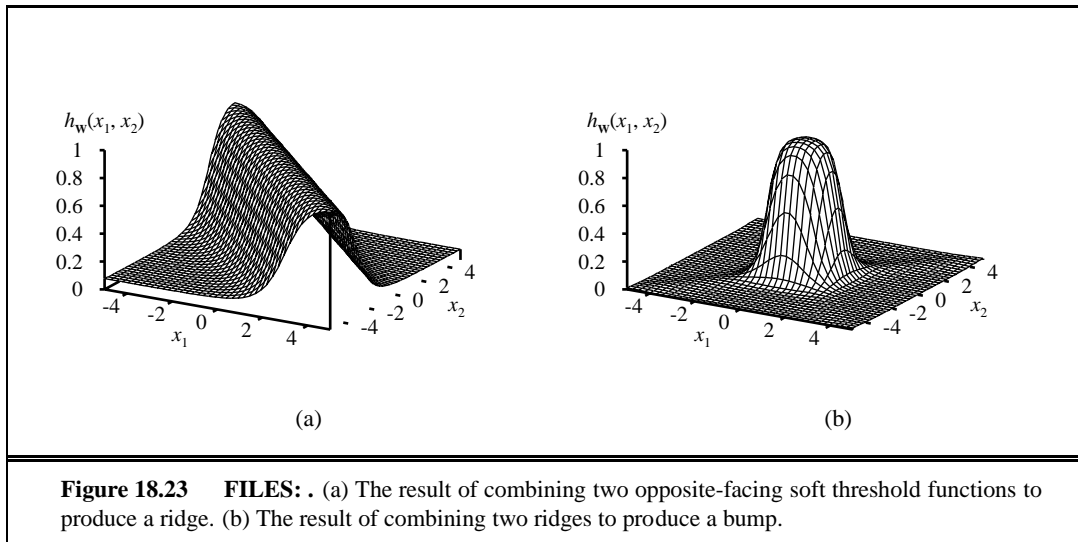


Figure 18.22 FILES: . Comparing the performance of perceptrons and decision trees. (a) Perceptrons are better at learning the majority function of 11 inputs. (b) Decision trees are better at learning the *WillWait* predicate in the restaurant example.





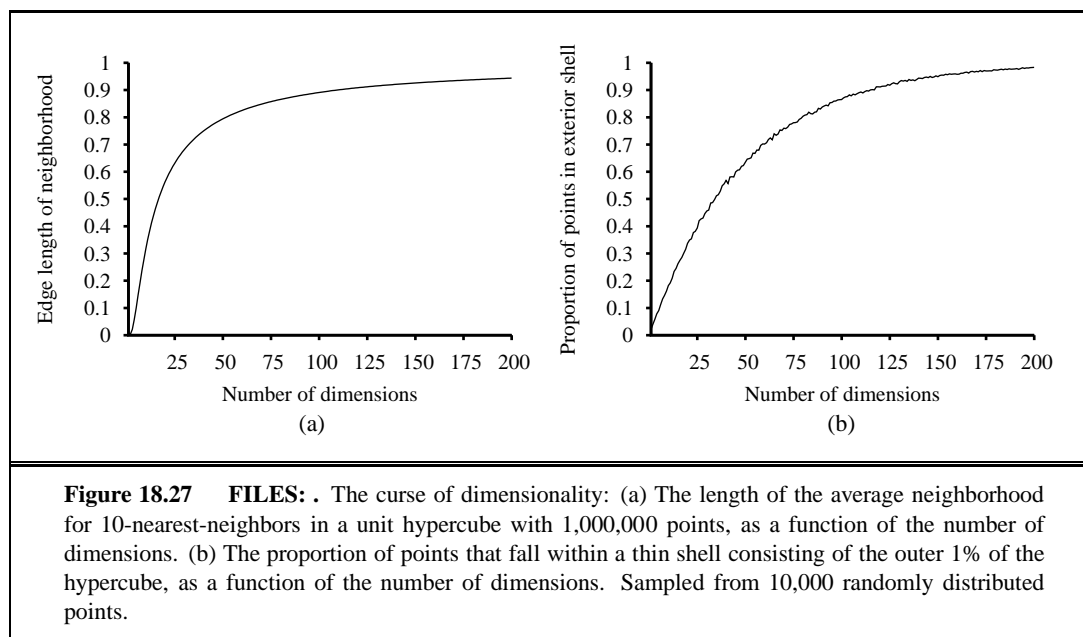
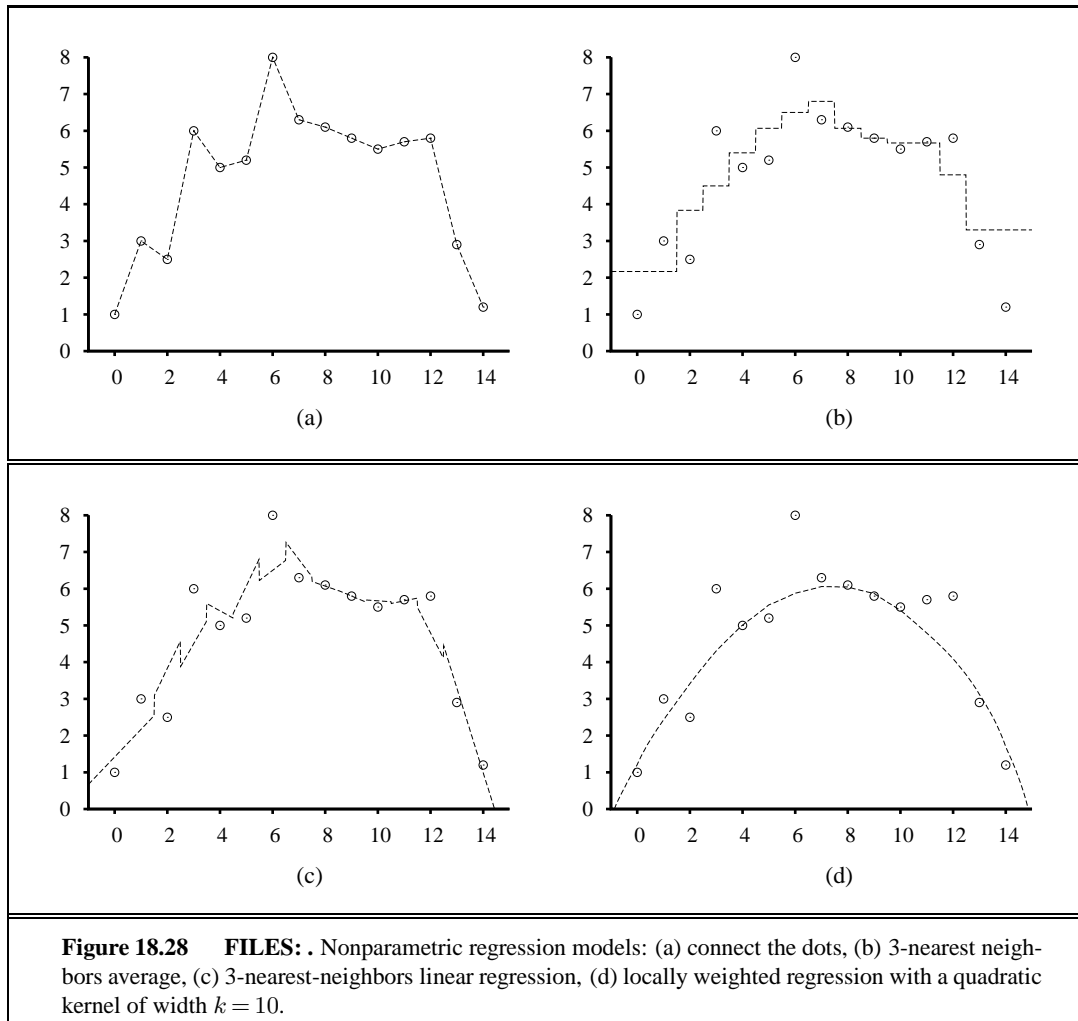
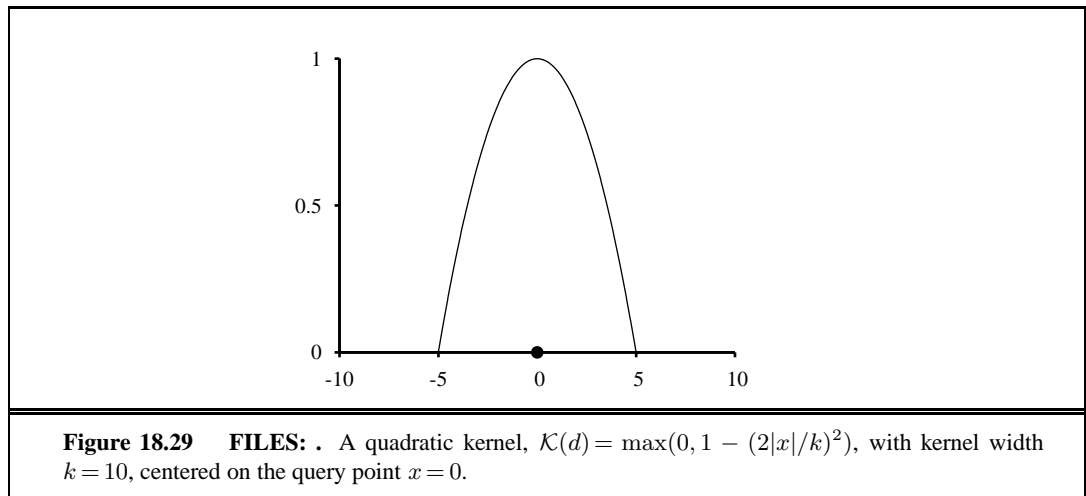



Figure 18.25 FILES: . (a) Training curve showing the gradual reduction in error as weights are modified over several epochs, for a given set of examples in the restaurant domain. (b) Comparative learning curves showing that decision-tree learning does slightly better on the restaurant problem than back-propagation in a multilayer network.











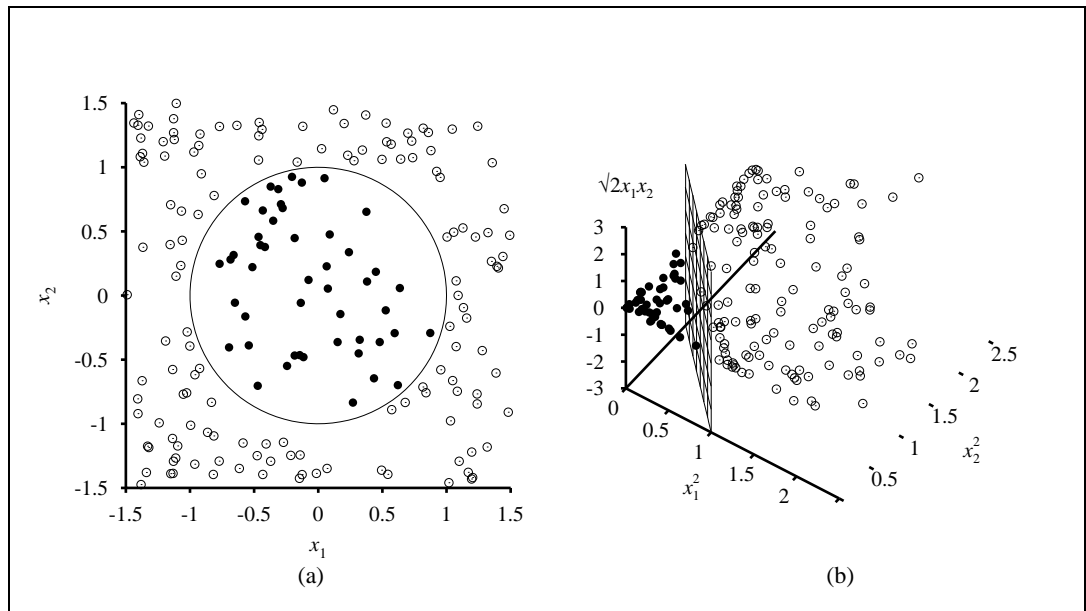


Figure 18.31 FILES: . (a) A two-dimensional training set with positive examples as black circles and negative examples as white circles. The true decision boundary, $x_1^2 + x_2^2 \leq 1$, is also shown. (b) The same data after mapping into a three-dimensional input space $(x_1^2, x_2^2, \sqrt{2}x_1x_2)$. The circular decision boundary in (a) becomes a linear decision boundary in three dimensions. Figure 18.29(b) gives a closeup of the separator in (b).




Figure 18.32 FILES: figures/ensemble-expressiveness.eps (Tue Nov 3 16:22:41 2009). Illustration of the increased expressive power obtained by ensemble learning. We take three linear threshold hypotheses, each of which classifies positively on the unshaded side, and classify as positive any example classified positively by all three. The resulting triangular region is a hypothesis not expressible in the original hypothesis space.

