Chapter 10

Artificial Intelligence

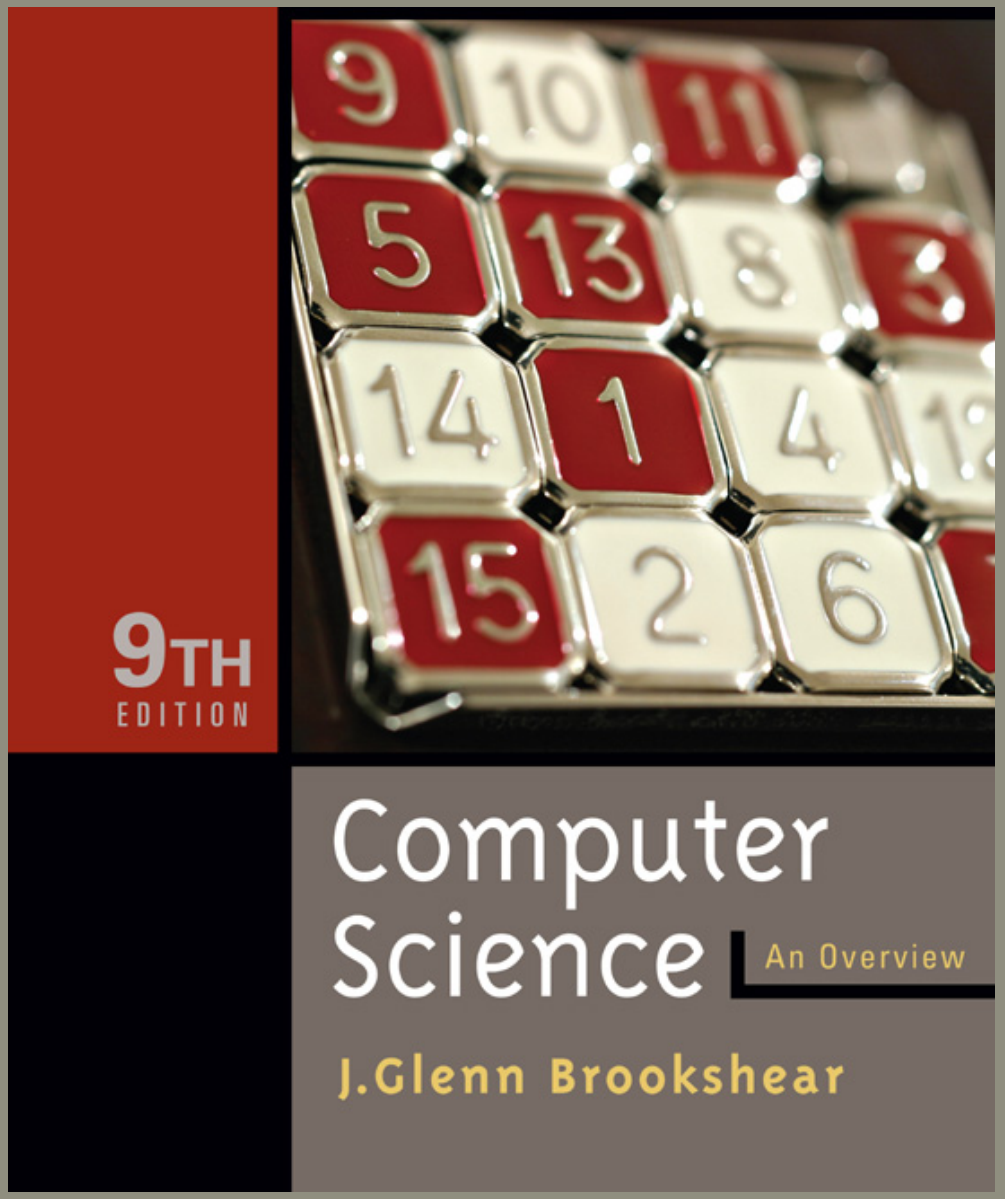

© 2007 Pearson Addison-Wesley.
All rights reserved

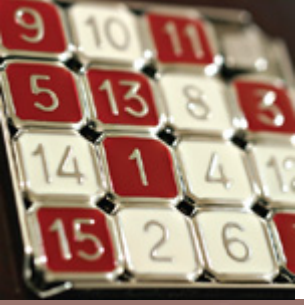

Chapter 10: Artificial Intelligence

- 10.1 Intelligence and Machines
- 10.2 Perception
- 10.3 Reasoning
- 10.4 Additional Areas of Research
- 10.5 Artificial Neural Networks
- 10.6 Robotics
- 10.7 Considering the Consequences

© 2007 Pearson Addison-Wesley. All rights reserved

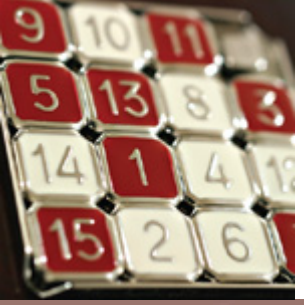

Intelligent Agents

- **Agent:** A “device” that responds to stimuli from its environment
  - Sensors
  - Actuators
- The goal of artificial intelligence is to build agents that behave intelligently

© 2007 Pearson Addison-Wesley. All rights reserved

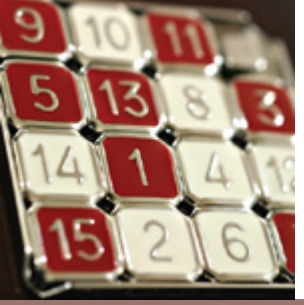

Levels of Intelligent Behavior

- Reflex: actions are predetermined responses to the input data
- Intelligent response: actions affected by knowledge of the environment
- Goal seeking
- Learning

© 2007 Pearson Addison-Wesley. All rights reserved

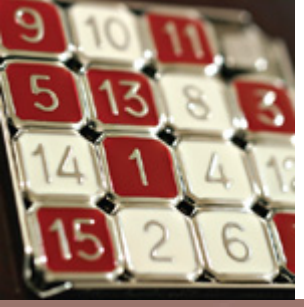

Figure 10.1 The eight-puzzle in its solved configuration

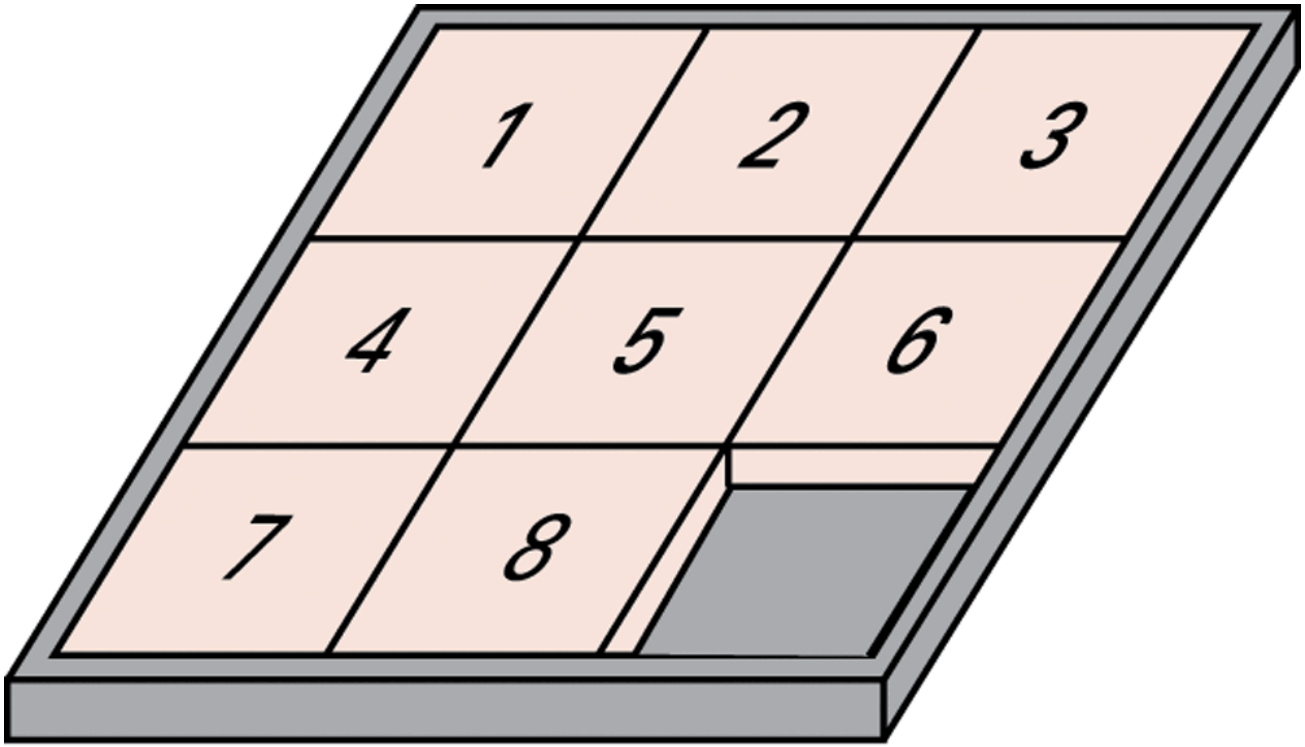

© 2007 Pearson Addison-Wesley. All rights reserved

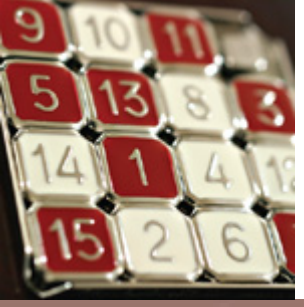

Figure 10.2 Our puzzle-solving machine

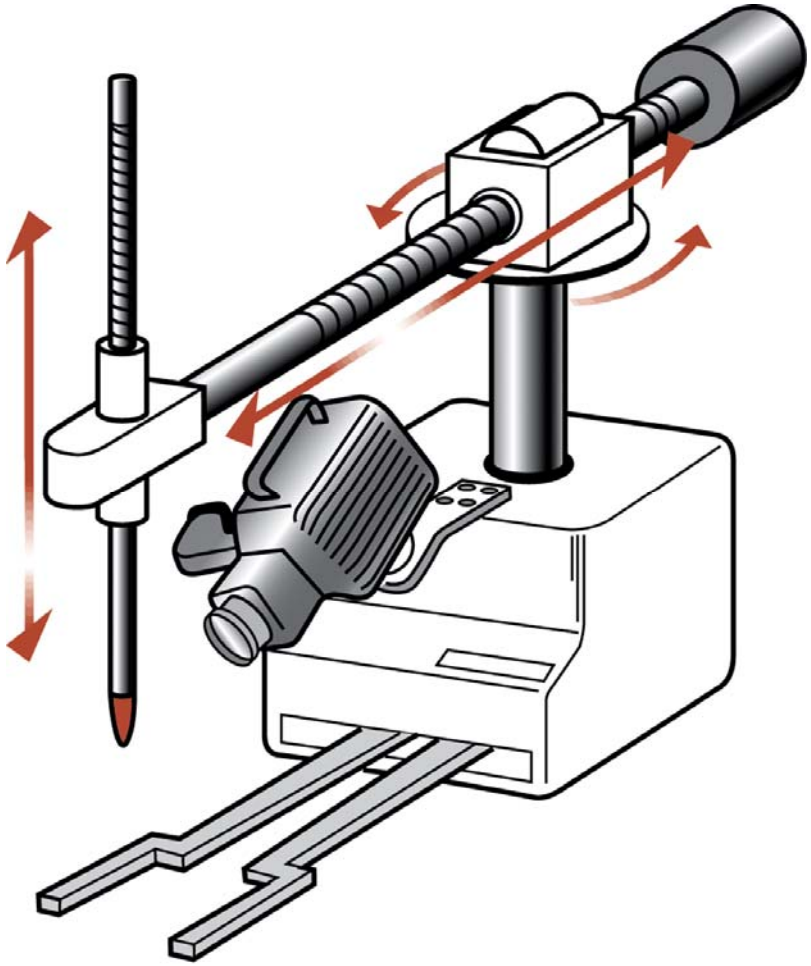

© 2007 Pearson Addison-Wesley. All rights reserved

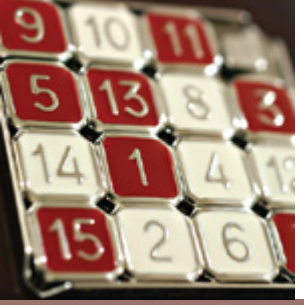

Approaches to Research in Artificial Intelligence

- Performance oriented: Researcher tries to maximize the performance of the agents.
- Simulation oriented: Researcher tries to understand how the agents produce responses.

