Computer Science and Information Technology
前瞻資訊科技
Spring 2009
Jane Hsu
http://agents.csie.ntu.edu.tw/~yjhsu/courses/CSIT/AI20090417.pdf

ARTIFICIAL INTELLIGENCE
人工智慧

AI: Fact or Fiction?

How far are we from building a robotic child?

Lecture Outline 04/17/2009

- What is intelligence?
- The agent approach to AI
- Problem solving as search
- Game search
- Productive games for semantic annotation

The Human Brain

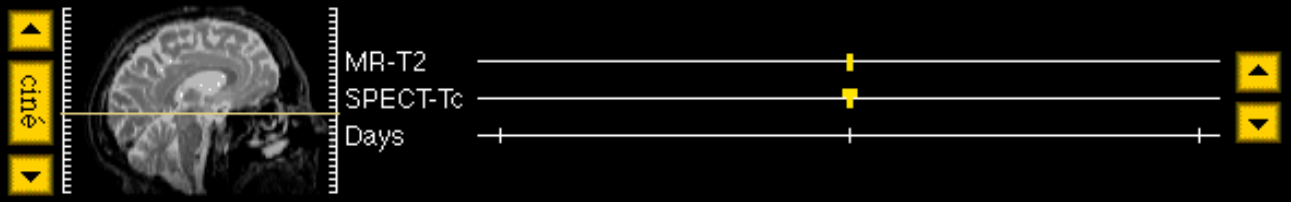

[Home][Help][Clinical]

Slice 23

the whole brain

Atlas

Click on sagittal image to select slice. Click on thin tickmark to change timepoint, or thick tickmark for overlay.

Keith A. Johnson (keith@bwh.harvard.edu), J. Alex Becker (jabecker@mit.edu)