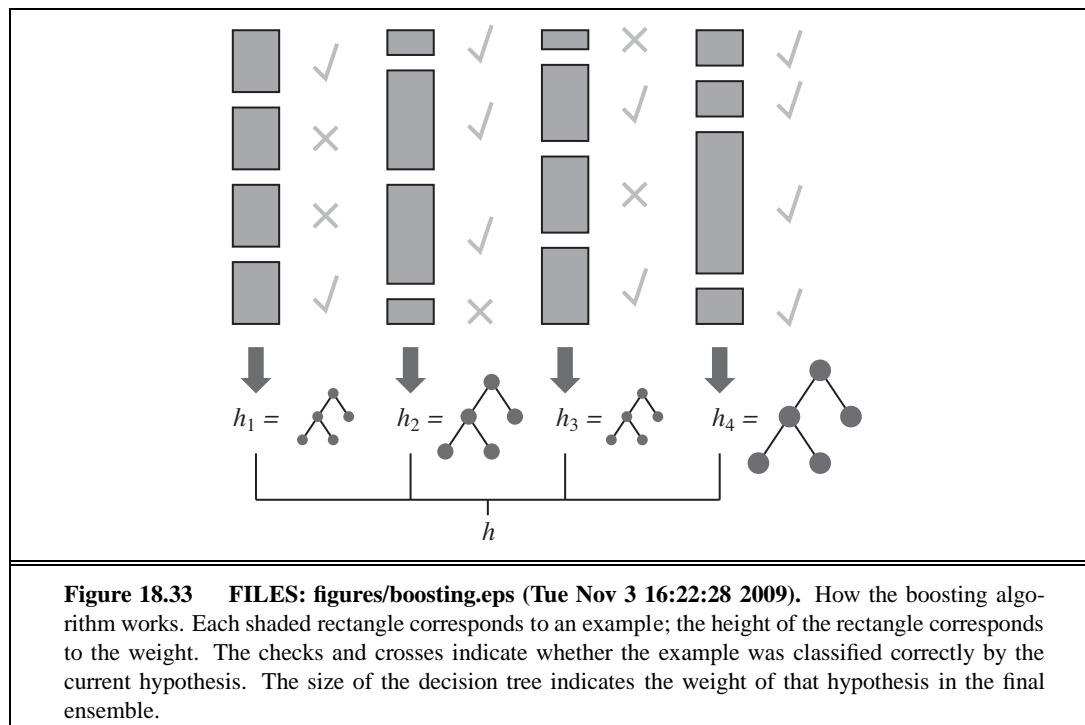



Figure 18.33 FILES: figures/boosting.eps (Tue Nov 3 16:22:28 2009). How the boosting algorithm works. Each shaded rectangle corresponds to an example; the height of the rectangle corresponds to the weight. The checks and crosses indicate whether the example was classified correctly by the current hypothesis. The size of the decision tree indicates the weight of that hypothesis in the final ensemble.






Figure 18.36 FILES: figures/easy-hard.eps (Wed Nov 4 15:38:34 2009). Examples from the NIST database of handwritten digits. Top row: examples of digits 0–9 that are easy to identify. Bottom row: more difficult examples of the same digits.




Figure 18.37 FILES: . Learning curves for five learning algorithms on a common task. Note that there appears to be more room for improvement in the horizontal direction (more training data) than in the vertical direction (different machine learning algorithm). Adapted from ? (?).

19 KNOWLEDGE IN LEARNING

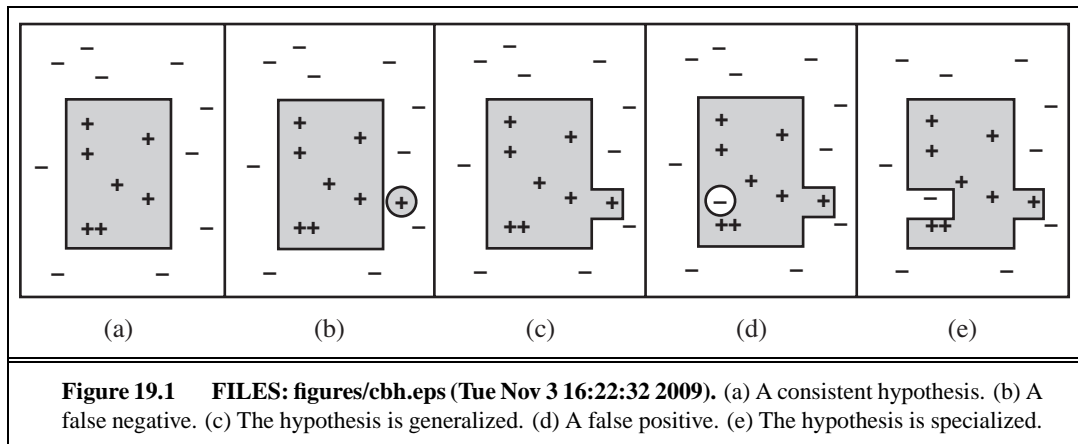


Figure 19.1 FILES: figures/cbh.eps (Tue Nov 3 16:22:32 2009). (a) A consistent hypothesis. (b) A false negative. (c) The hypothesis is generalized. (d) A false positive. (e) The hypothesis is specialized.




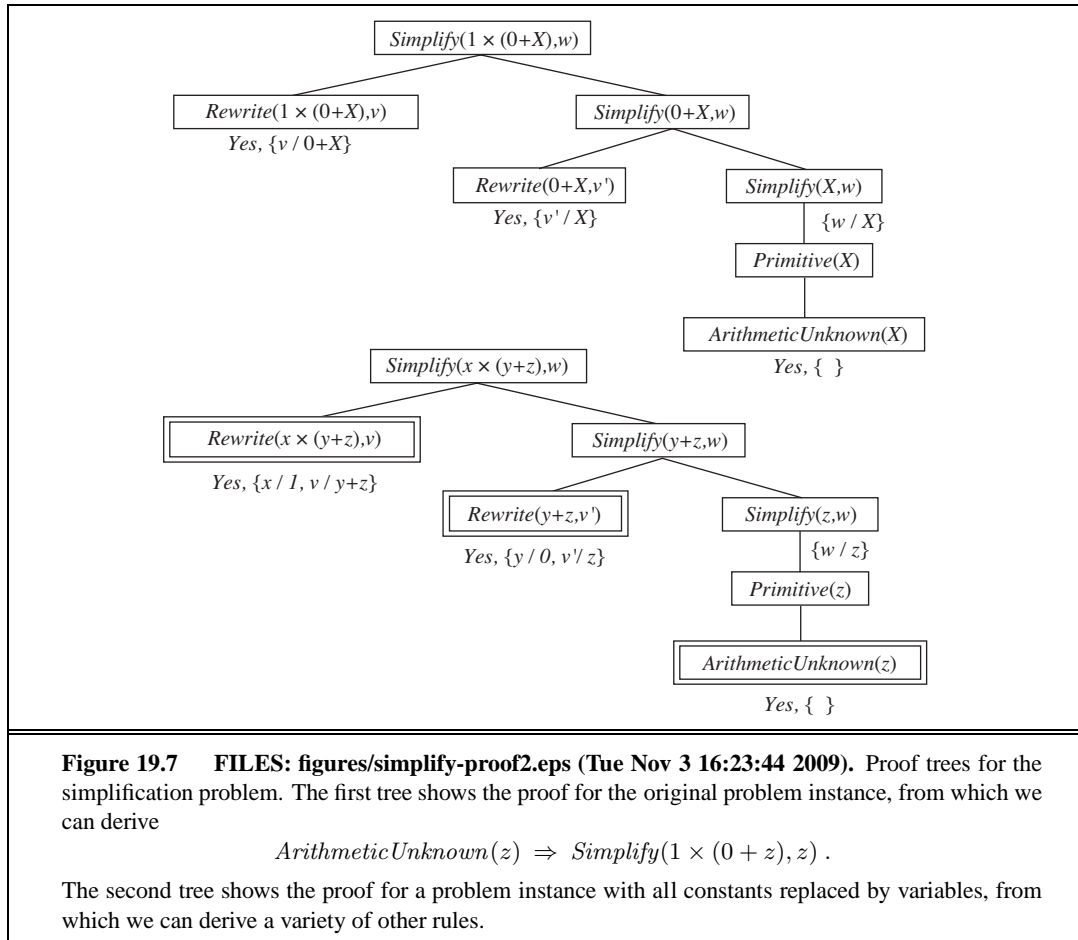
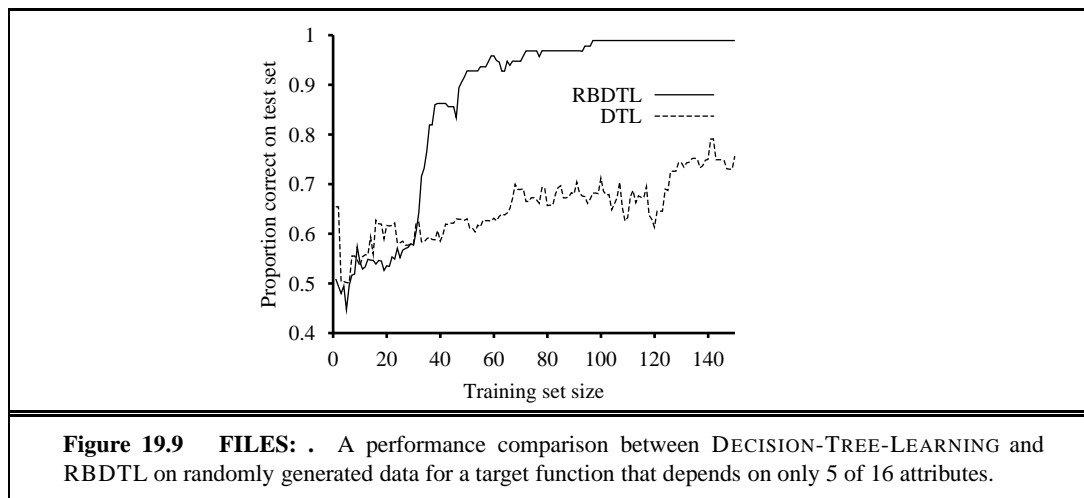


Figure 19.4 FILES: figures/version-space.eps (Tue Nov 3 16:24:02 2009). The version space contains all hypotheses consistent with the examples.









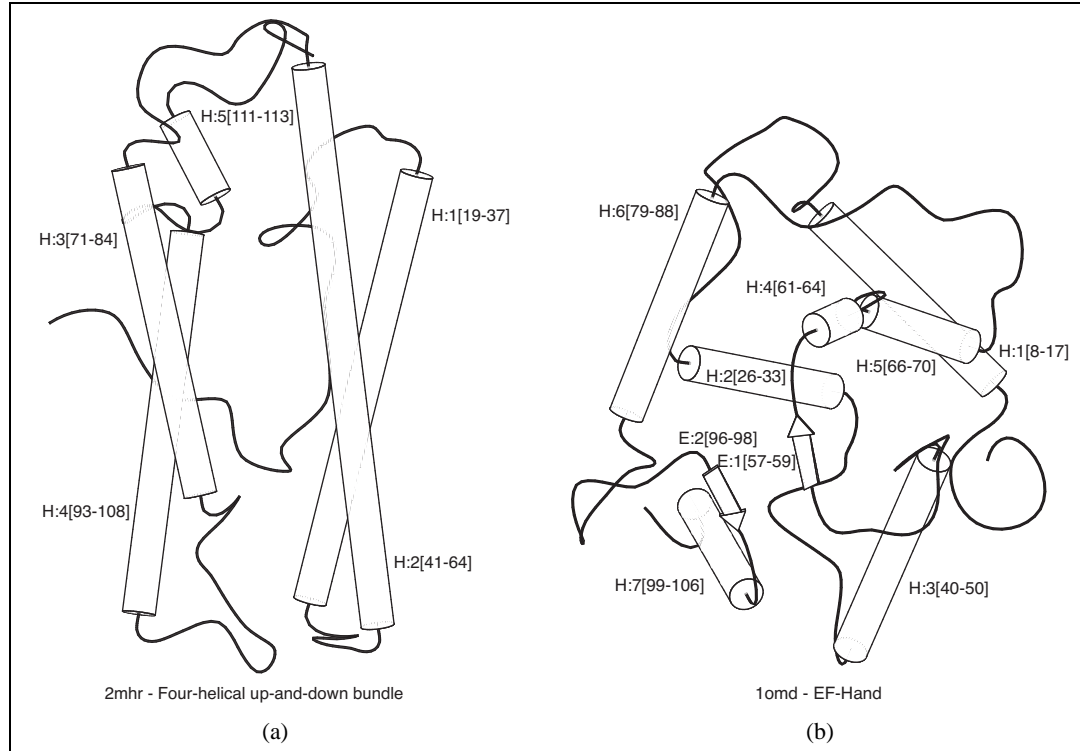
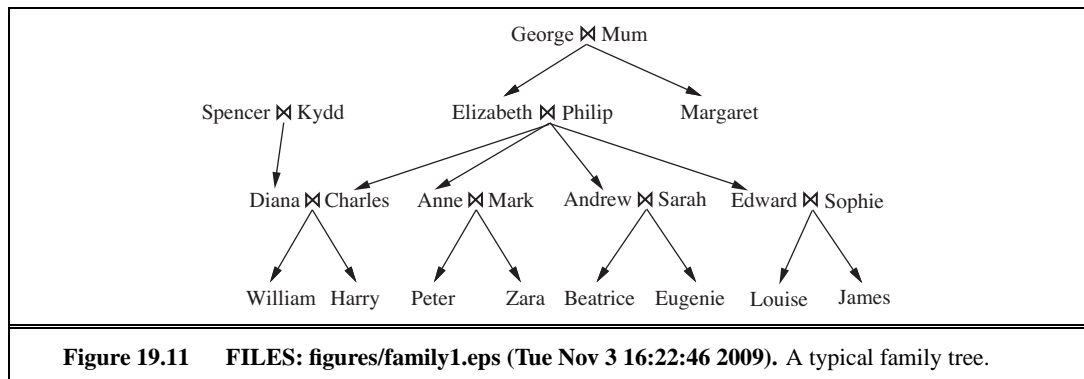
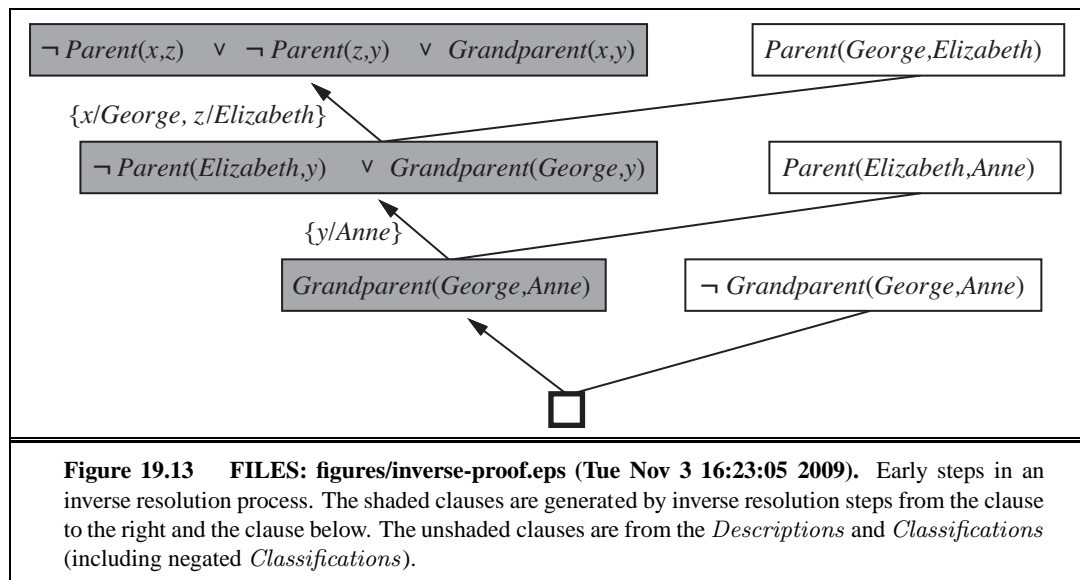