© 2007 Pearson Addison-Wesley. All rights reserved

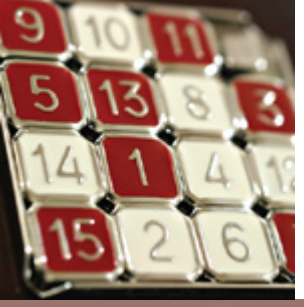

Turing Test

- Proposed by Alan Turing in 1950
- Benchmark for progress in artificial intelligence
- Test setup: Human interrogator communicates with test subject by typewriter.
- Test: Can the human interrogator distinguish whether the test subject is human or machine?

© 2007 Pearson Addison-Wesley. All rights reserved

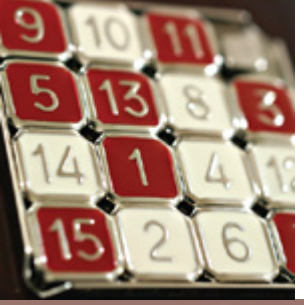

Techniques for Understanding Images

- Template matching
- Image processing
  - edge enhancement
  - region finding
  - smoothing
- Image analysis

© 2007 Pearson Addison-Wesley. All rights reserved

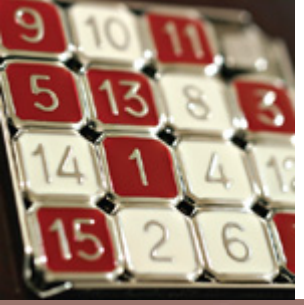

Language Processing

- Syntactic(语法) Analysis
- Semantic (语义) Analysis
- Contextual (上下文) Analysis

（P438 examples）

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.3 A semantic net

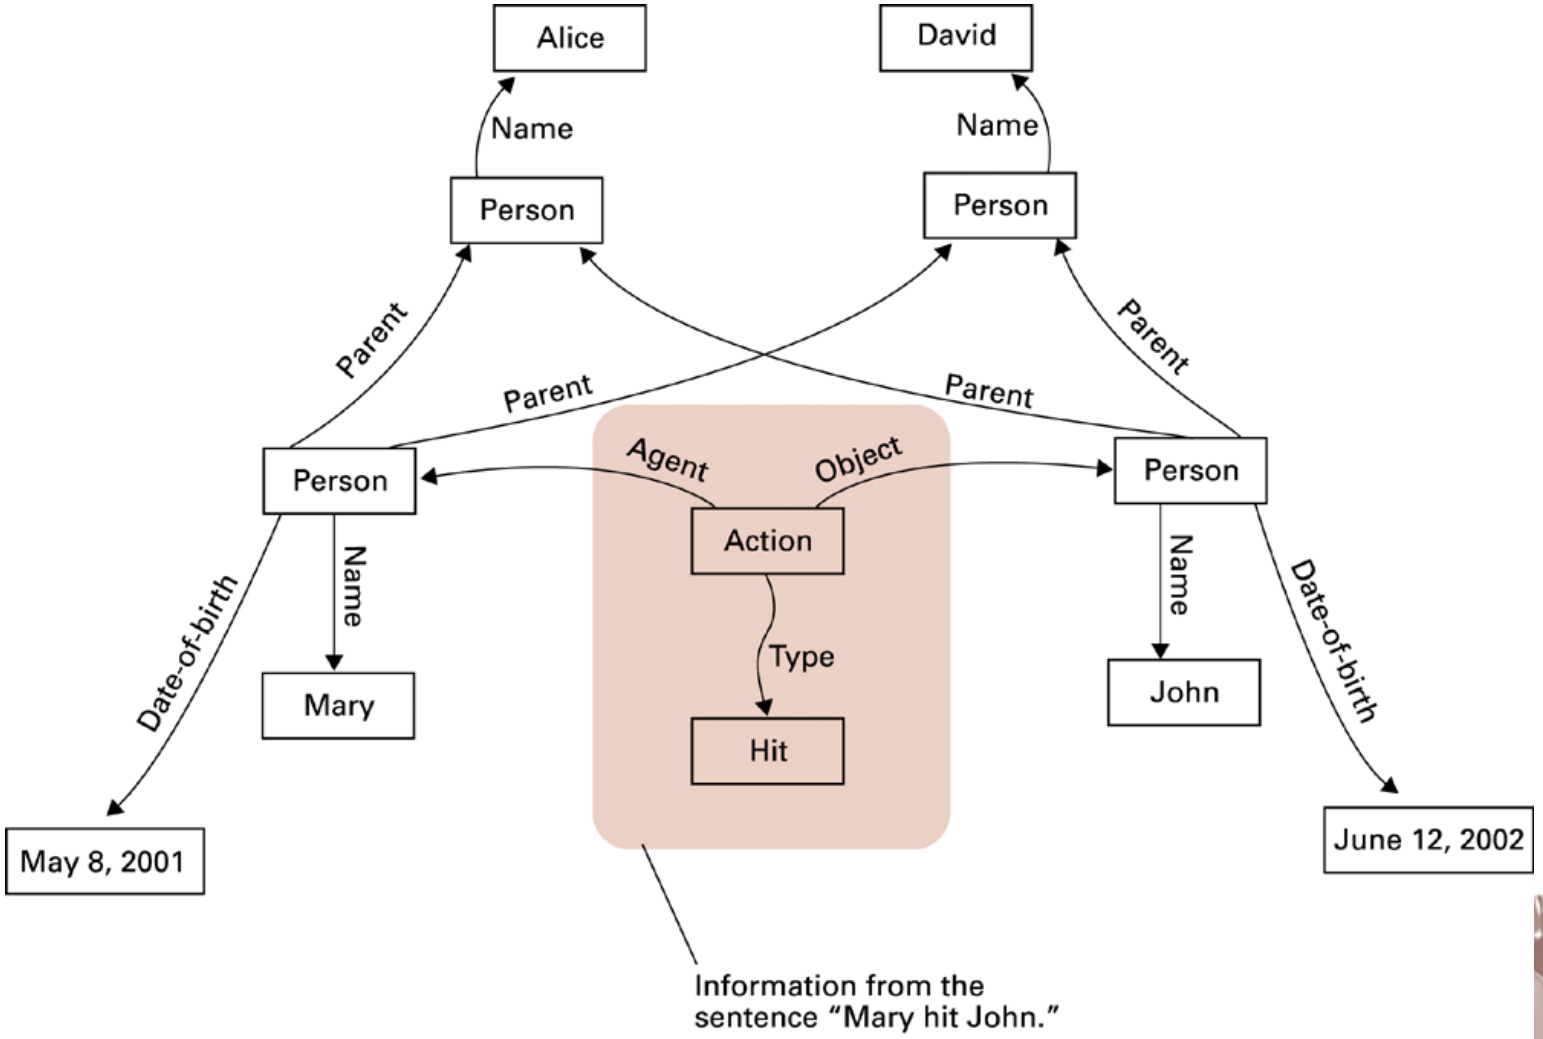

© 2007 Pearson Addison-Wesley. All rights reserved

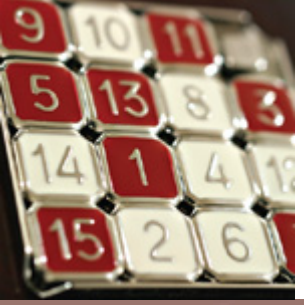

Components of a Production Systems

1. Collection of states
   - Start (or initial) state
   - Goal state (or states)
2. Collection of productions: rules or moves
   - Each production may have preconditions
3. ***Control system:*** decides which production to apply next

The task

© 2007 Pearson Addison-Wesley. All rights reserved

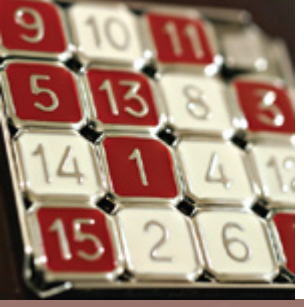

Reasoning by Searching

An important concept in the development of a control system

- **State Graph:** All states and productions
- **Search Tree:** A record of state transitions explored while searching for a goal state
  - Breadth-first search
  - Depth-first search