Figure 19.10 FILES: figures/pdb2mhr.eps (Tue Nov 3 16:23:15 2009) figures/pdb1omd.eps (Tue Nov 3 16:23:15 2009). (a) and (b) show positive and negative examples, respectively, of the “four-helical up-and-down bundle” concept in the domain of protein folding. Each example structure is coded into a logical expression of about 100 conjuncts such as $TotalLength(D2mhr, 118) \wedge NumberHelices(D2mhr, 6) \wedge \dots$. From these descriptions and from classifications such as $Fold(FOUR-HELICAL-UP-AND-DOWN-BUNDLE, D2mhr)$, the ILP system PROGOL (?) learned the following rule:

$$\begin{aligned}
 Fold(FOUR-HELICAL-UP-AND-DOWN-BUNDLE, p) \Leftarrow & \\
 & Helix(p, h_1) \wedge Length(h_1, HIGH) \wedge Position(p, h_1, n) \\
 & \wedge (1 \leq n \leq 3) \wedge Adjacent(p, h_1, h_2) \wedge Helix(p, h_2) .
 \end{aligned}$$


This kind of rule could not be learned, or even represented, by an attribute-based mechanism such as we saw in previous chapters. The rule can be translated into English as “Protein p has fold class “Four-helical up-and-down-bundle” if it contains a long helix h_1 at a secondary structure position between 1 and 3 and h_1 is next to a second helix.”

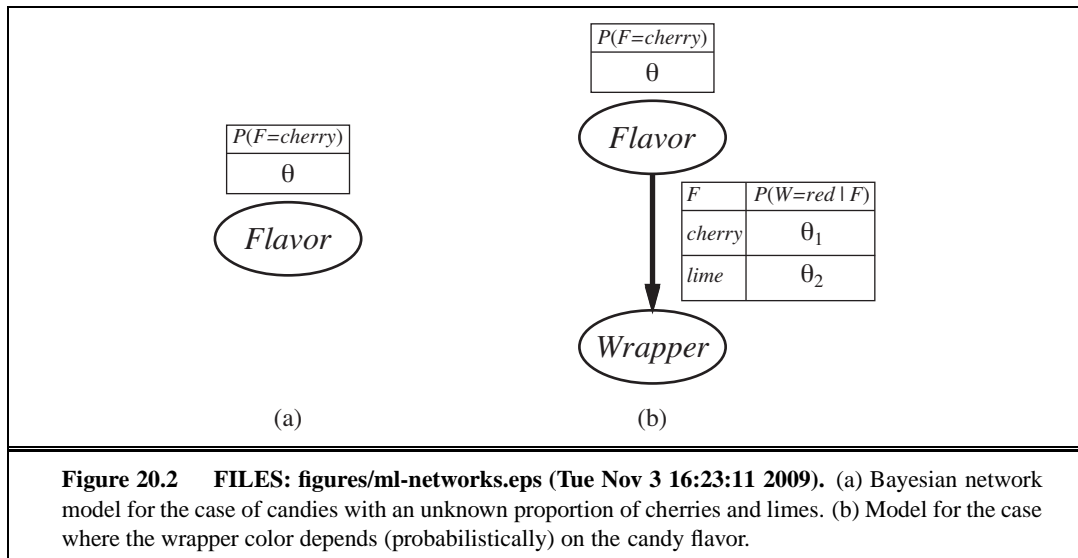






20 LEARNING PROBABILISTIC MODELS









Figure 20.3 FILES: . The learning curve for naive Bayes learning applied to the restaurant problem from Chapter 18; the learning curve for decision-tree learning is shown for comparison.






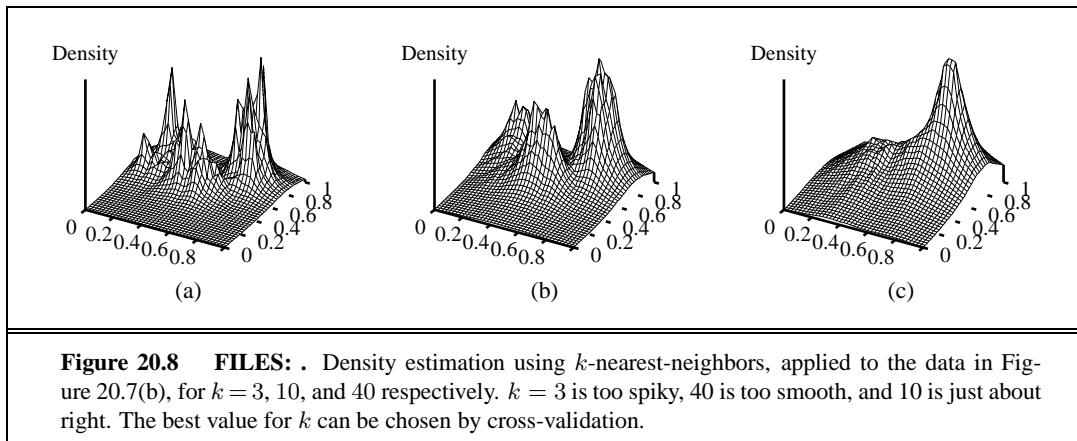



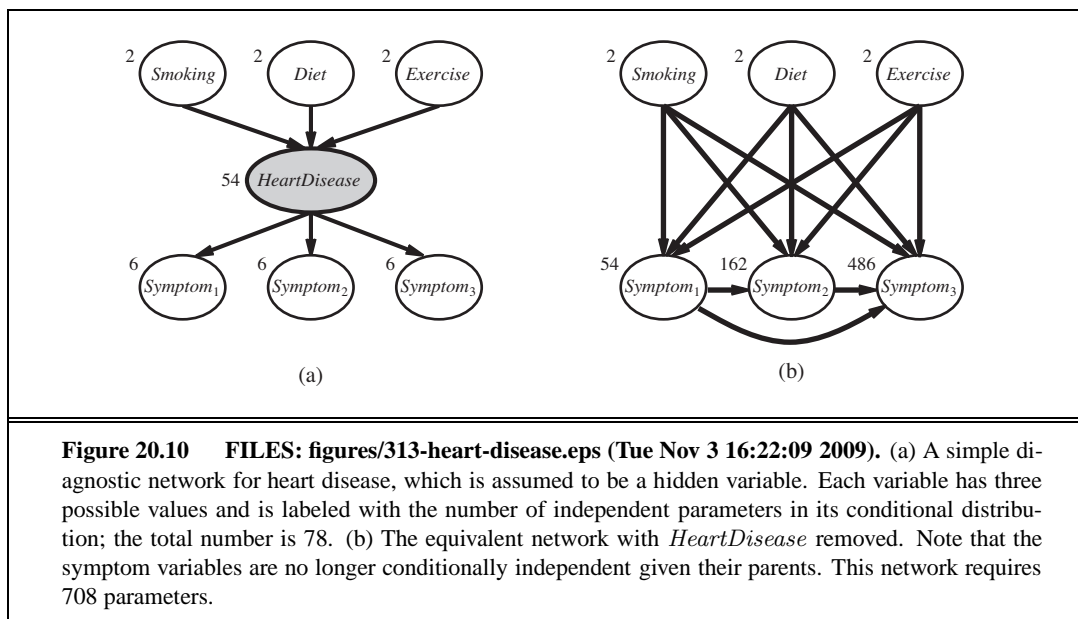
Figure 20.5 FILES: . Examples of the beta[a, b] distribution for different values of $[a, b]$.















Figure 20.11 FILES: . (a) A Gaussian mixture model with three components; the weights (left-to-right) are 0.2, 0.3, and 0.5. (b) 500 data points sampled from the model in (a). (c) The model reconstructed by EM from the data in (b).



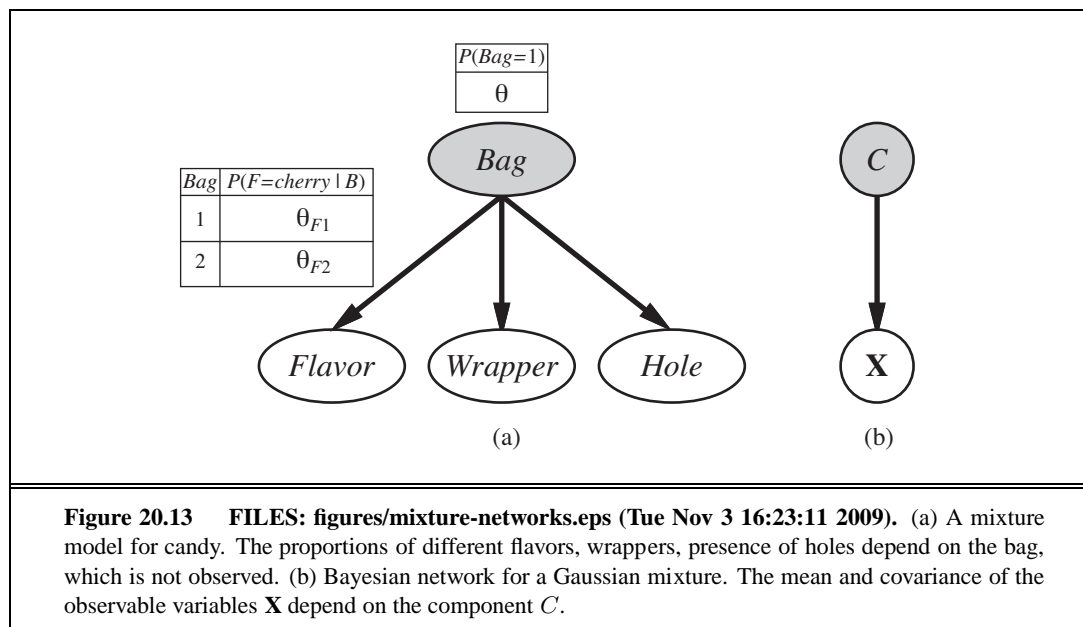
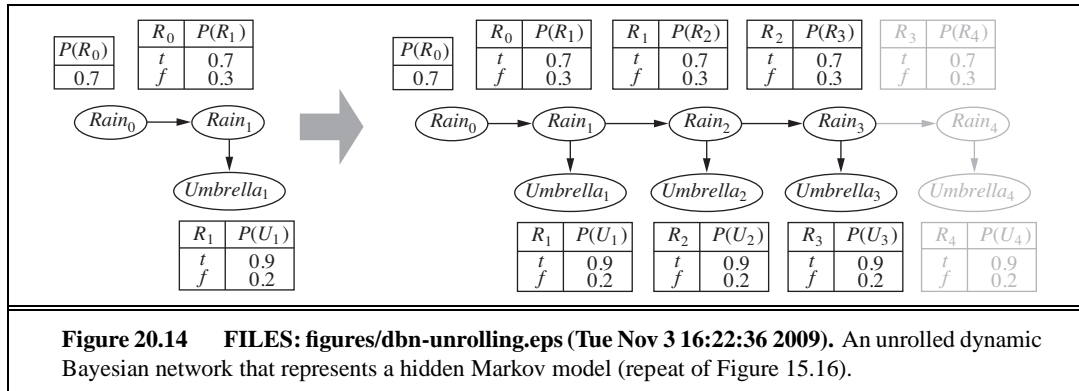
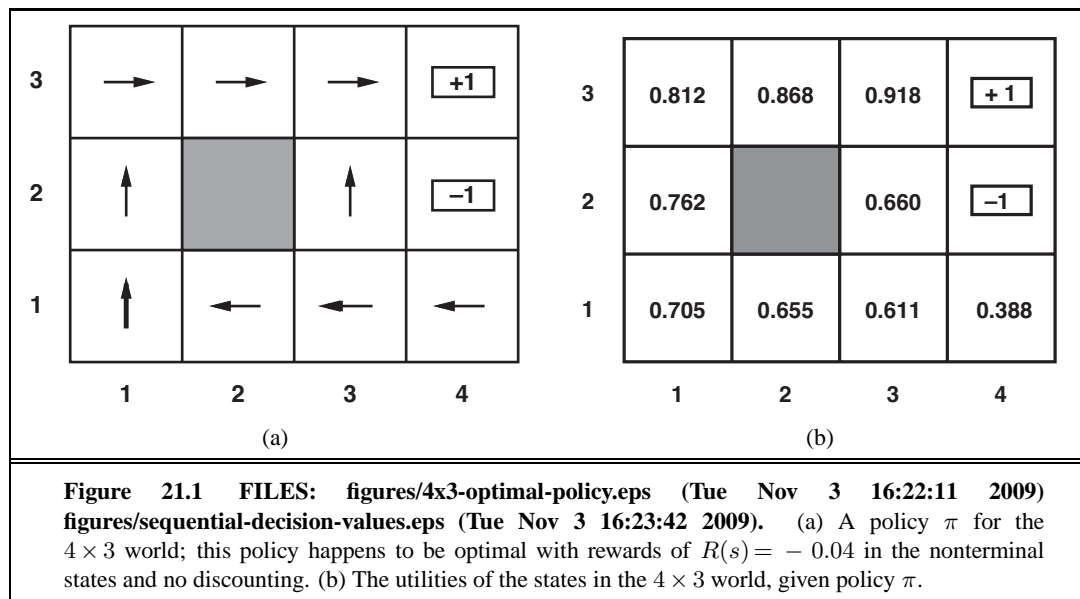


Figure 20.13 FILES: figures/mixture-networks.eps (Tue Nov 3 16:23:11 2009). (a) A mixture model for candy. The proportions of different flavors, wrappers, presence of holes depend on the bag, which is not observed. (b) Bayesian network for a Gaussian mixture. The mean and covariance of the observable variables \mathbf{X} depend on the component C .