© 2007 Pearson Addison-Wesley. All rights reserved

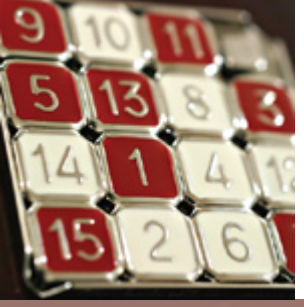

Figure 10.4 A small portion of the eight-puzzle's state graph

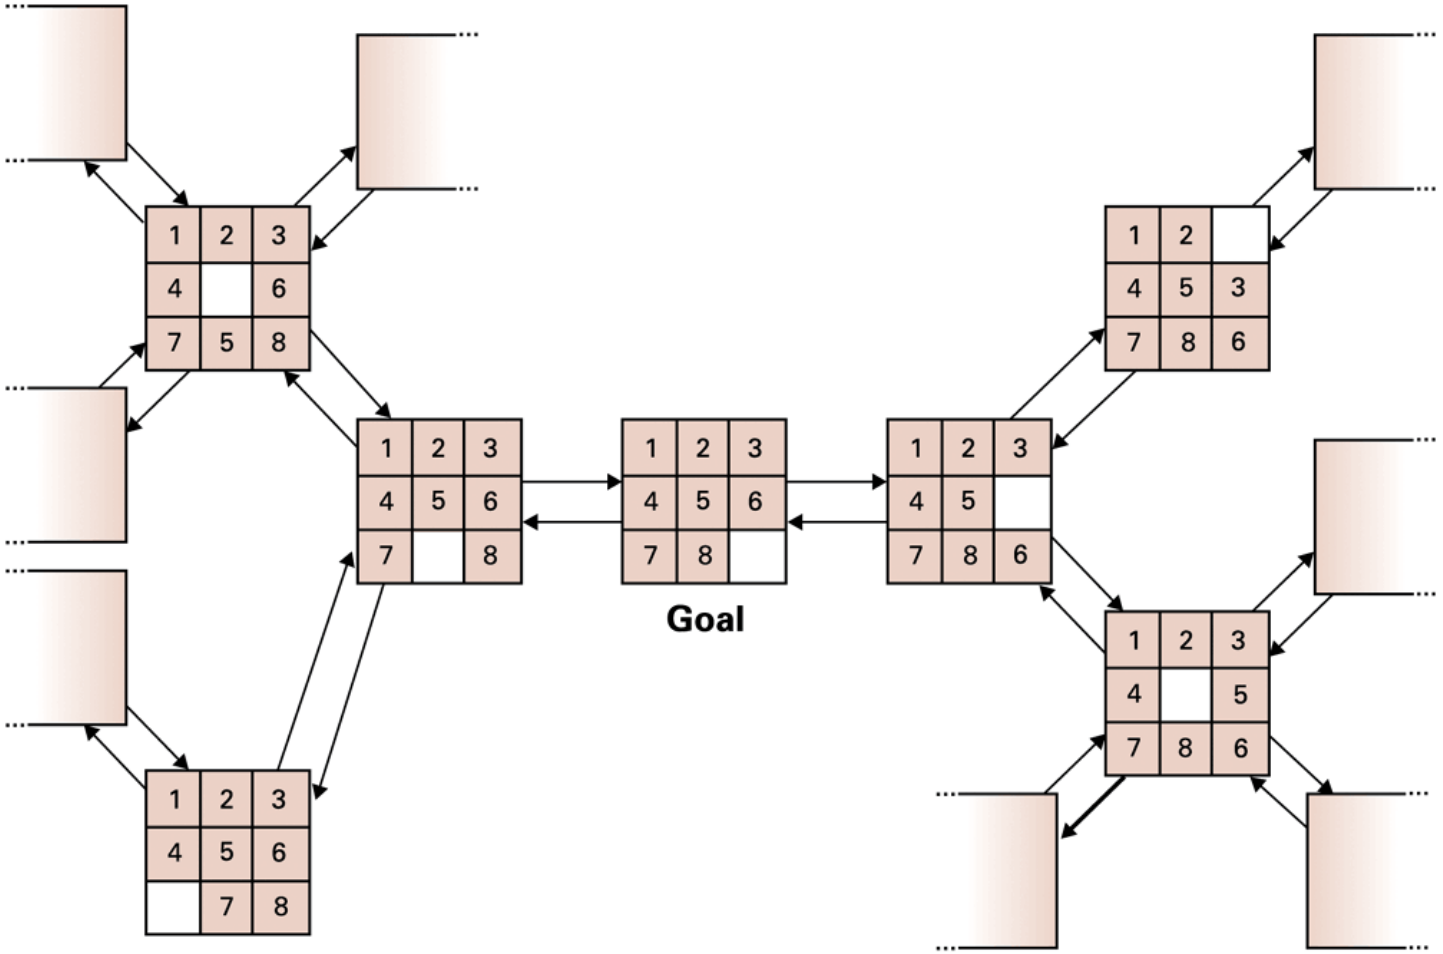

© 2007 Pearson Addison-Wesley. All rights reserved

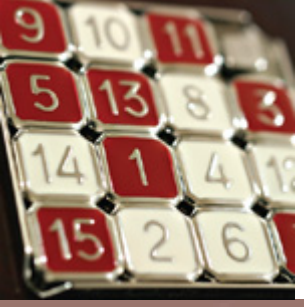

Figure 10.5 Deductive reasoning in the context of a production system

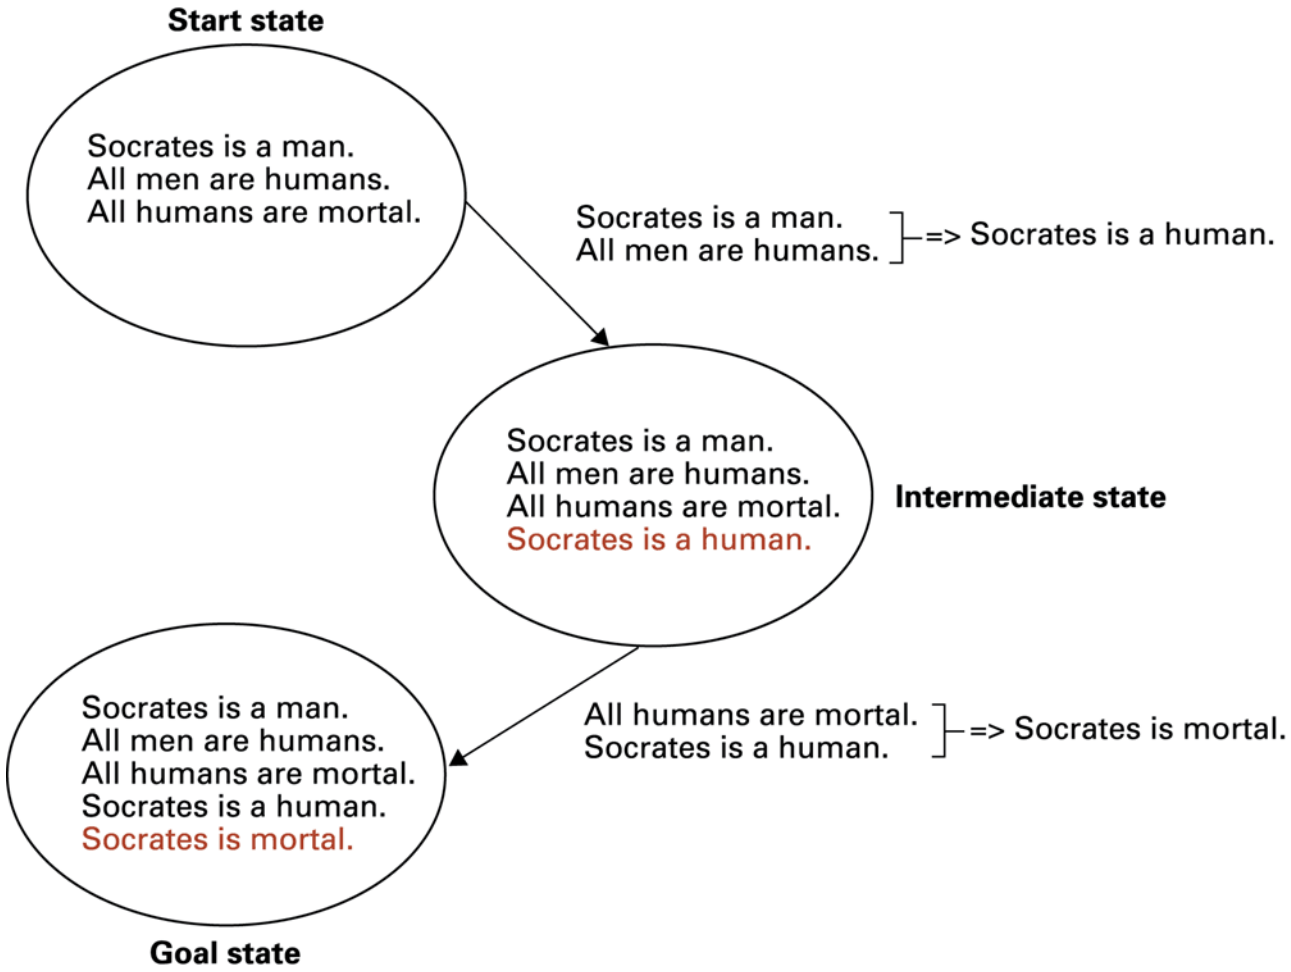

© 2007 Pearson Addison-Wesley. All rights reserved

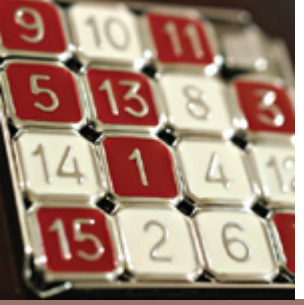

Figure 10.6 An unsolved eight-puzzle

1	3	5
4	2	
7	8	6

© 2007 Pearson Addison-Wesley. All rights reserved

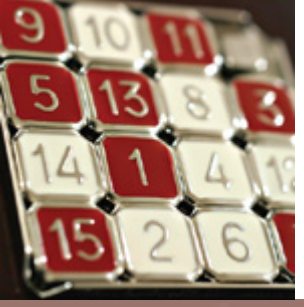

Figure 10.7 A sample search tree

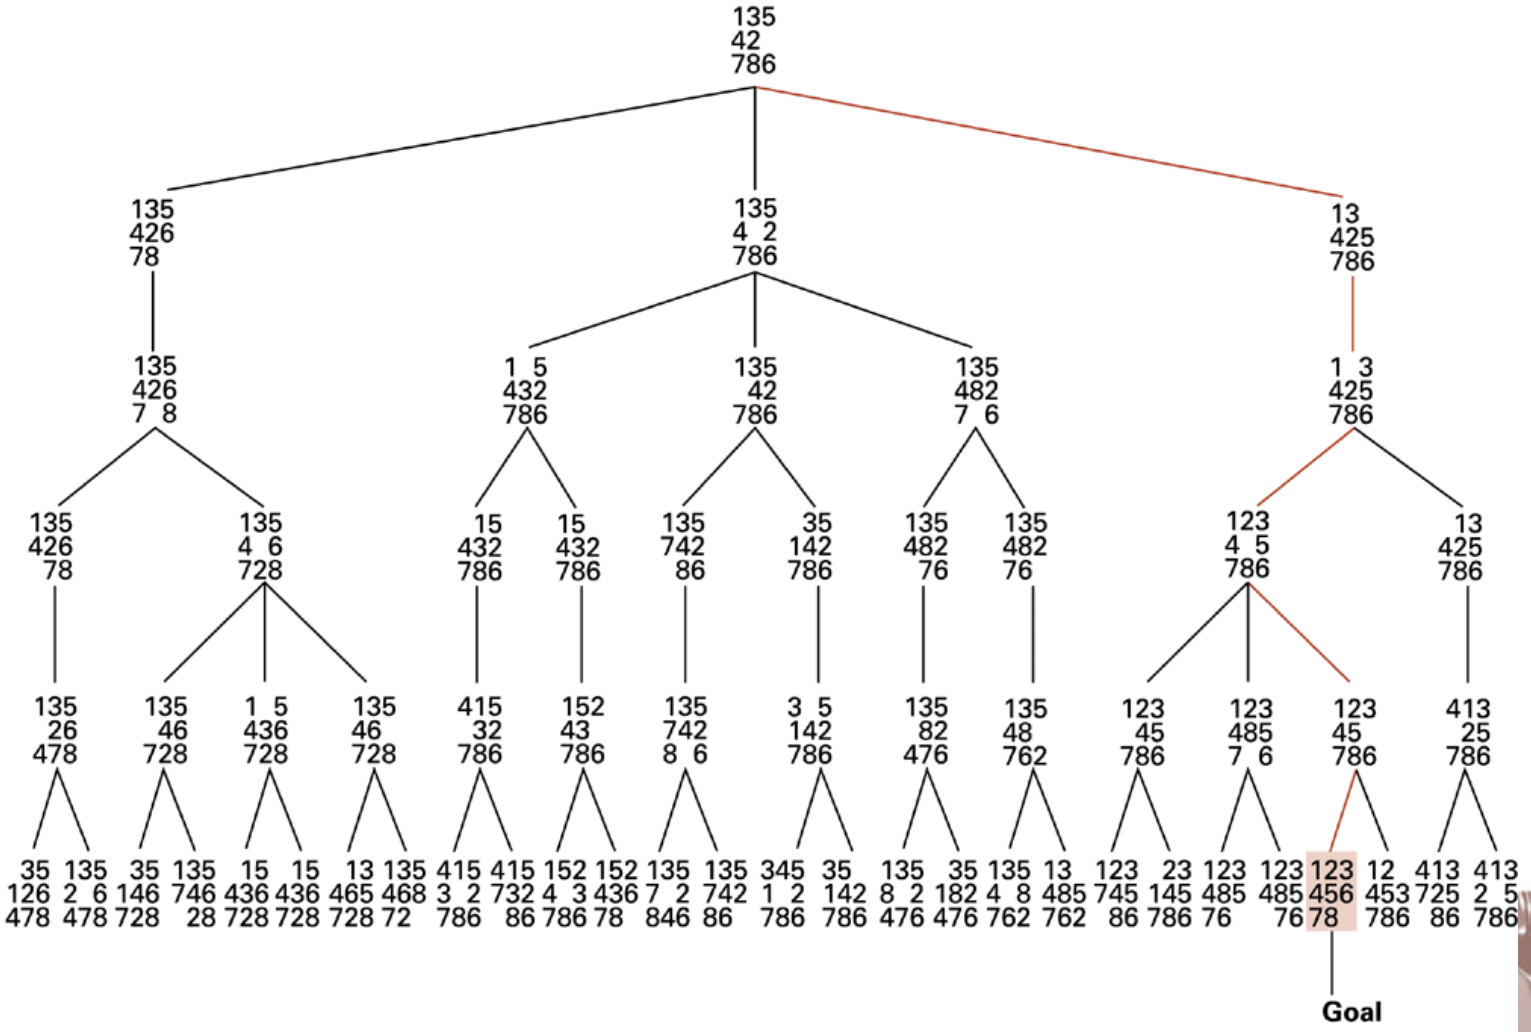

© 2007 Pearson Addison-Wesley. All rights reserved

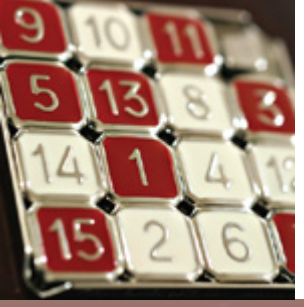

Figure 10.8 Productions stacked for later execution

Top of stack

Move the 5 tile down.

Move the 3 tile right.

Move the 2 tile up.

Move the 5 tile left.

Move the 6 tile up.