21 REINFORCEMENT LEARNING



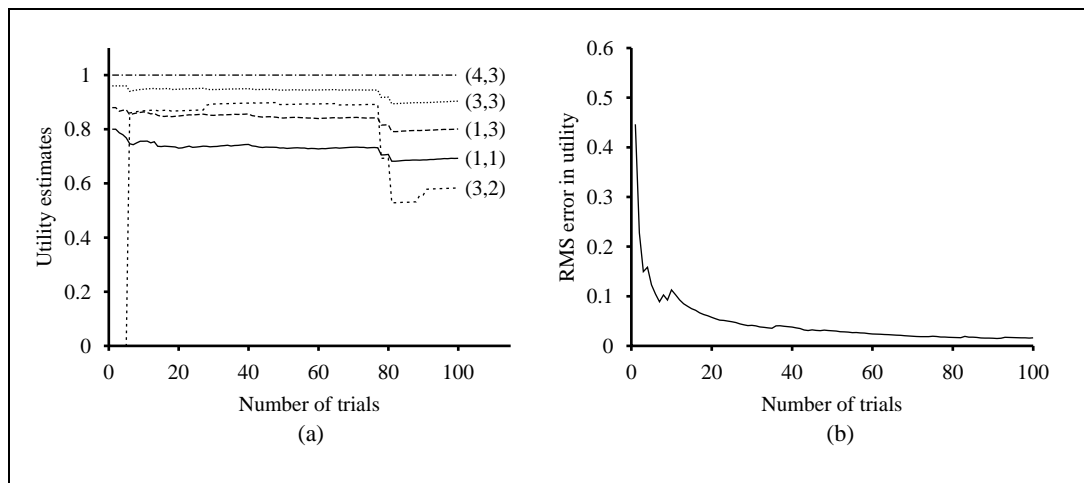

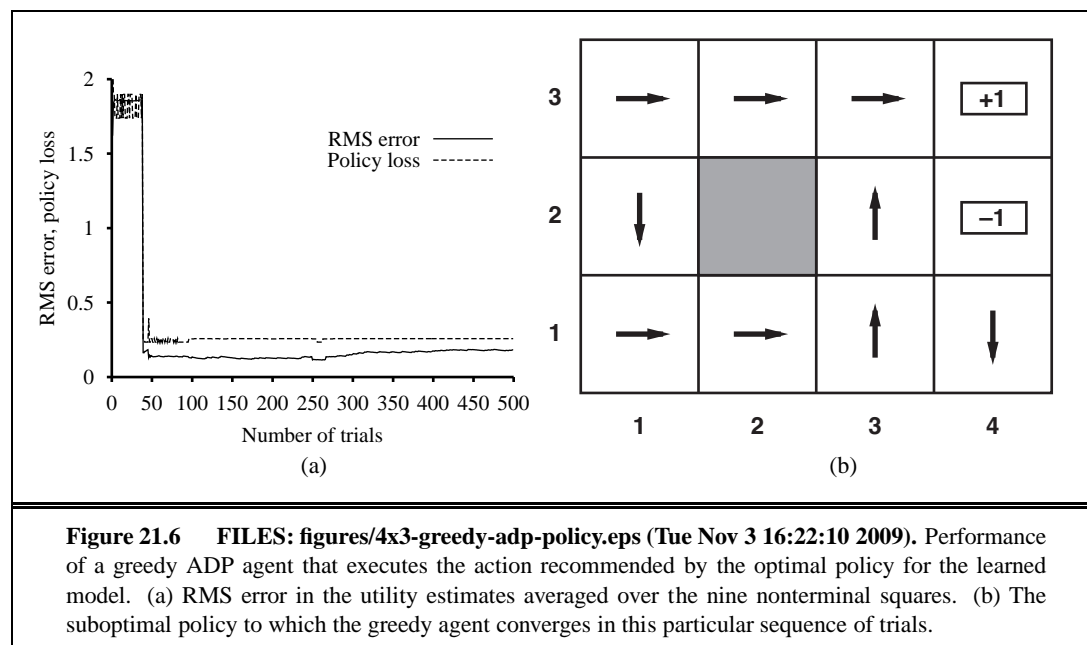



Figure 21.3 FILES: . The passive ADP learning curves for the 4×3 world, given the optimal policy shown in Figure 21.1. (a) The utility estimates for a selected subset of states, as a function of the number of trials. Notice the large changes occurring around the 78th trial—this is the first time that the agent falls into the -1 terminal state at (4,2). (b) The root-mean-square error (see Appendix A) in the estimate for $U(1,1)$, averaged over 20 runs of 100 trials each.








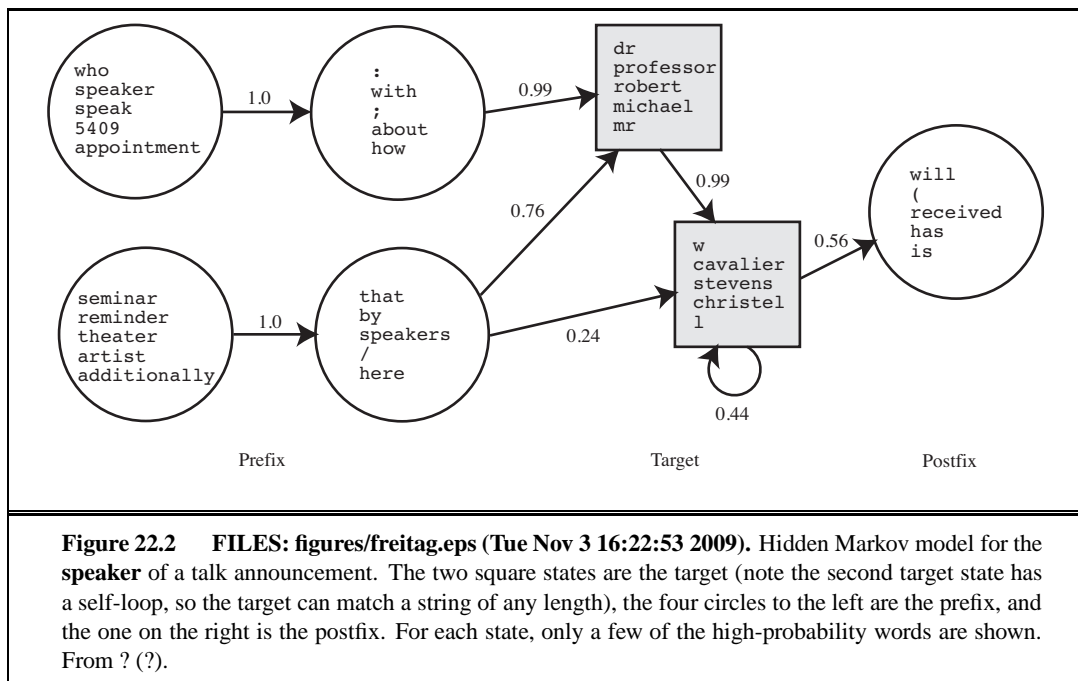




Figure 21.10 FILES: figures/heliComposite.eps (Tue Nov 3 16:23:02 2009). Superimposed time-lapse images of an autonomous helicopter performing a very difficult “nose-in circle” maneuver. The helicopter is under the control of a policy developed by the PEGASUS policy-search algorithm. A simulator model was developed by observing the effects of various control manipulations on the real helicopter; then the algorithm was run on the simulator model overnight. A variety of controllers were developed for different maneuvers. In all cases, performance far exceeded that of an expert human pilot using remote control. (Image courtesy of Andrew Ng.)

22 NATURAL LANGUAGE PROCESSING



23 NATURAL LANGUAGE FOR COMMUNICATION


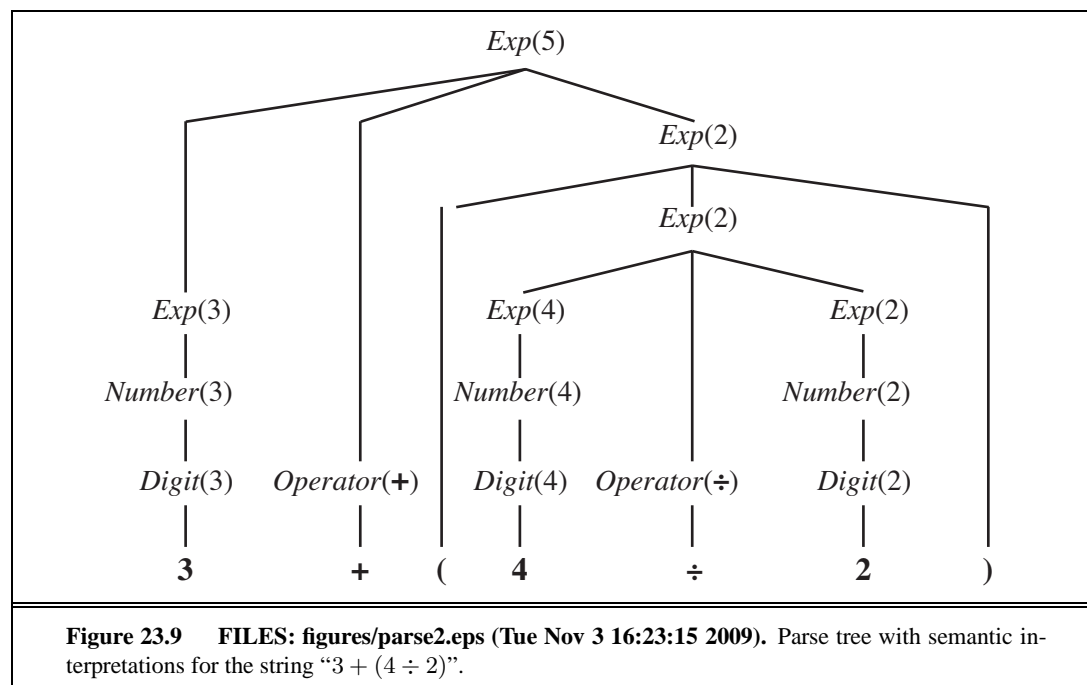
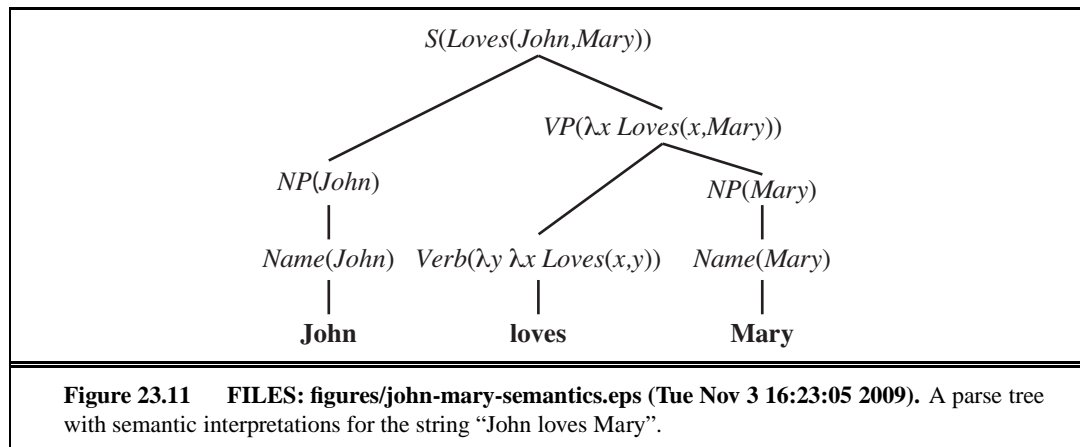


Figure 23.3 FILES: figures/parse-pcfg.eps (Tue Nov 3 16:23:15 2009). Parse tree for the sentence “Every wumpus smells” according to the grammar \mathcal{E}_0 . Each interior node of the tree is labeled with its probability. The probability of the tree as a whole is $0.9 \times 0.25 \times 0.05 \times 0.15 \times 0.40 \times 0.10 = 0.0000675$. Since this tree is the only parse of the sentence, that number is also the probability of the sentence. The tree can also be written in linear form as $[S [NP [Article \textbf{every}] [Noun \textbf{wumpus}]] [VP [Verb \textbf{smells}]]]$.












Figure 23.12 FILES: figures/mt-interlingua.eps (Tue Nov 3 16:23:11 2009). The Vauquois triangle: schematic diagram of the choices for a machine translation system (?). We start with English text at the top. An interlingua-based system follows the solid lines, parsing English first into a syntactic form, then into a semantic representation and an interlingua representation, and then through generation to a semantic, syntactic, and lexical form in French. A transfer-based system uses the dashed lines as a shortcut. Different systems make the transfer at different points; some make it at multiple points.










24 PERCEPTION






Figure 24.2 FILES: figures/newpinhole.eps (Tue Nov 3 16:23:14 2009). Each light-sensitive element in the image plane at the back of a pinhole camera receives light from a the small range of directions that passes through the pinhole. If the pinhole is small enough, the result is a focused image at the back of the pinhole. The process of projection means that large, distant objects look the same as smaller, nearby objects. Note that the image is projected upside down.





Figure 24.3 FILES: figures/lens-eye.eps (Tue Nov 3 16:23:06 2009). Lenses collect the light leaving a scene point in a range of directions, and steer it all to arrive at a single point on the image plane. Focusing works for points lying close to a focal plane in space; other points will not be focused properly. In cameras, elements of the lens system move to change the focal plane, whereas in the eye, the shape of the lens is changed by specialized muscles.








Figure 24.5 FILES: figures/lambert.eps (Tue Nov 3 13:41:38 2009). Two surface patches are illuminated by a distant point source, whose rays are shown as gray arrowheads. Patch A is tilted away from the source (θ is close to 90°) and collects less energy, because it cuts fewer light rays per unit surface area. Patch B, facing the source (θ is close to 0°), collects more energy.






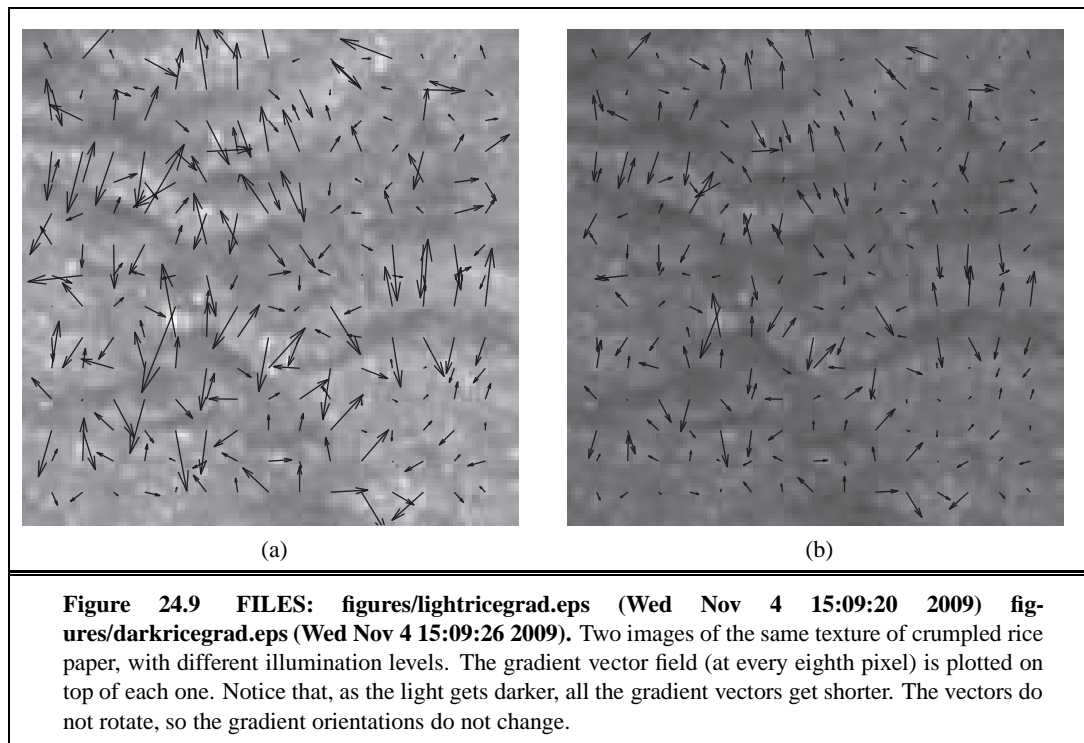
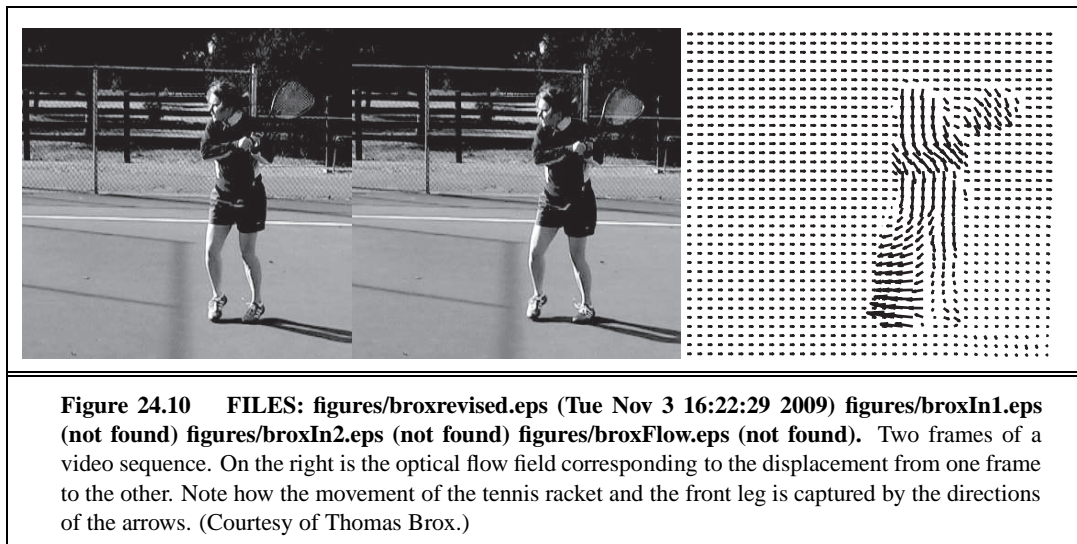


Figure 24.8 FILES: figures/edgewderiv.eps (Tue Nov 3 16:22:41 2009). Top: Intensity profile $I(x)$ along a one-dimensional section across an edge at $x=50$. Middle: The derivative of intensity, $I'(x)$. Large values of this function correspond to edges, but the function is noisy. Bottom: The derivative of a smoothed version of the intensity, $(I * G_\sigma)'$, which can be computed in one step as the convolution $I * G'_\sigma$. The noisy candidate edge at $x=75$ has disappeared.








Figure 24.11 FILES: figures/101087.eps (Tue Nov 3 16:22:07 2009) figures/101087-ucm-th0.eps (not found) figures/101087-seg-th0.eps (not found) figures/101087-seg-th0-5.eps (not found). (a) Original image. (b) Boundary contours, where the higher the P_b value, the darker the contour. (c) Segmentation into regions, corresponding to a fine partition of the image. Regions are rendered in their mean colors. (d) Segmentation into regions, corresponding to a coarser partition of the image, resulting in fewer regions. (Courtesy of Pablo Arbelaez, Michael Maire, Charles Fowlkes, and Jitendra Malik)

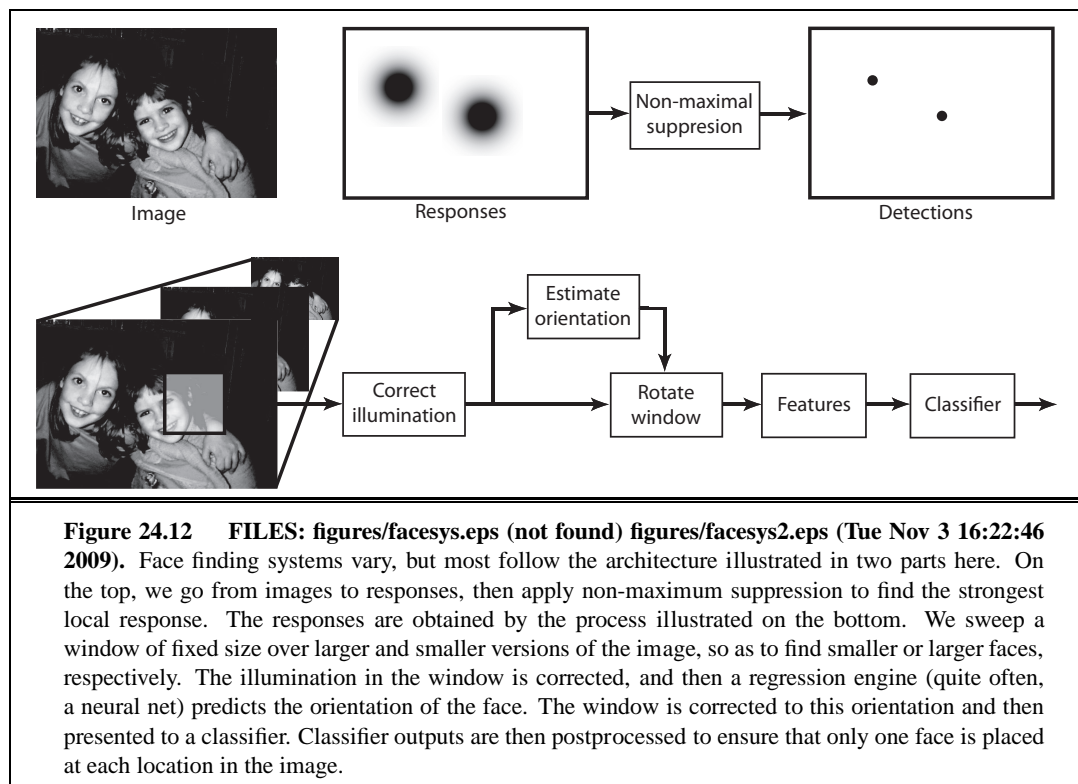
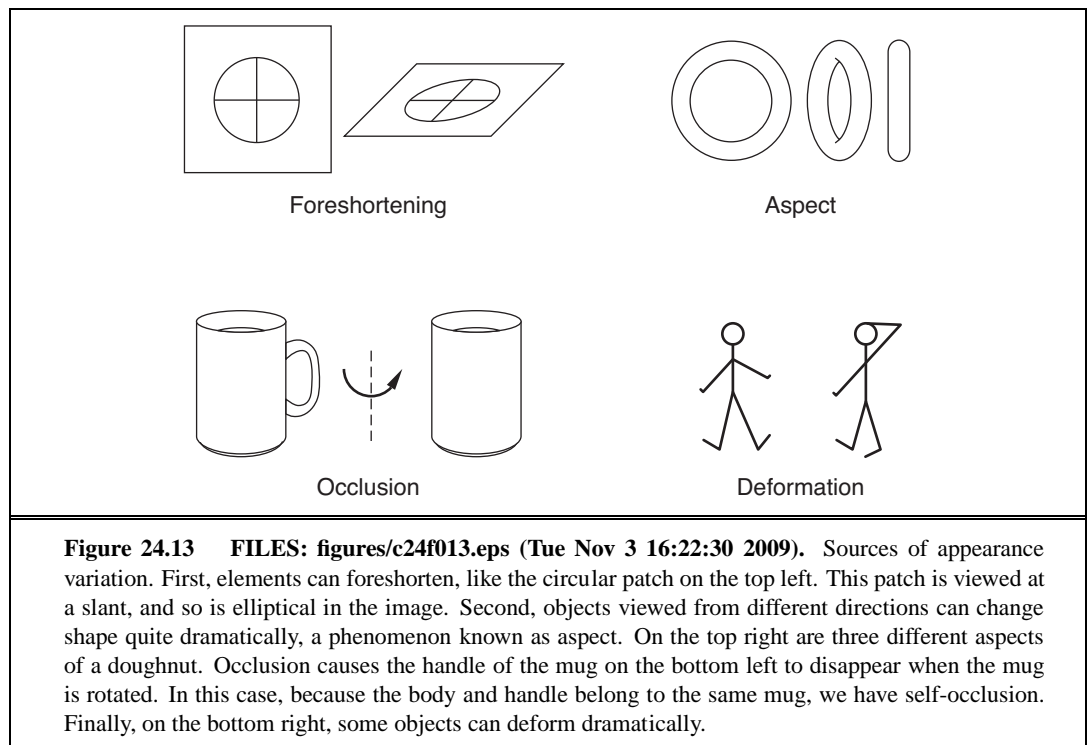