© 2007 Pearson Addison-Wesley. All rights reserved

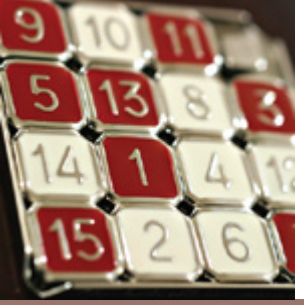

Heuristic Strategies

- **Heuristic:** A quantitative estimate of the distance to a goal
- Requirements for good heuristics
  - Must be much easier to compute than a complete solution
  - Must provide a reasonable estimate of proximity to a goal

© 2007 Pearson Addison-Wesley. All rights reserved

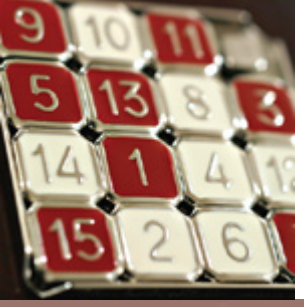

Figure 10.9 An unsolved eight-puzzle

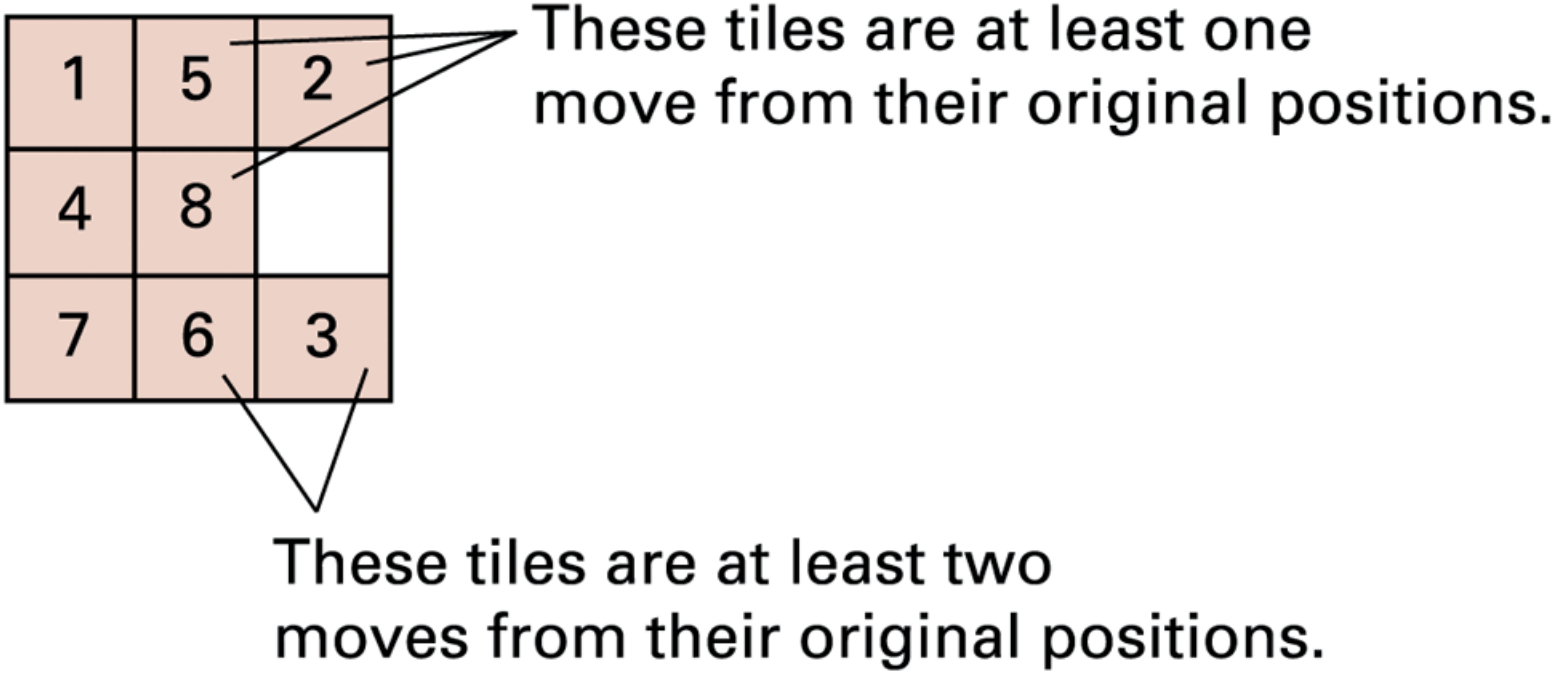

© 2007 Pearson Addison-Wesley. All rights reserved

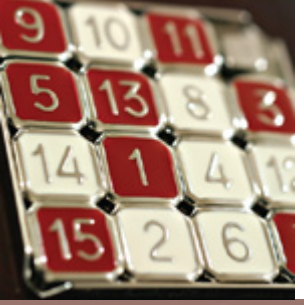

Figure 10.10 An algorithm for a control system using heuristics

Establish the start node of the state graph as the root of the search tree and record its heuristic value.
while (the goal node has not been reached) **do**
[Select the leftmost leaf node with the smallest heuristic value of all leaf nodes.
To this selected node attach as children those nodes that can be reached by a single production.
Record the heuristic of each of these new nodes next to the node in the search tree
]
Traverse the search tree from the goal node up to the root, pushing the production associated with each arc traversed onto a stack.
Solve the original problem by executing the productions as they are popped off the stack.

© 2007 Pearson Addison-Wesley. All rights reserved

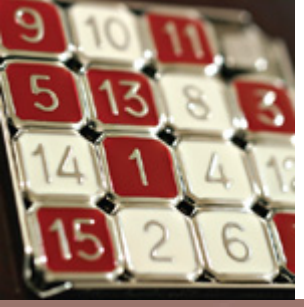

Figure 10.11 The beginnings of our heuristic search

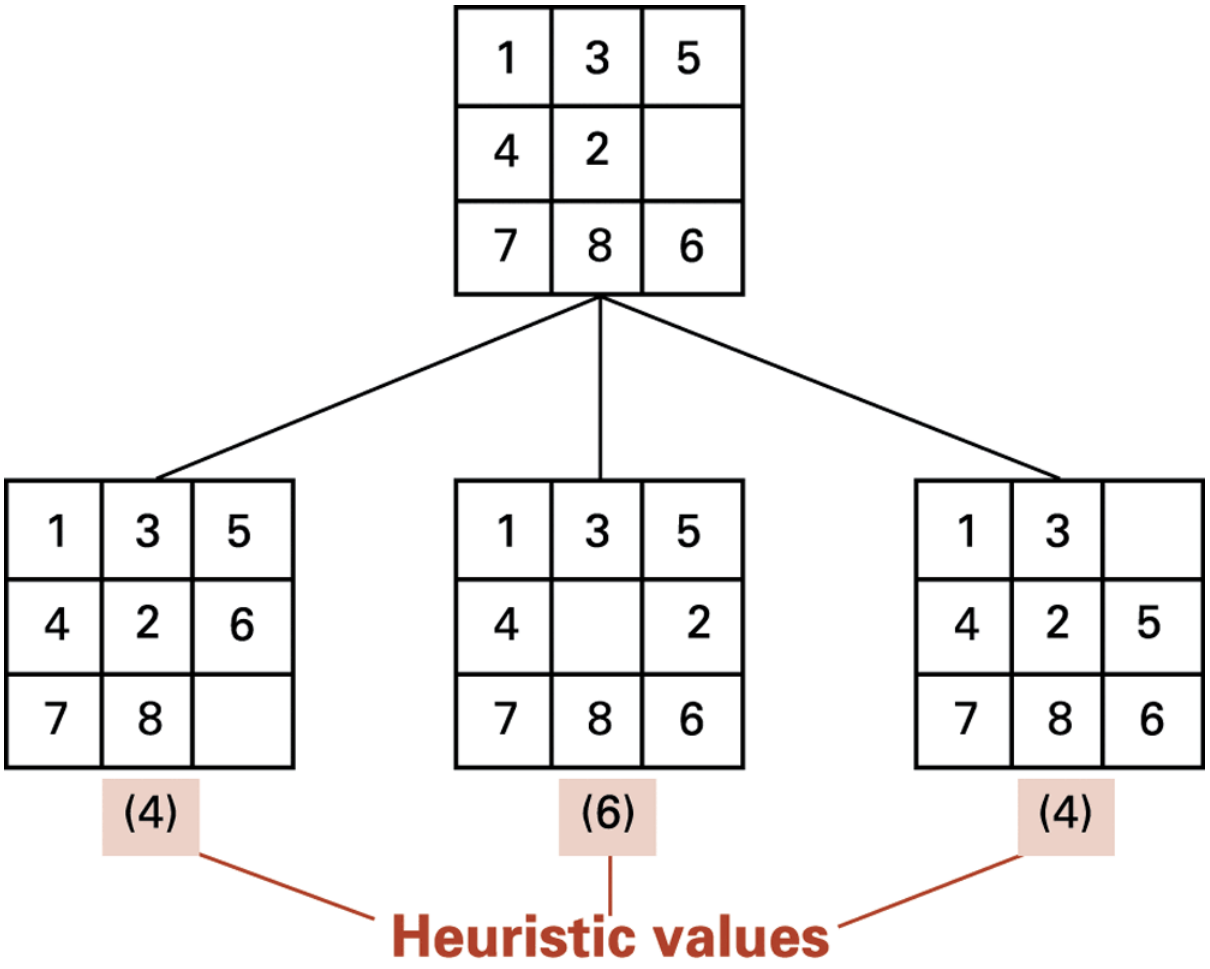

© 2007 Pearson Addison-Wesley. All rights reserved

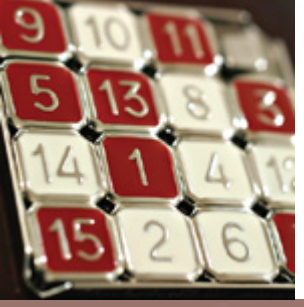

Figure 10.12 The search tree after two passes

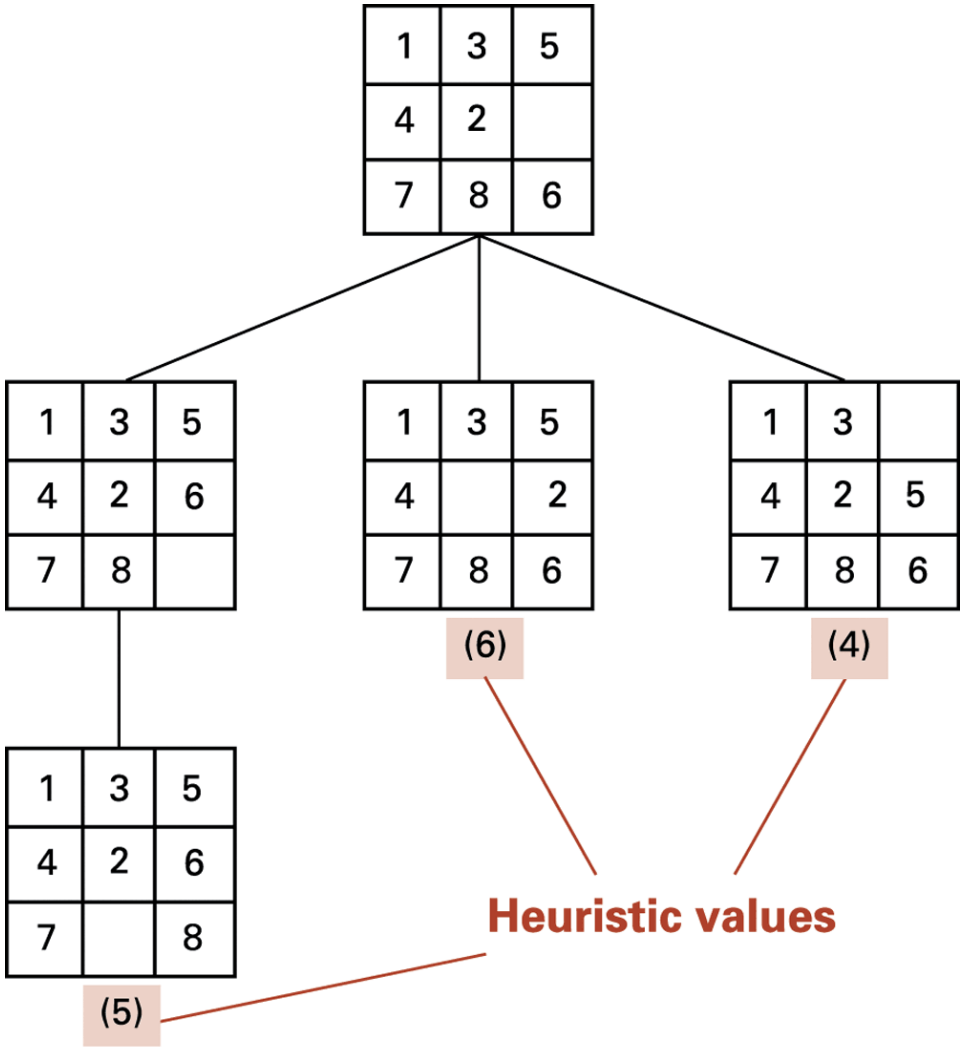

© 2007 Pearson Addison-Wesley. All rights reserved

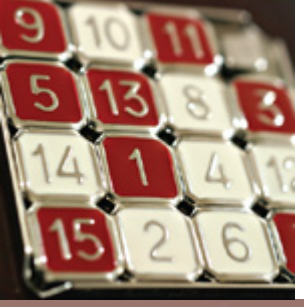

Figure 10.13 The search tree after three passes

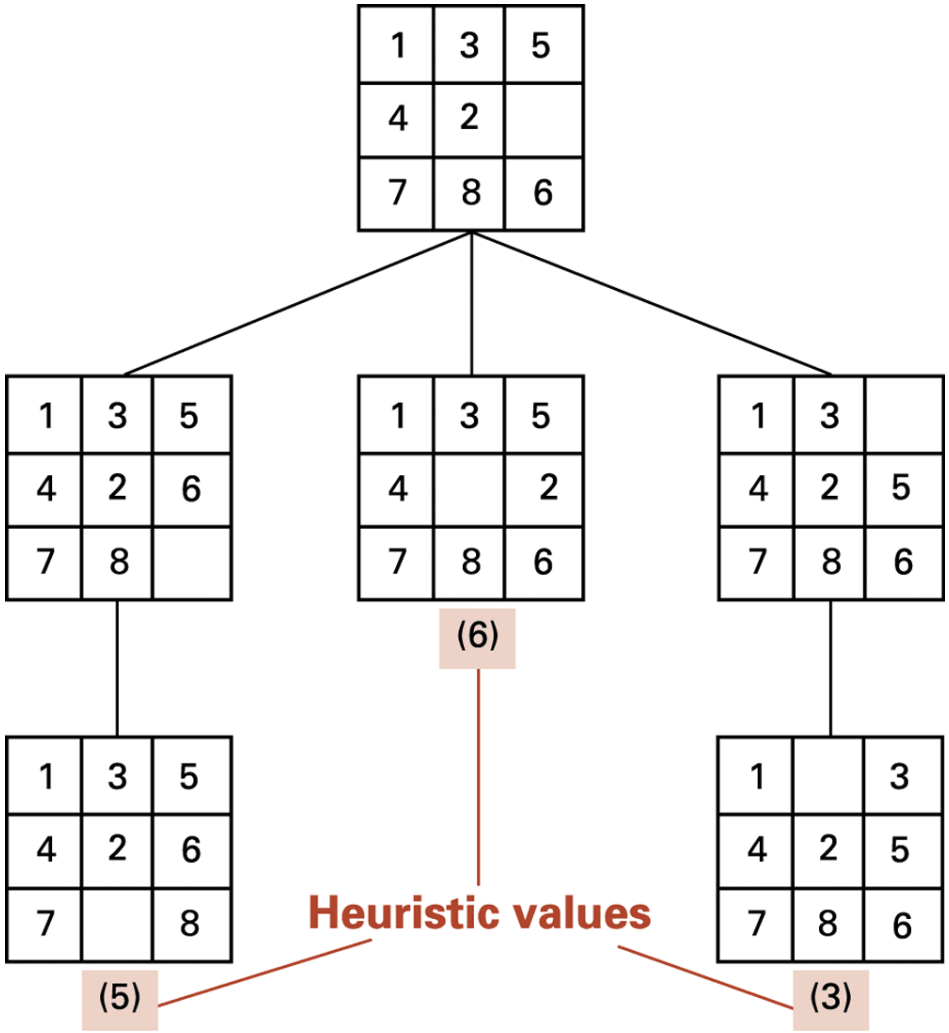

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.14
The complete search tree formed by our heuristic system

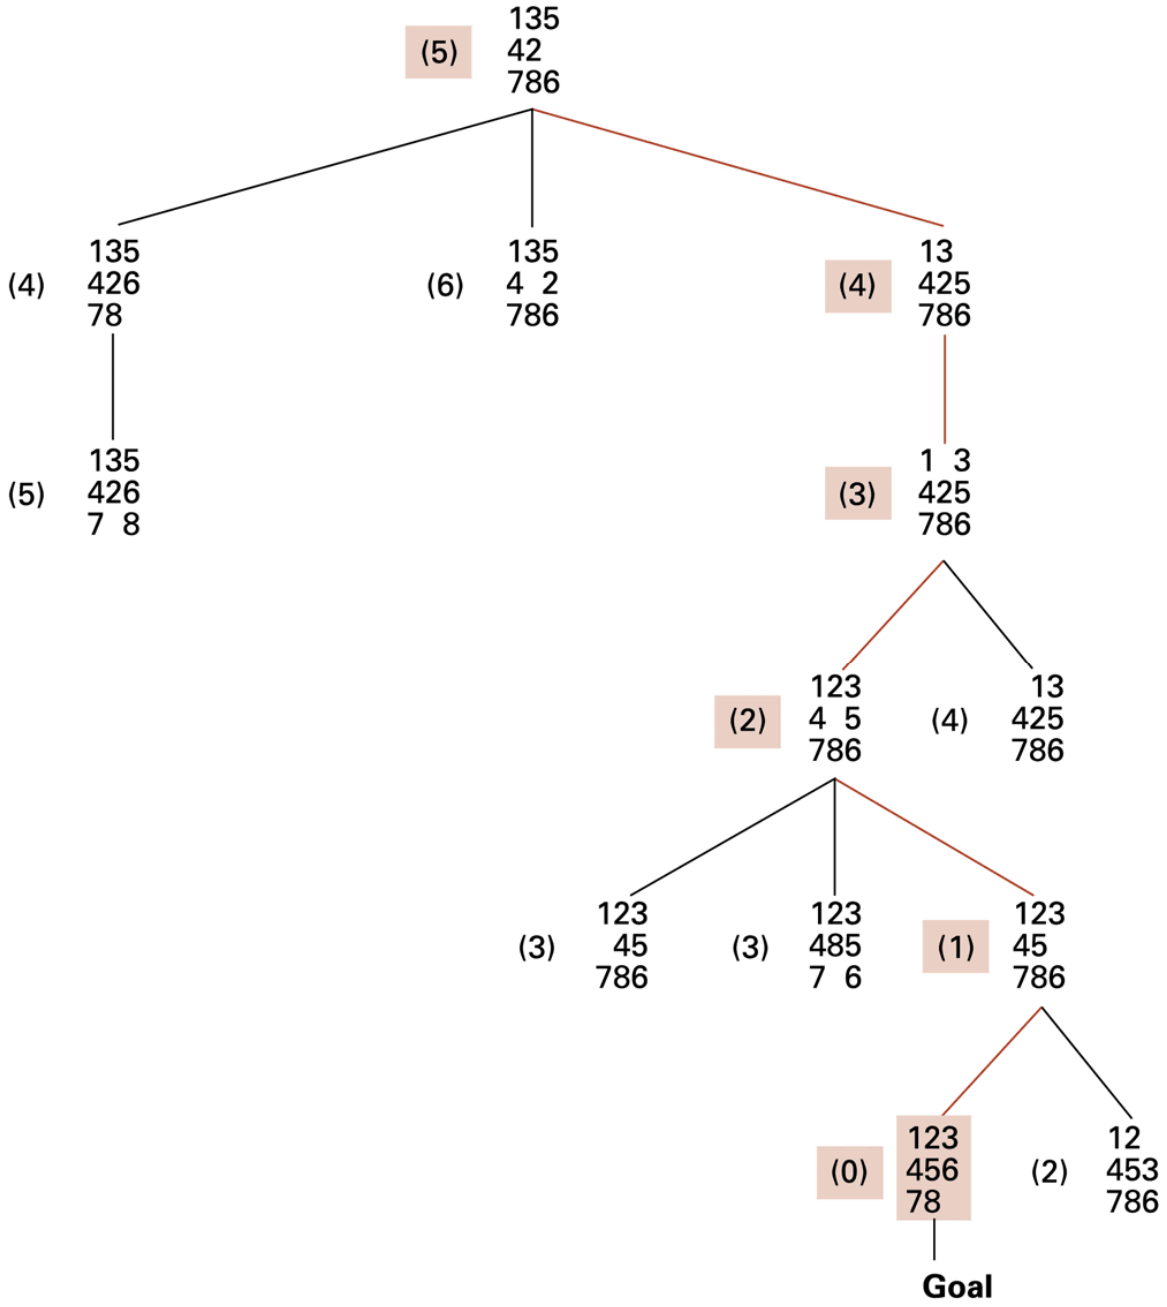

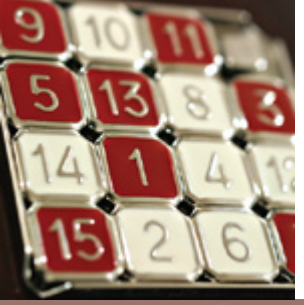

Handling Real-World Knowledge

- Representation and storage
- Accessing relevant information
  - Meta-Reasoning
  - Closed-World Assumption
- Frame problem

© 2007 Pearson Addison-Wesley. All rights reserved

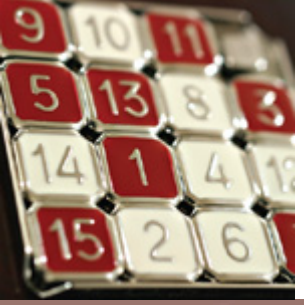

Learning

- Imitation
- Supervised Training
- Reinforcement
- Evolutionary Techniques

© 2007 Pearson Addison-Wesley. All rights reserved

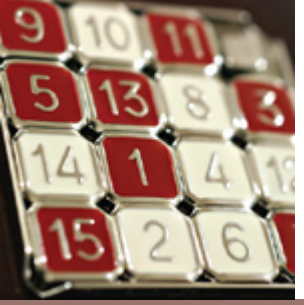

Artificial Neural Networks

- Artificial Neuron
  - Each input is multiplied by a weighting factor.
  - Output is 1 if sum of weighted inputs exceeds the threshold value; 0 otherwise.
- Network is programmed by adjusting weights using feedback from examples.