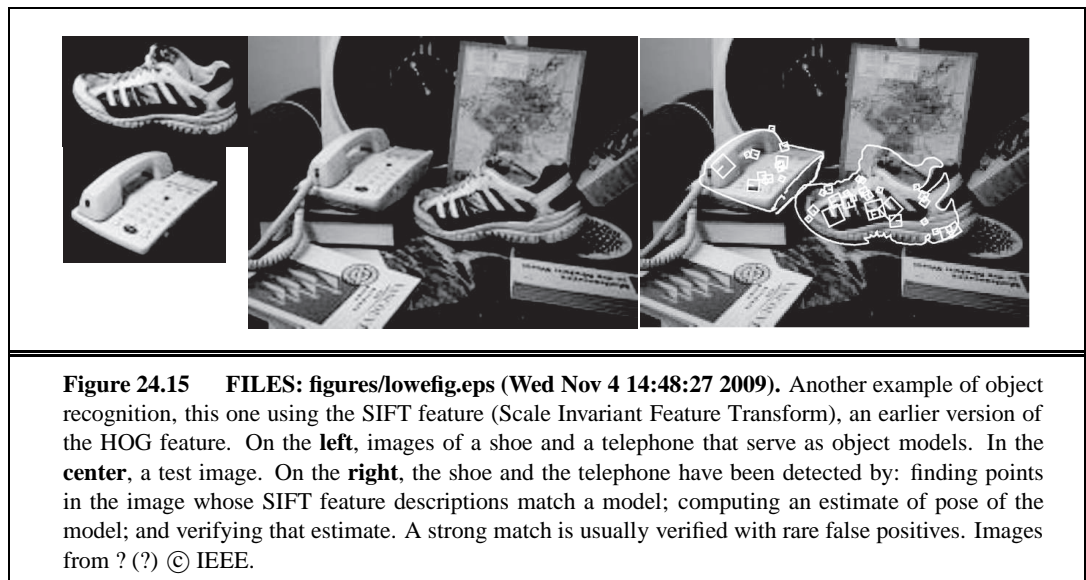


Figure 24.12 FILES: figures/facesys.eps (not found) figures/facesys2.eps (Tue Nov 3 16:22:46 2009). Face finding systems vary, but most follow the architecture illustrated in two parts here. On the top, we go from images to responses, then apply non-maximum suppression to find the strongest local response. The responses are obtained by the process illustrated on the bottom. We sweep a window of fixed size over larger and smaller versions of the image, so as to find smaller or larger faces, respectively. The illumination in the window is corrected, and then a regression engine (quite often, a neural net) predicts the orientation of the face. The window is corrected to this orientation and then presented to a classifier. Classifier outputs are then postprocessed to ensure that only one face is placed at each location in the image.










Figure 24.16 FILES: figures/c24f017.eps (Tue Nov 3 16:22:31 2009) figures/stereo-1.eps (not found) figures/stereo-2.eps (not found). Translating a camera parallel to the image plane causes image features to move in the camera plane. The disparity in positions that results is a cue to depth. If we superimpose left and right image, as in (b), we see the disparity.


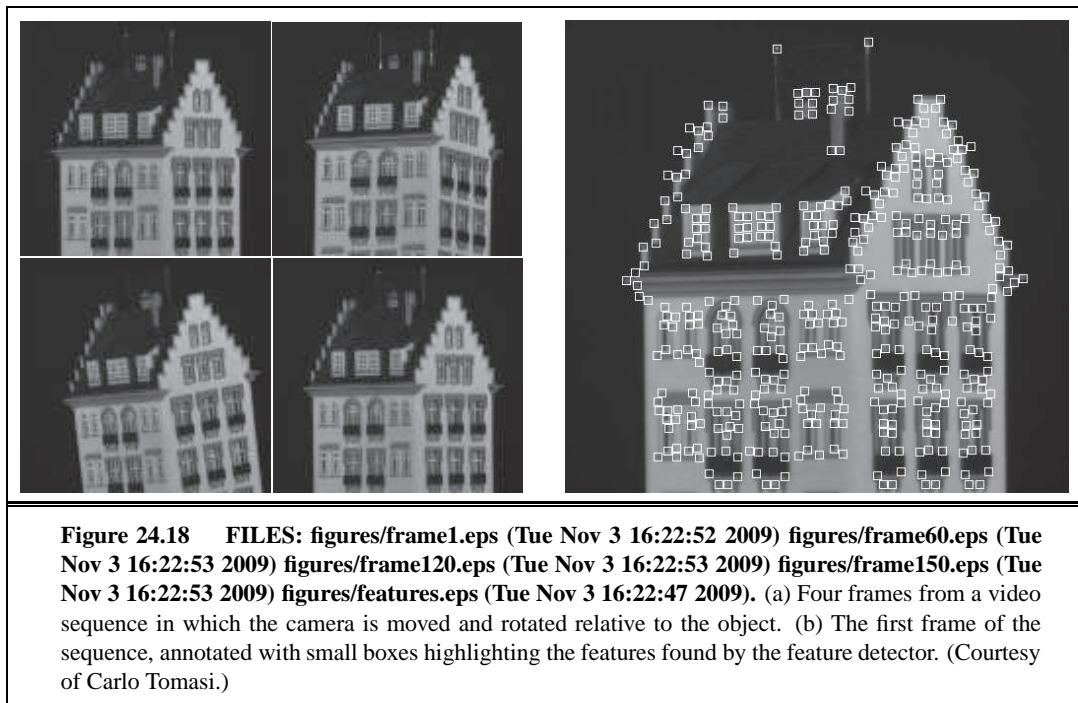



Figure 24.17 FILES: figures/stereopsis.eps (Tue Nov 3 16:23:49 2009). The relation between disparity and depth in stereopsis. The centers of projection of the two eyes are b apart, and the optical axes intersect at the fixation point P_0 . The point P in the scene projects to points P_L and P_R in the two eyes. In angular terms, the disparity between these is $\delta\theta$. See text.









Figure 24.20 FILES: figures/camp-test.eps (Tue Nov 3 16:22:32 2009) figures/chem-test.eps (Tue Nov 3 16:22:32 2009). (a) A textured scene. Assuming that the real texture is uniform allows recovery of the surface orientation. The computed surface orientation is indicated by overlaying a black circle and pointer, transformed as if the circle were painted on the surface at that point. (b) Recovery of shape from texture for a curved surface (white circle and pointer this time). Images courtesy of Jitendra Malik and Ruth Rosenholtz (?).



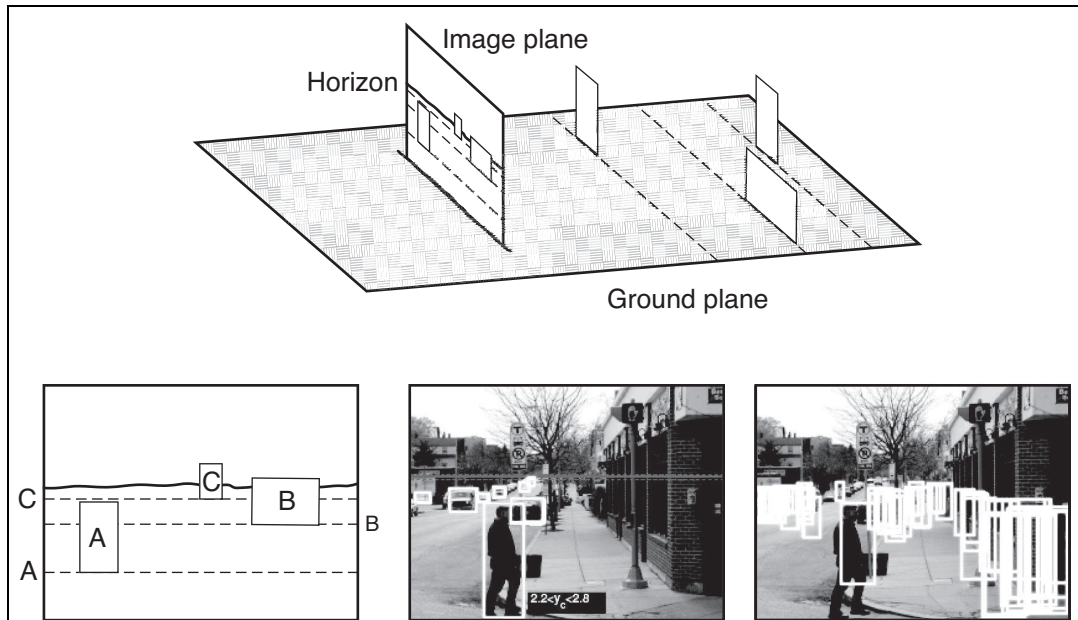


Figure 24.22 FILES: figures/c24f022-a.eps (Tue Nov 3 16:22:31 2009). In an image of people standing on a ground plane, the people whose feet are closer to the horizon in the image must be farther away (top drawing). This means they must look smaller in the image (left lower drawing). This means that the size and location of real pedestrians in an image depend upon one another and on the location of the horizon. To exploit this, we need to identify the ground plane, which is done using shape-from-texture methods. From this information, and from some likely pedestrians, we can recover a horizon as shown in the center image. On the right, acceptable pedestrian boxes given this geometric context. Notice that pedestrians who are higher in the scene must be smaller. If they are not, then they are false positives. Images from ? (?) © IEEE.


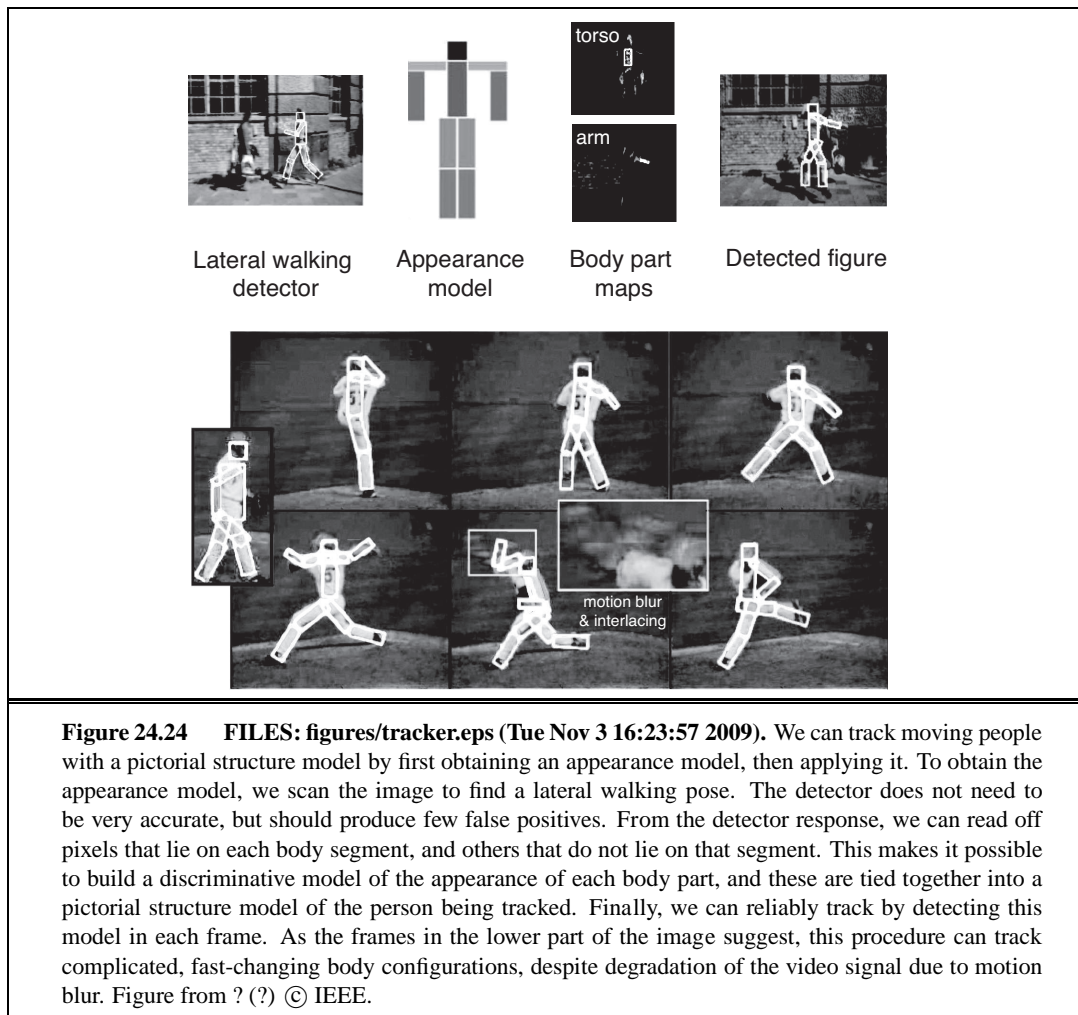


Figure 24.23 FILES: figures/armslegs.eps (Tue Nov 3 16:22:24 2009). A pictorial structure model evaluates a match between a set of image rectangles and a cardboard person (shown on the left) by scoring the similarity in appearance between body segments and image segments and the spatial relations between the image segments. Generally, a match is better if the image segments have about the right appearance and are in about the right place with respect to one another. The appearance model uses average colors for hair, head, torso, and upper and lower arms and legs. The relevant relations are shown as arrows. On the right, the best match for a particular image, obtained using dynamic programming. The match is a fair estimate of the configuration of the body. Figure from ? (?) © IEEE.










Figure 24.25 FILES: figures/drinking-2.eps (Tue Nov 3 16:22:38 2009). Some complex human actions produce consistent patterns of appearance and motion. For example, drinking involves movements of the hand in front of the face. The first three images are correct detections of drinking; the fourth is a false-positive (the cook is looking into the coffee pot, but not drinking from it). Figure from ? (?) © IEEE.