© 2007 Pearson Addison-Wesley. All rights reserved

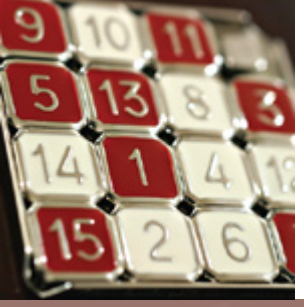

Figure 10.15 A neuron in a living biological system

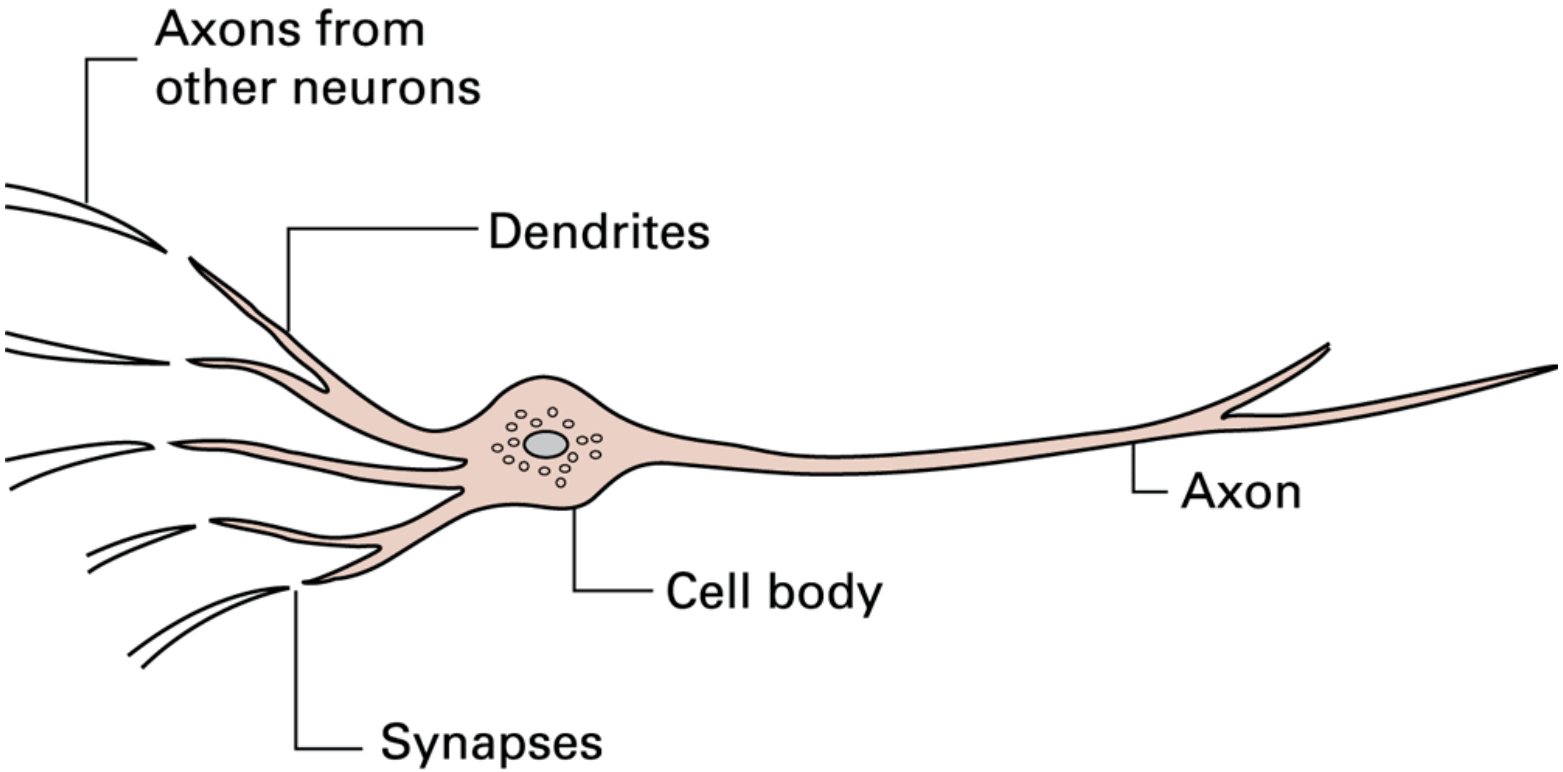

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.16 The activities within a processing unit

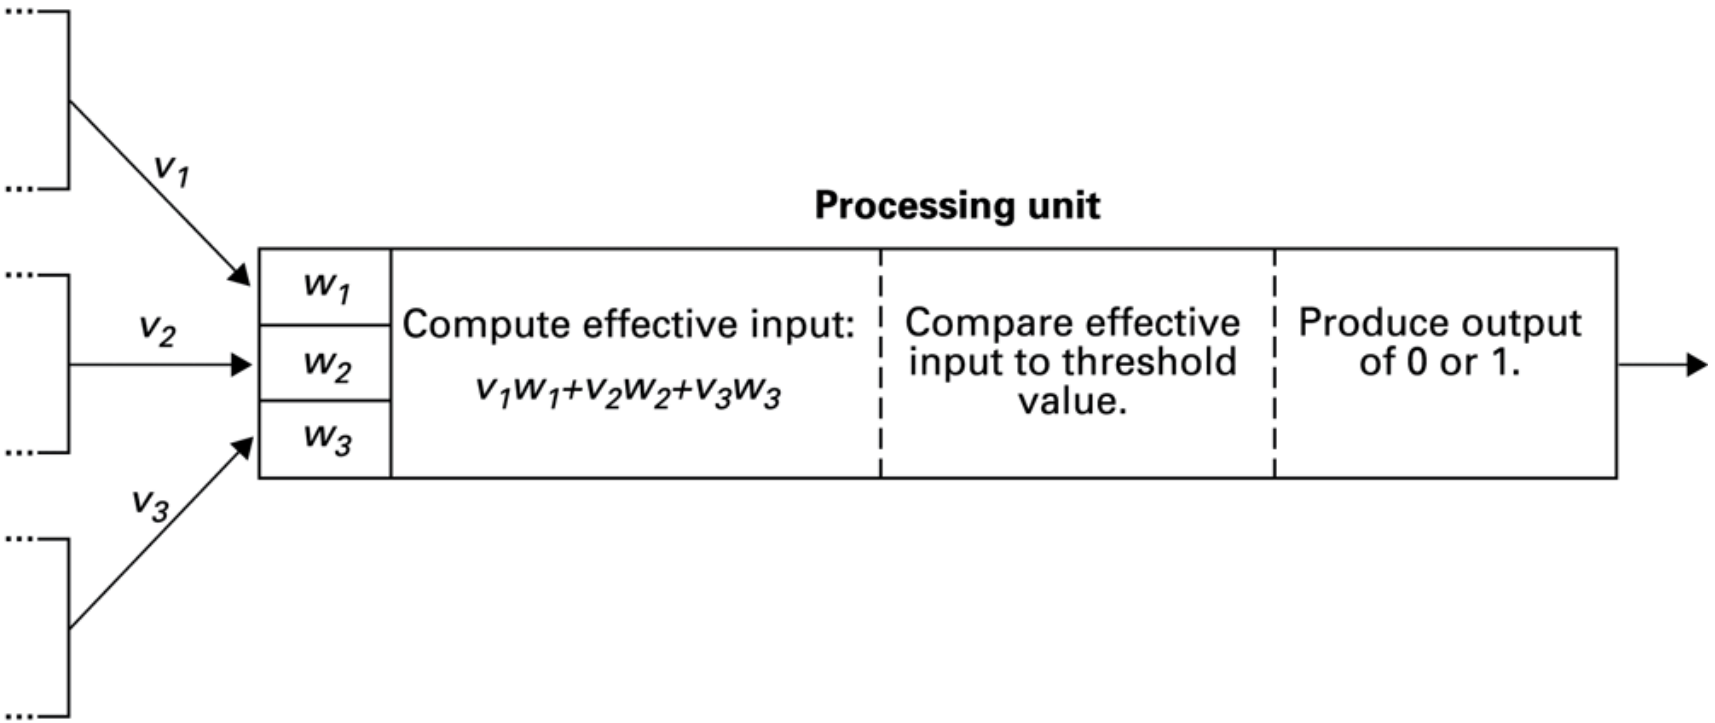

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.17 Representation of a processing unit