25 ROBOTICS













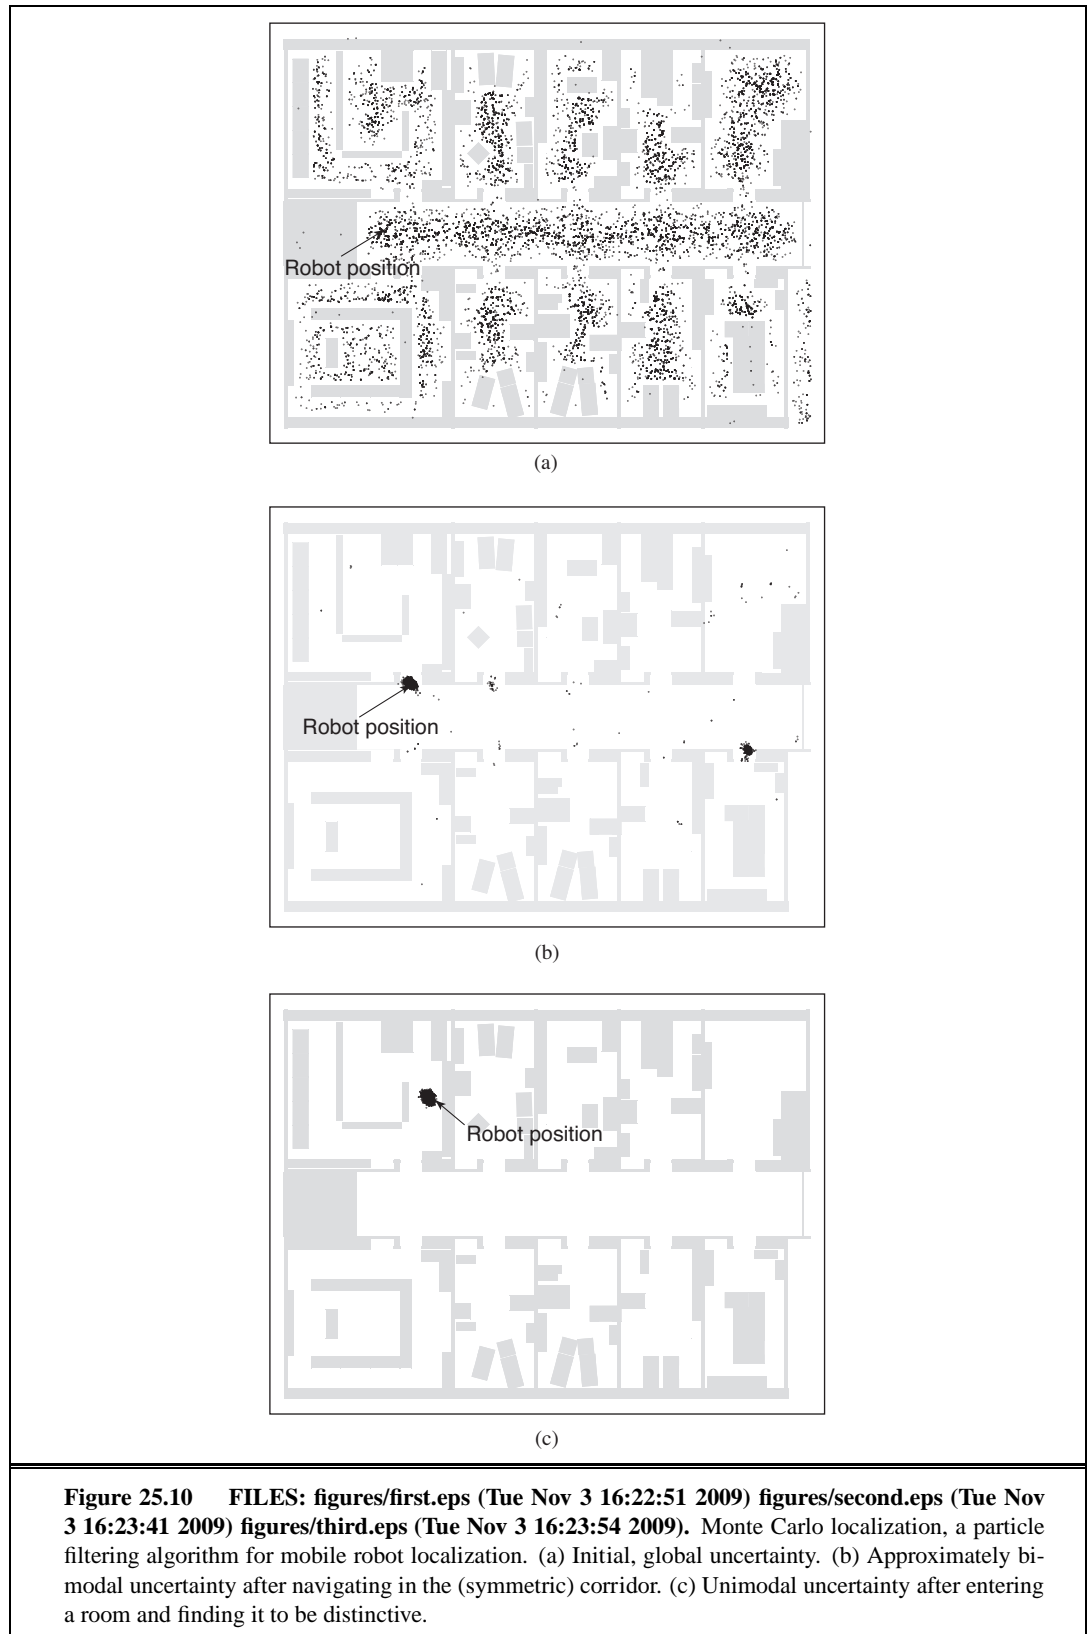
Figure 25.4 FILES: `figures/stanford-arm.eps` (Tue Nov 3 16:23:46 2009) `figures/car-like.eps` (Tue Nov 3 16:22:32 2009). (a) The Stanford Manipulator, an early robot arm with five revolute joints (R) and one prismatic joint (P), for a total of six degrees of freedom. (b) Motion of a nonholonomic four-wheeled vehicle with front-wheel steering.












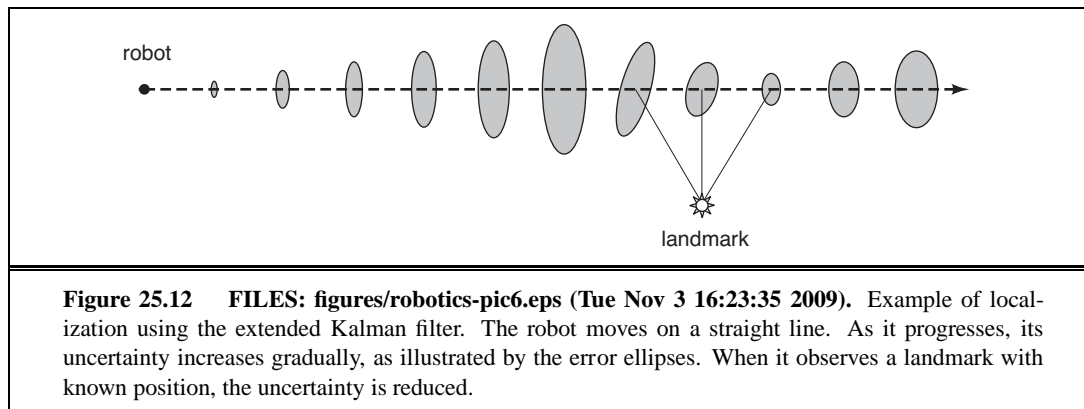



Figure 25.11 FILES: figures/robotics-pic3.eps (Tue Nov 3 16:23:34 2009) figures/robotics-pic4.eps (Tue Nov 3 16:23:34 2009). One-dimensional illustration of a linearized motion model: (a) The function f , and the projection of a mean μ_t and a covariance interval (based on Σ_t) into time $t+1$. (b) The linearized version is the tangent of f at μ_t . The projection of the mean μ_t is correct. However, the projected covariance $\tilde{\Sigma}_{t+1}$ differs from Σ_{t+1} .






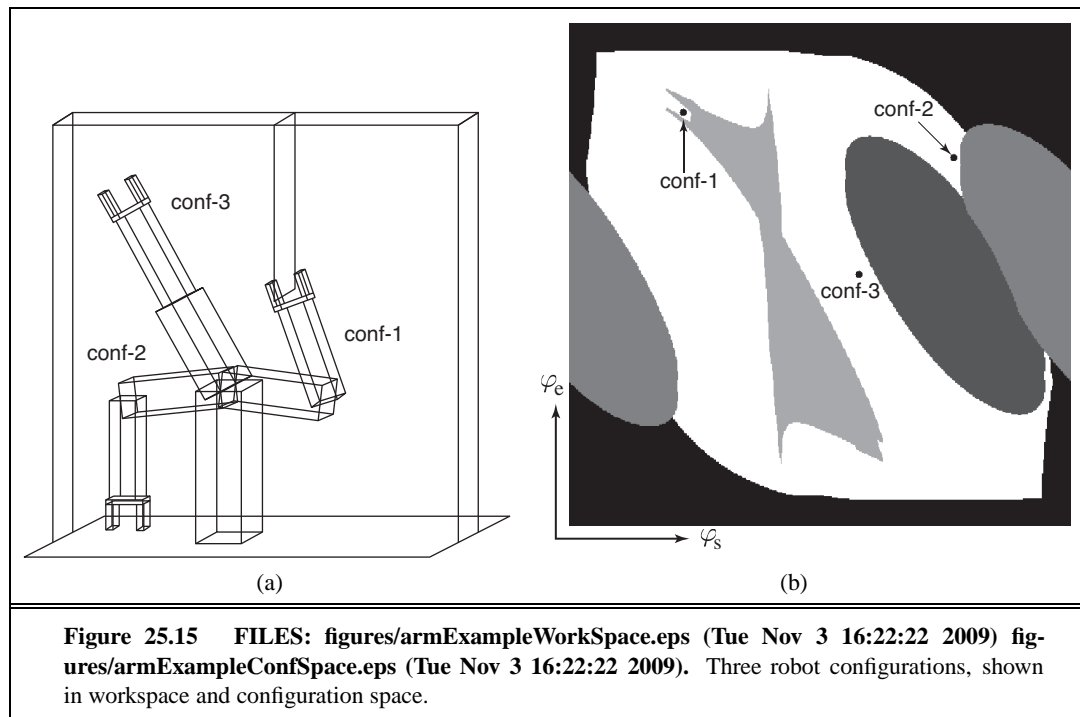




Figure 25.14 FILES: figures/armPlain.eps (Tue Nov 3 16:22:22 2009) figures/armPlainConfSpace.eps (Tue Nov 3 16:22:23 2009). (a) Workspace representation of a robot arm with 2 DOFs. The workspace is a box with a flat obstacle hanging from the ceiling. (b) Configuration space of the same robot. Only white regions in the space are configurations that are free of collisions. The dot in this diagram corresponds to the configuration of the robot shown on the left.










Figure 25.18 FILES: `figures/armVoronoi.eps` (Tue Nov 3 16:22:23 2009) `figures/armRoadmap.eps` (Tue Nov 3 16:22:23 2009). (a) The Voronoi graph is the set of points equidistant to two or more obstacles in configuration space. (b) A probabilistic roadmap, composed of 400 randomly chosen points in free space.

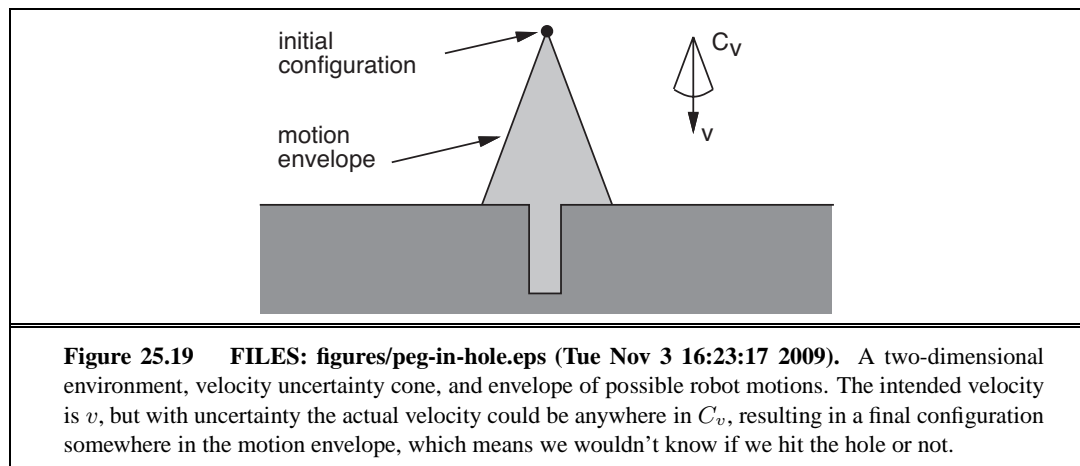
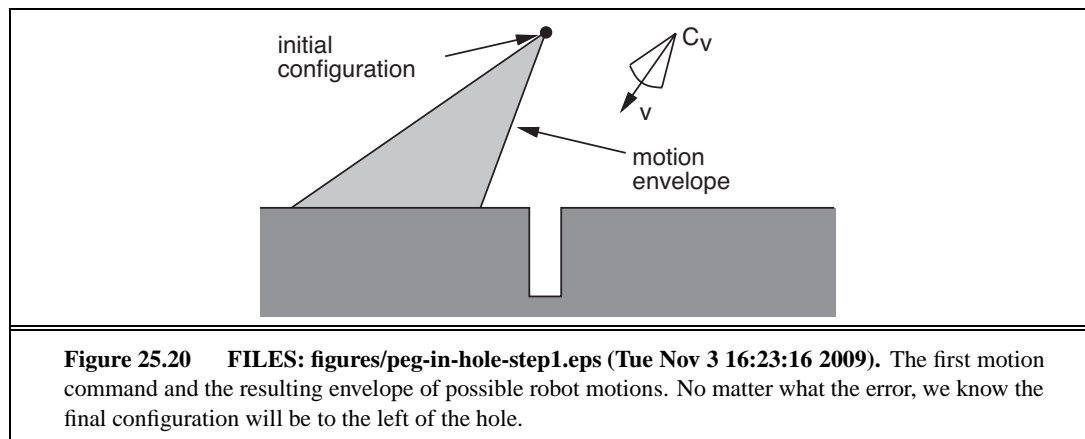