© 2007 Pearson Addison-Wesley. All rights reserved

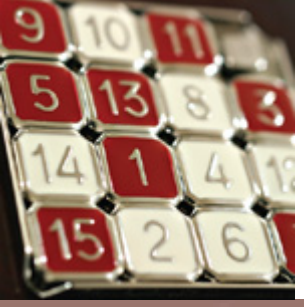

Figure 10.18 A neural network with two different programs

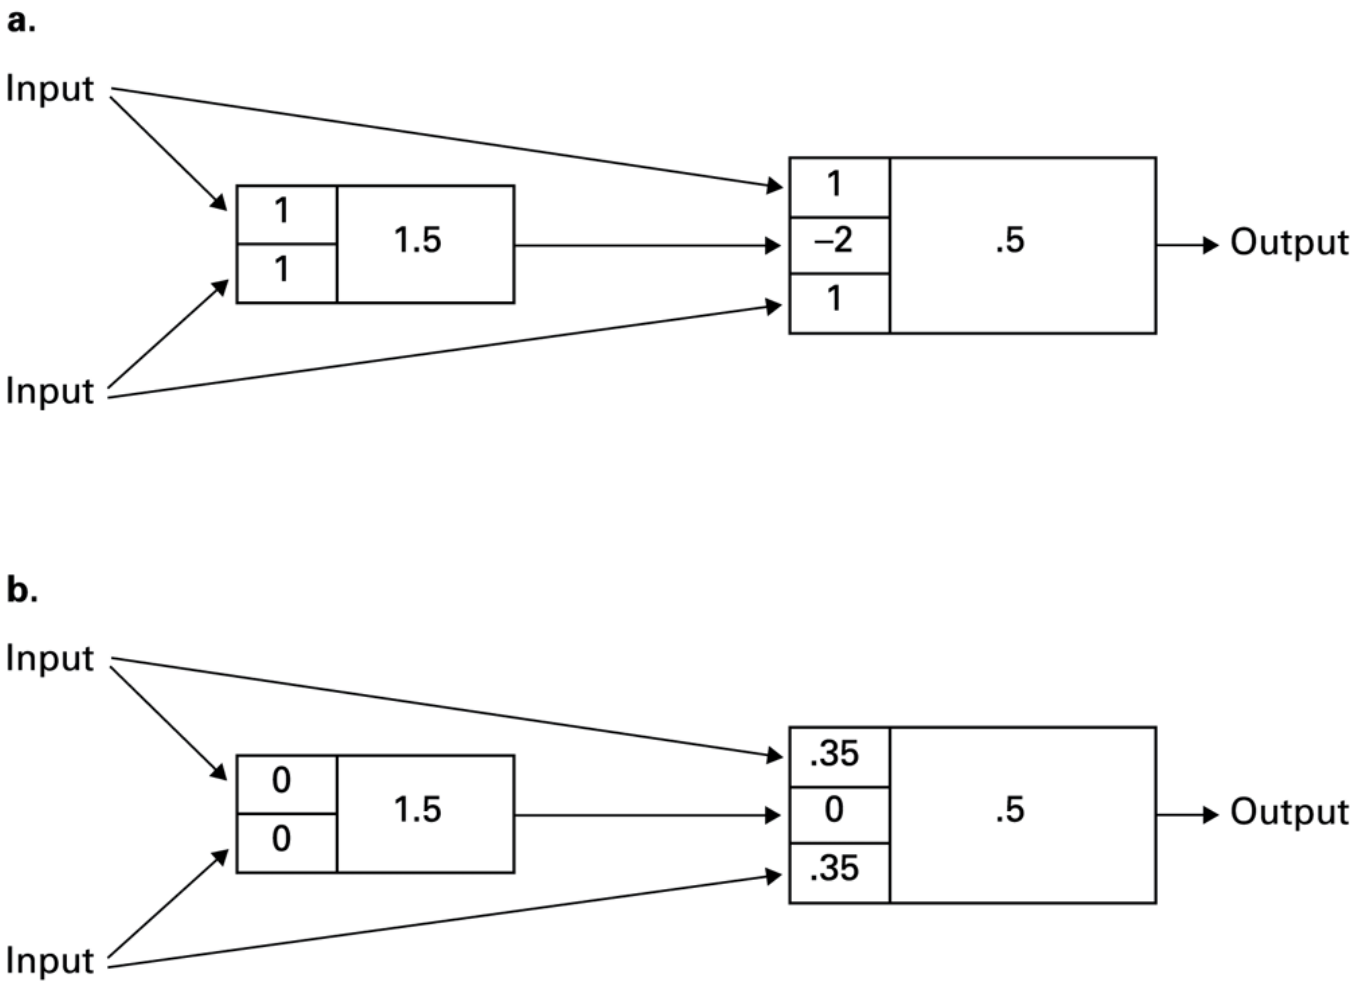

© 2007 Pearson Addison-Wesley. All rights reserved

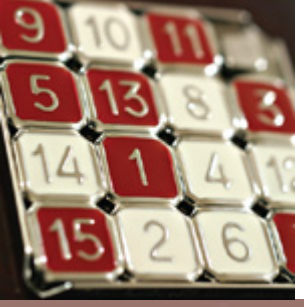

Figure 10.19 An artificial neural network

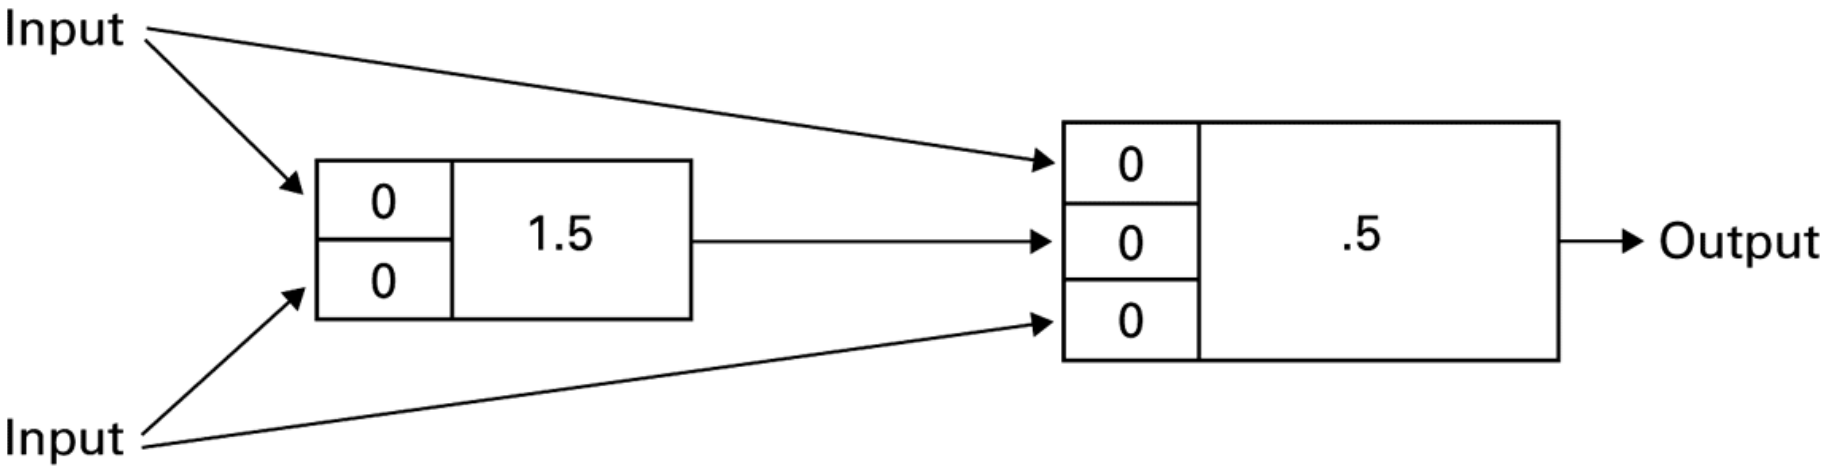

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.20 Training an artificial neural network

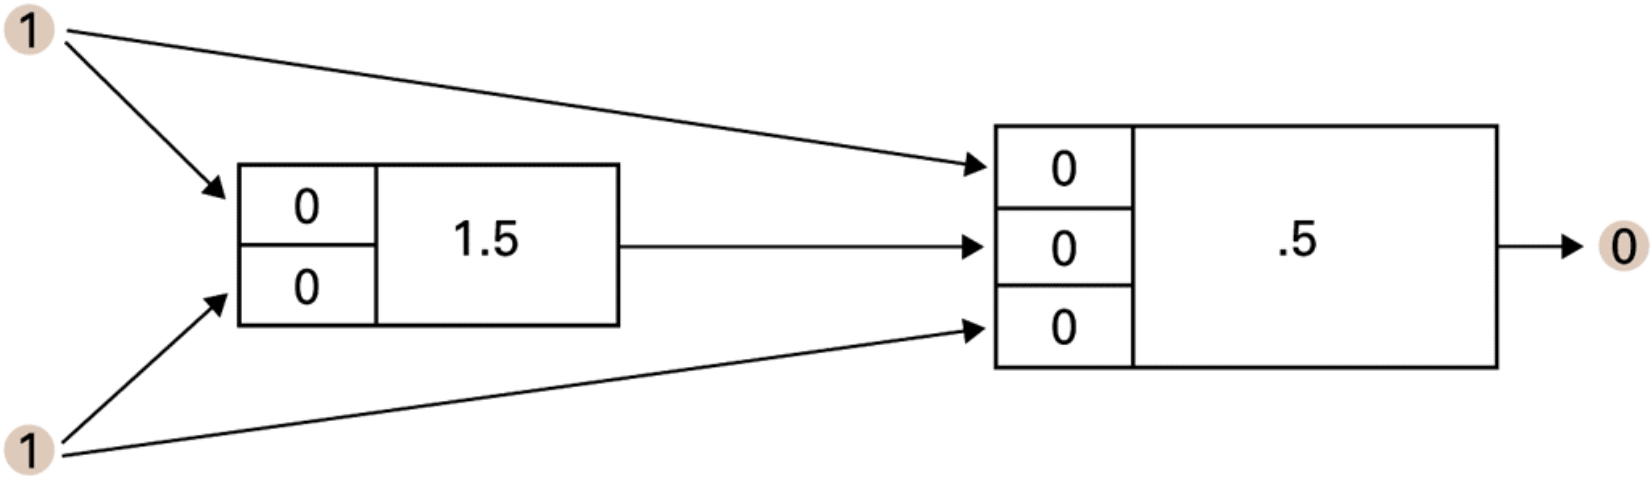

a. The network performs correctly for the input pattern 1, 1.

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.20 Training an artificial neural network (continued)

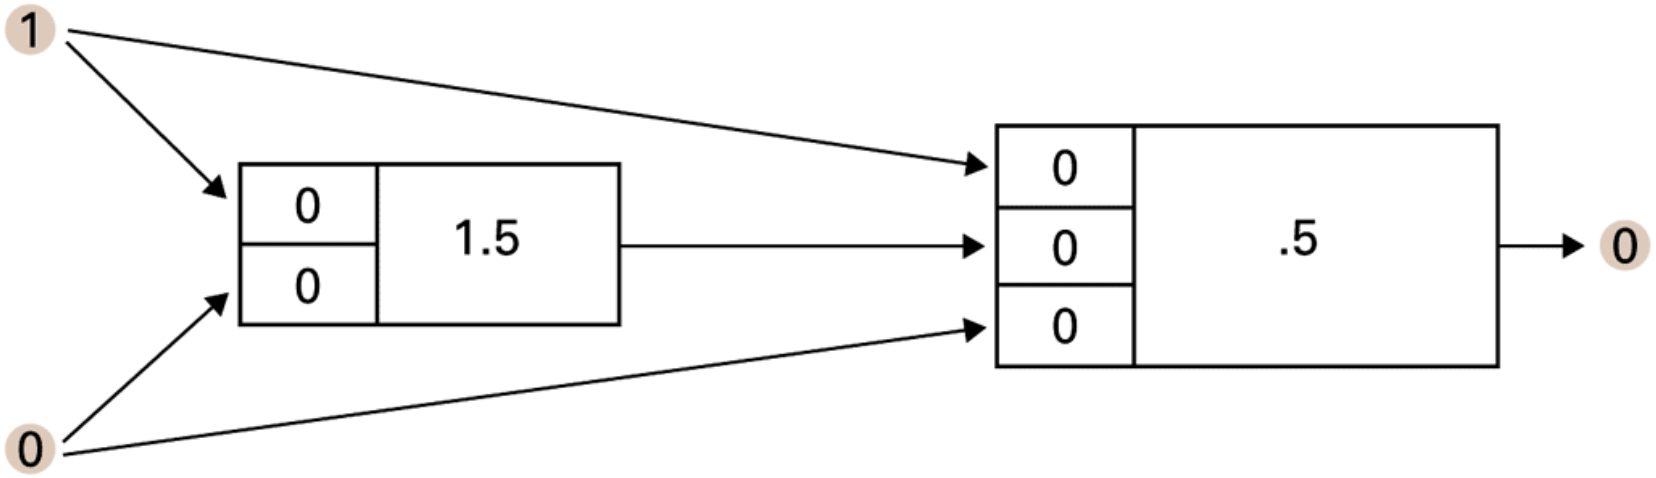

b. The network performs incorrectly for the input pattern 1, 0.

© 2007 Pearson Addison-Wesley. All rights reserved

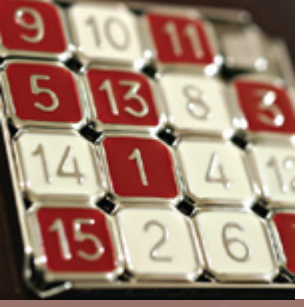

Figure 10.20 Training an artificial neural network (continued)

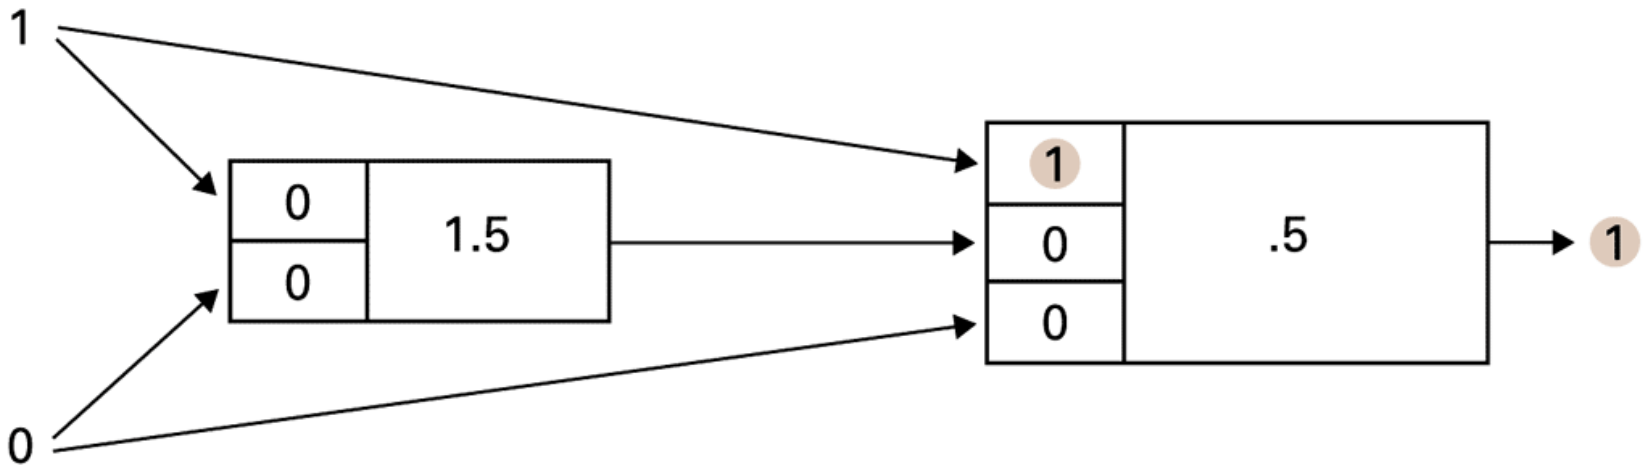

c. The upper weight in the second processing unit is adjusted.

© 2007 Pearson Addison-Wesley. All rights reserved

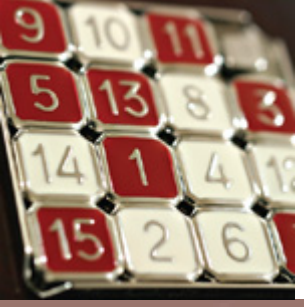

Figure 10.20 Training an artificial neural network (continued)

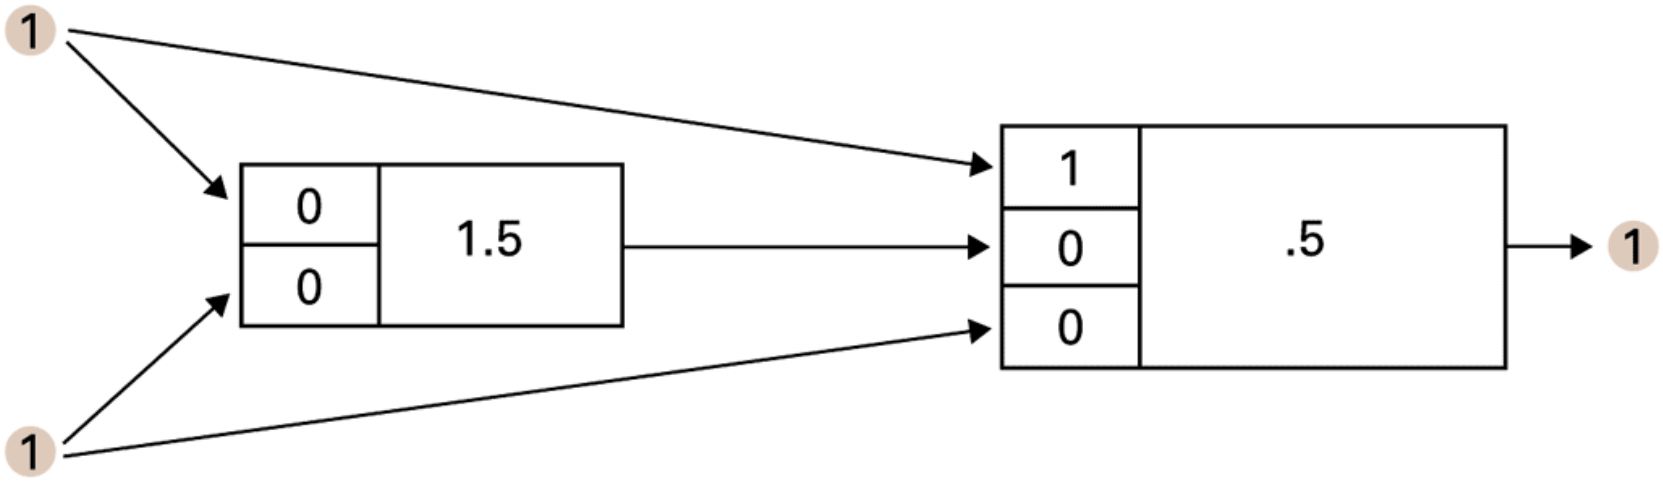

d. However, the network no longer performs correctly for the input pattern 1, 1.

© 2007 Pearson Addison-Wesley. All rights reserved

Figure 10.21 The structure of ALVINN

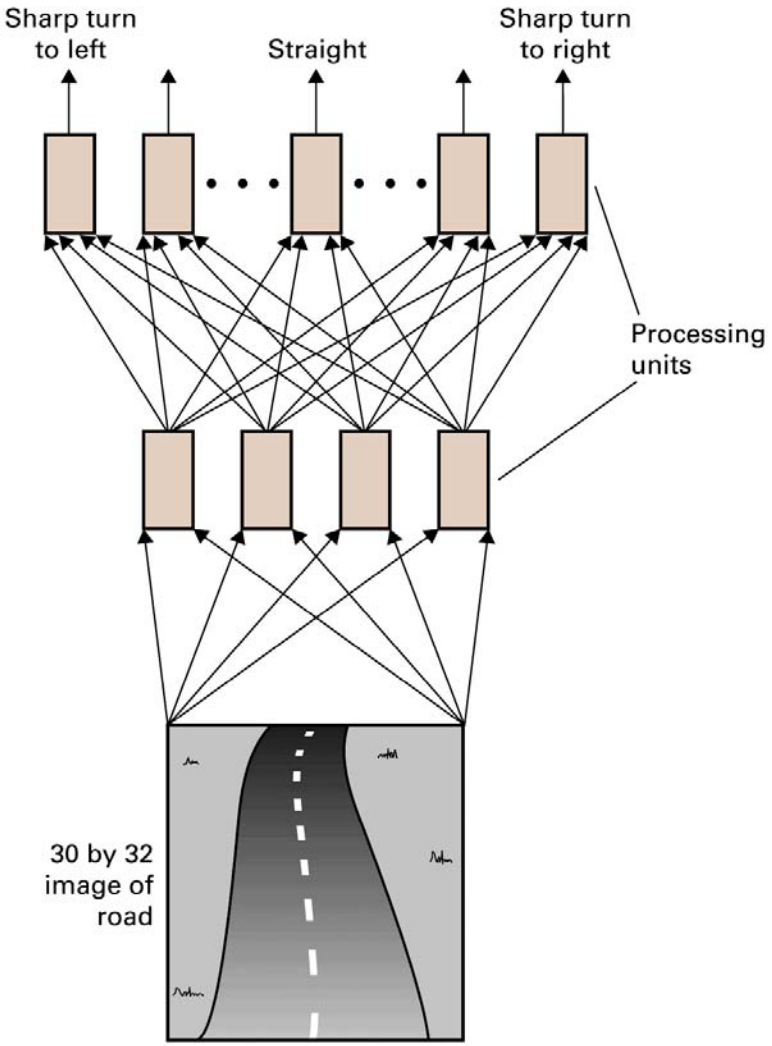

© 2007 Pearson Addison-Wesley. All rights reserved

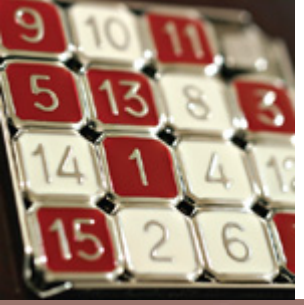

Associative Memory

- **Associative memory:** The retrieval of information relevant to the information at hand
- One direction of research seeks to build associative memory using neural networks that when given a partial pattern, transition themselves to a completed pattern.

© 2007 Pearson Addison-Wesley. All rights reserved

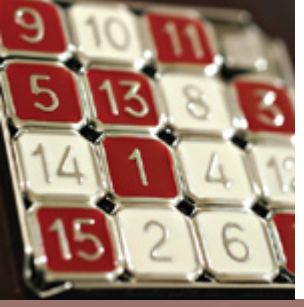

Figure 10.22 An artificial neural network implementing an associative memory

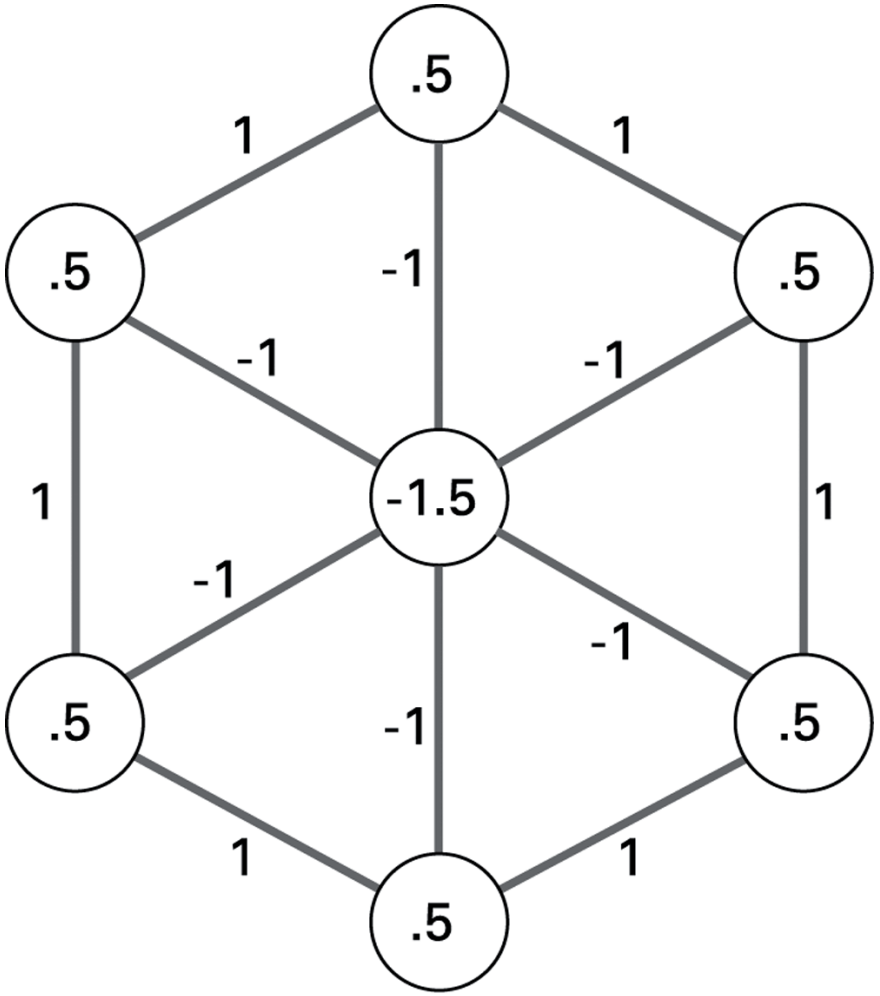

© 2007 Pearson Addison-Wesley. All rights reserved

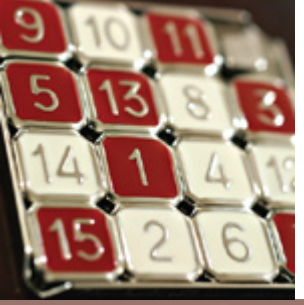

Figure 10.23 The steps leading to a stable configuration

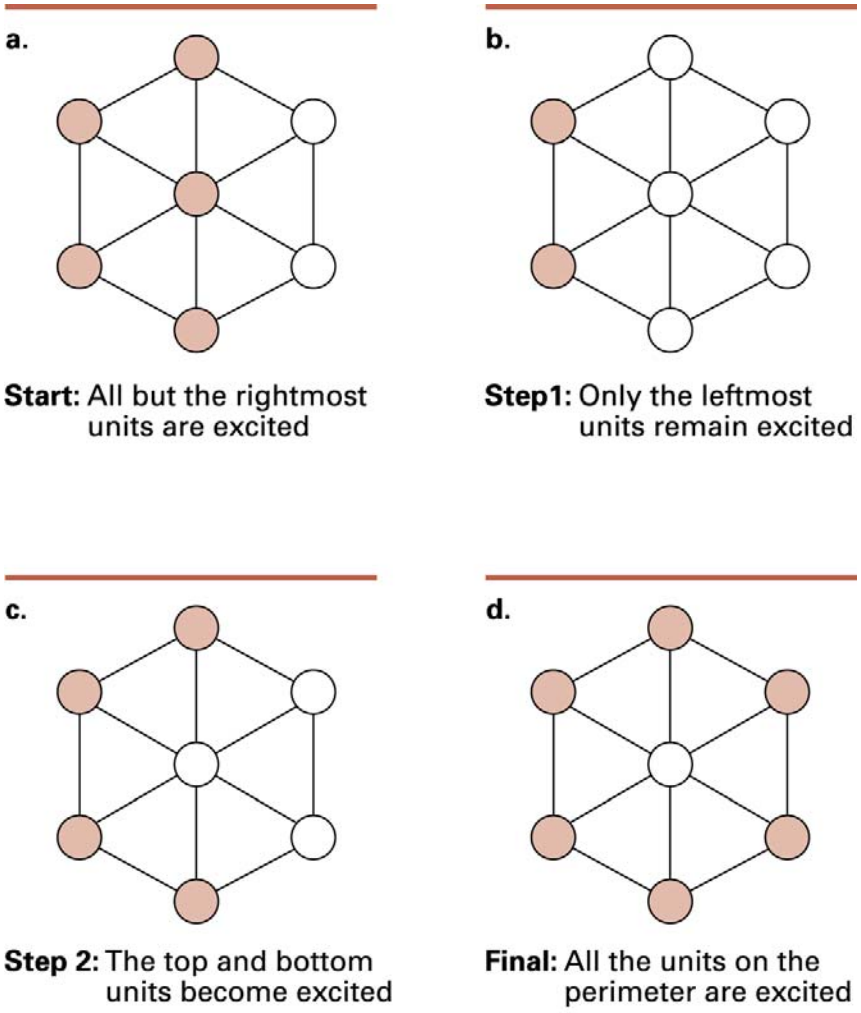

© 2007 Pearson Addison-Wesley. All rights reserved

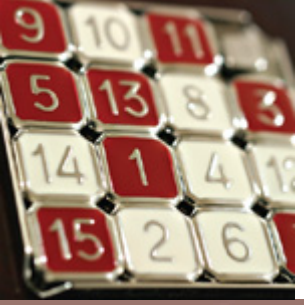

Robotics

- Truly autonomous robots require progress in perception and reasoning.
- Major advances being made in mobility
- Plan development versus reactive responses
- Evolutionary robotics

© 2007 Pearson Addison-Wesley. All rights reserved

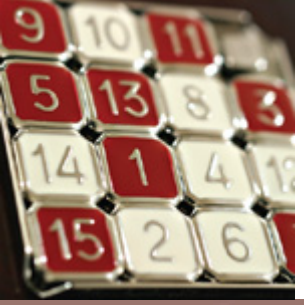

Issues Raised by Artificial Intelligence

- When should a computer's decision be trusted over a human's?
- If a computer can do a job better than a human, when should a human do the job anyway?
- What would be the social impact if computer "intelligence" surpasses that of many humans?

© 2007 Pearson Addison-Wesley. All rights reserved