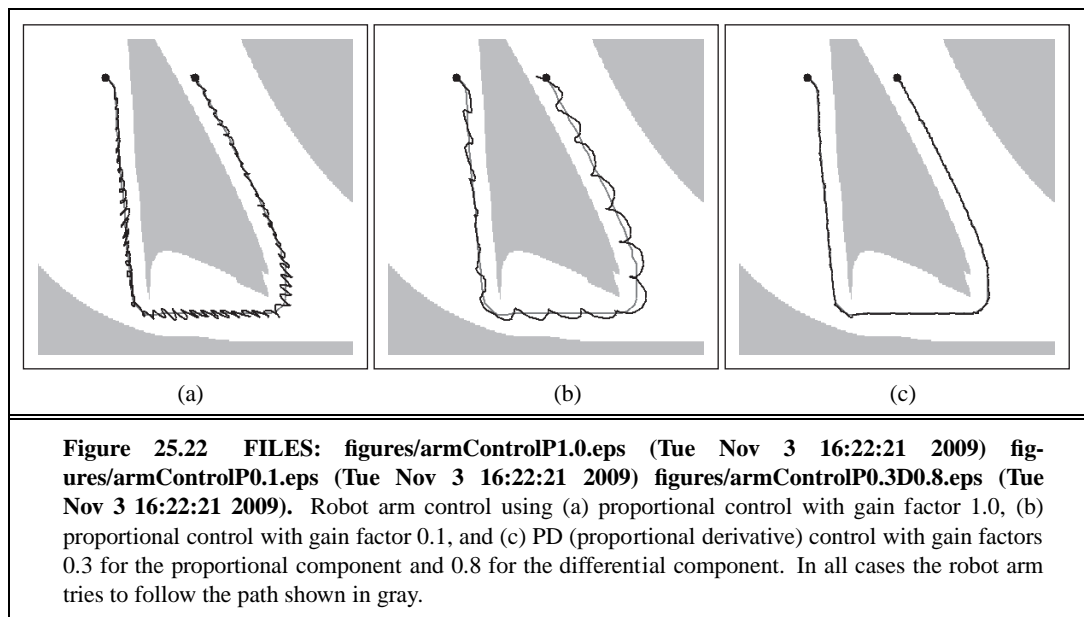
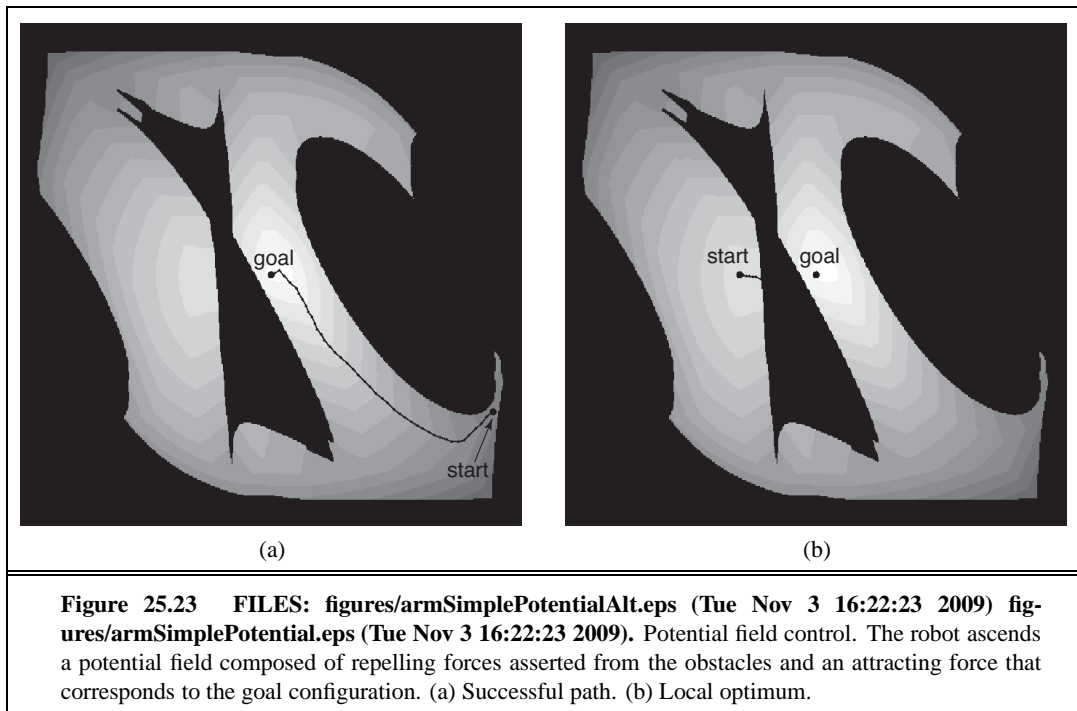


Figure 25.19 FILES: figures/peg-in-hole.eps (Tue Nov 3 16:23:17 2009). A two-dimensional environment, velocity uncertainty cone, and envelope of possible robot motions. The intended velocity is v , but with uncertainty the actual velocity could be anywhere in C_v , resulting in a final configuration somewhere in the motion envelope, which means we wouldn't know if we hit the hole or not.














Figure 25.25 FILES: figures/flip-mosaic.eps (not found). Multiple exposures of an RC helicopter executing a flip based on a policy learned with reinforcement learning. Images courtesy of Andrew Ng, Stanford University.


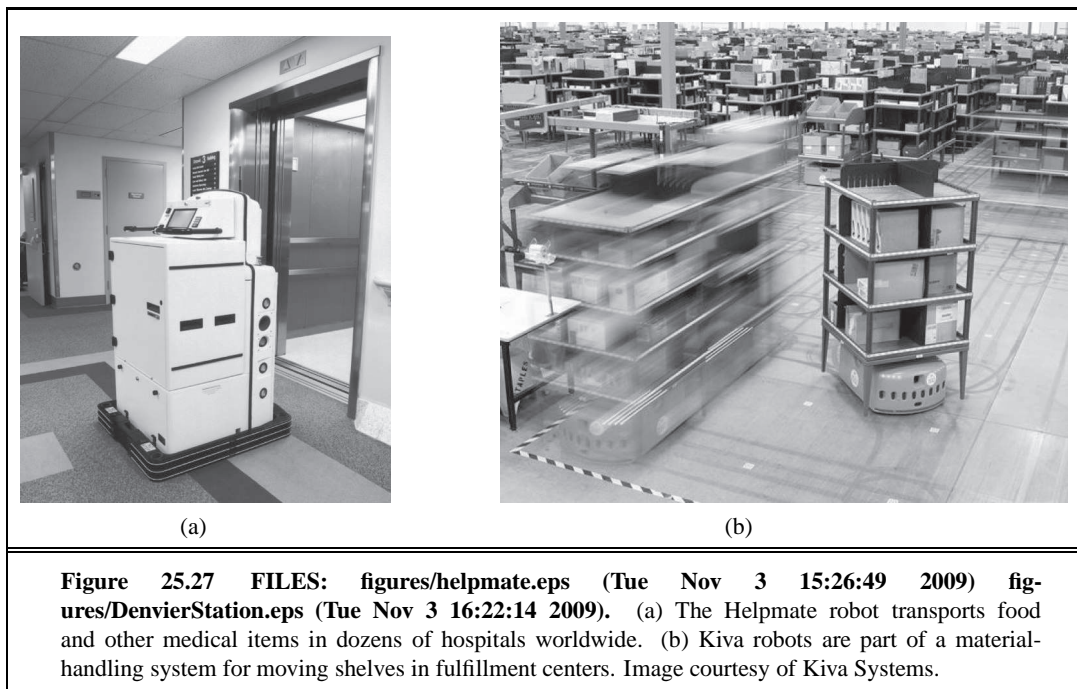



Figure 25.26 FILES: figures/stanley-processes.eps (Wed Nov 4 11:17:26 2009). Software architecture of a robot car. This software implements a data pipeline, in which all modules process data simultaneously.





(a)



(b)

Figure 25.28 FILES: **figures/race12.eps** (Wed Nov 4 15:18:14 2009) **figures/munich-ORsmall.eps** (not found). (a) Robotic car BOSS, which won the DARPA Urban Challenge. Courtesy of Carnegie Mellon University. (b) Surgical robots in the operating room. Image courtesy of da Vinci Surgical Systems.




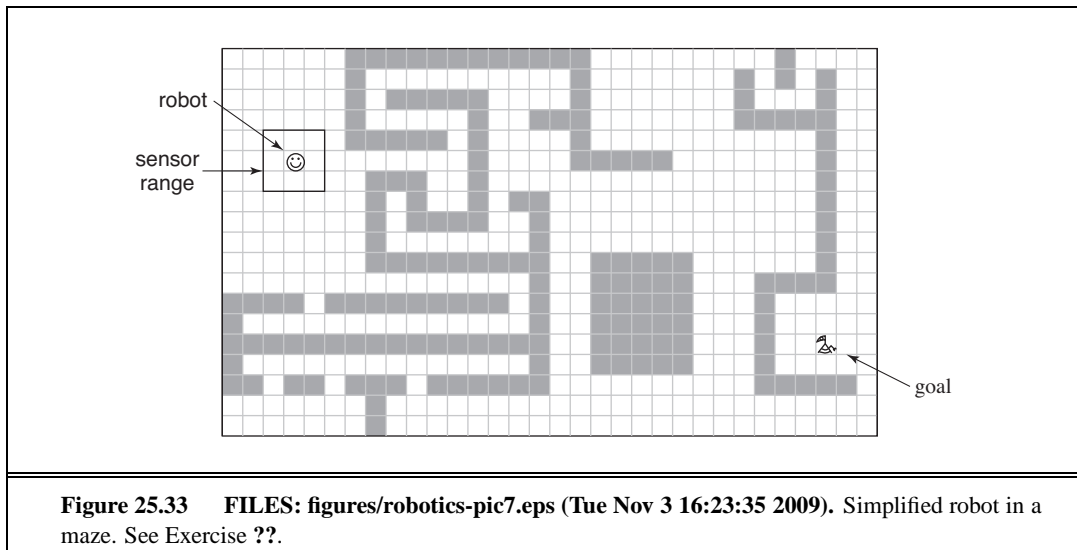




Figure 25.30 FILES: figures/roomba1.eps (Wed Nov 4 15:22:23 2009) figures/icra-cordless-phone-gray3.eps (not found). (a) Roomba, the world's best-selling mobile robot, vacuums floors. Image courtesy of iRobot, © 2009. (b) Robotic hand modeled after human hand. Image courtesy of University of Washington and Carnegie Mellon University.







26 PHILOSOPHICAL
FOUNDATIONS

27 AI: THE PRESENT AND
FUTURE

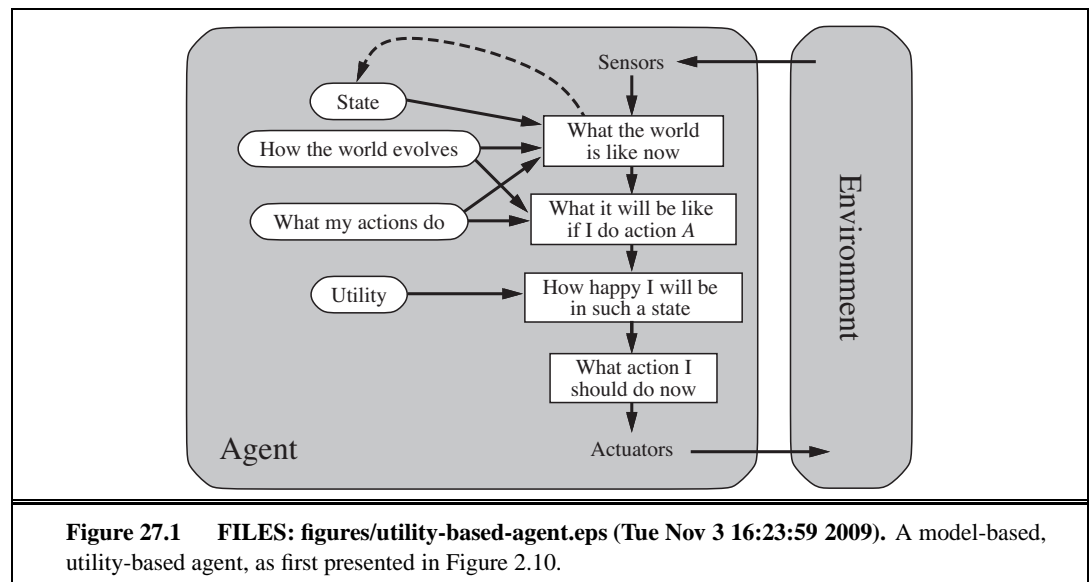



Figure 27.1 FILES: figures/utility-based-agent.eps (Tue Nov 3 16:23:59 2009). A model-based, utility-based agent, as first presented in Figure 2.10.



28 MATHEMATICAL BACKGROUND

29

NOTES ON LANGUAGES AND ALGORITHMS