Brain Structure

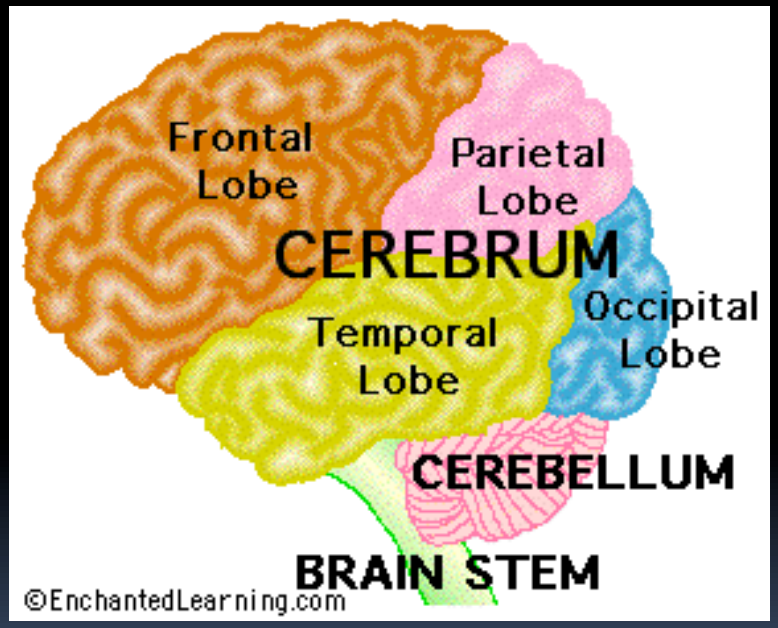

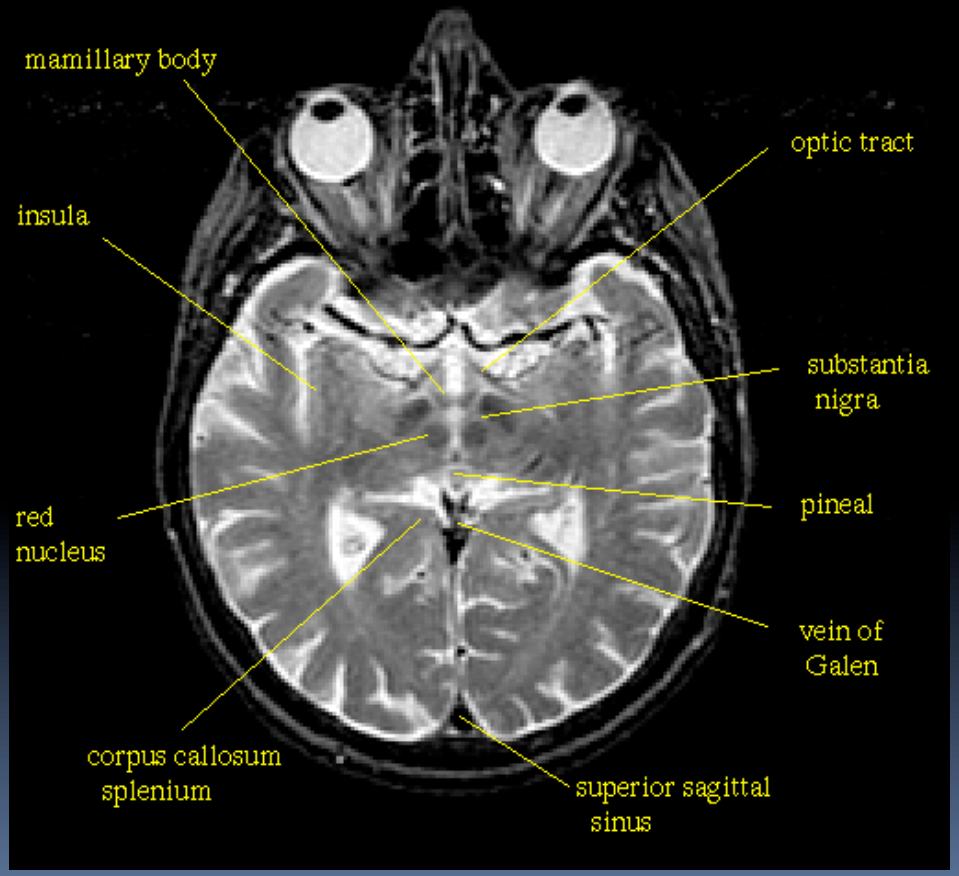

Brain Functions: The Cerebrum

Frontal Lobe

- Problem solving
- Creative/abstract thought
- Judgement
- Skilled movements

Occipital Lobe

- Vision
- Reading

Occipital Lobe

- Stereognosis: form from tough
- Sensory combination

Temporal Lobe

- Auditory memories
- Visual memories
- Music, some speech/language
- Sense of identity

Right Hemisphere

- Temporal/spatial relationships
- Communicating emotion

Left Hemisphere

- Produce/understand language

Brain Functions: Others

The Cerebellum

- Balance
- Posture
- Cardiac, respiratory, vasomotor centers

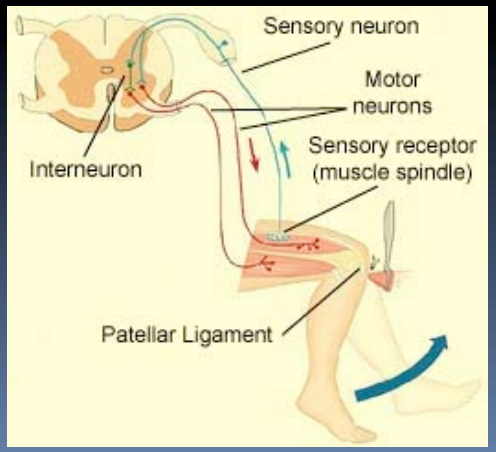

The Brain Stem

- Motor/sensory pathway
- Vital centers

Hypothalamus

- Moods and motivation
- Sexual maturation

Spinal Cord

- Conduit/sense of sensation

Dimensions of Human Intelligence

- Linguistic
- Logico-mathematical
- Spatial
- Musical
- Kinesthetic
- Intrapersonal
- Interpersonal

The Turing Test

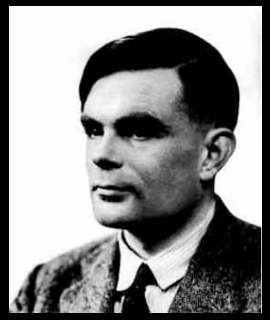

- "Can machines think?" → "Can machines behave intelligently?"
- Imagine that you are typing into a computer terminal. At the other end of the line is either another person or an artificial system of some sort. You have thirty minutes to ask whatever you want; if, at the end of that time, you cannot reliably distinguish the human from the artificial respondent, the artificial system is deemed to be generally intelligent.

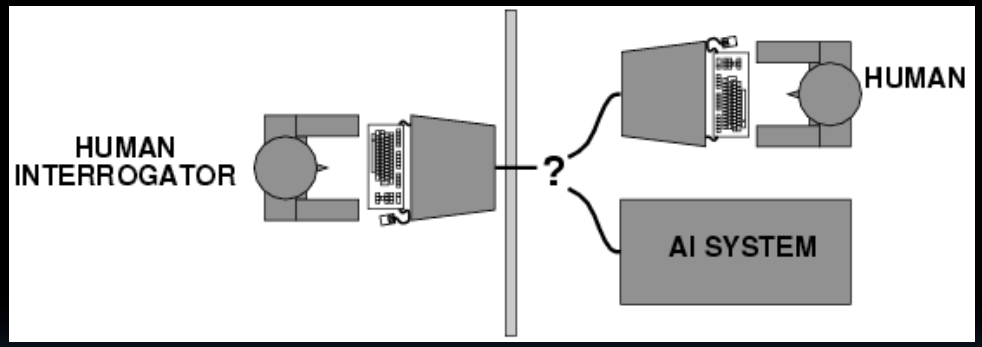

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problems with The Turing Test

- The human interrogator may be incompetent.
- The human interrogator is too lazy to ask any questions.
- The human at the other end may try to trick the interrogator.
- The machine may store (if possible) all possible sequences of up to 18,000 characters together with appropriate responses.
- The communication channel is too narrow. No gestures, facial expressions, or physical contacts may be used.
- The test equates intelligence with conversational human-like behavior.
- Philosophical objections
- The implication of such a test is: A program doesn't have to think like a human. Intelligence is really decided by what a program (or other agent) does, not how it does it.

Taxonomy of AI

- **Acting humanly:** Turing test
  - Natural language
  - Automated reasoning
  - Machine learning
  - Computer vision
  - Humanoid Robots
- **Thinking humanly:** cognitive modeling
  - Introspection
  - Psychological experiments
- **Acting rationally:** rational agent
  - Knowledge representation and reasoning
  - Natural language
  - Learning
  - Visual perception
  - Limited rationality
- **Thinking rationally:** laws of thought
  - Formal logic
  - (Correct) Inference

Agent

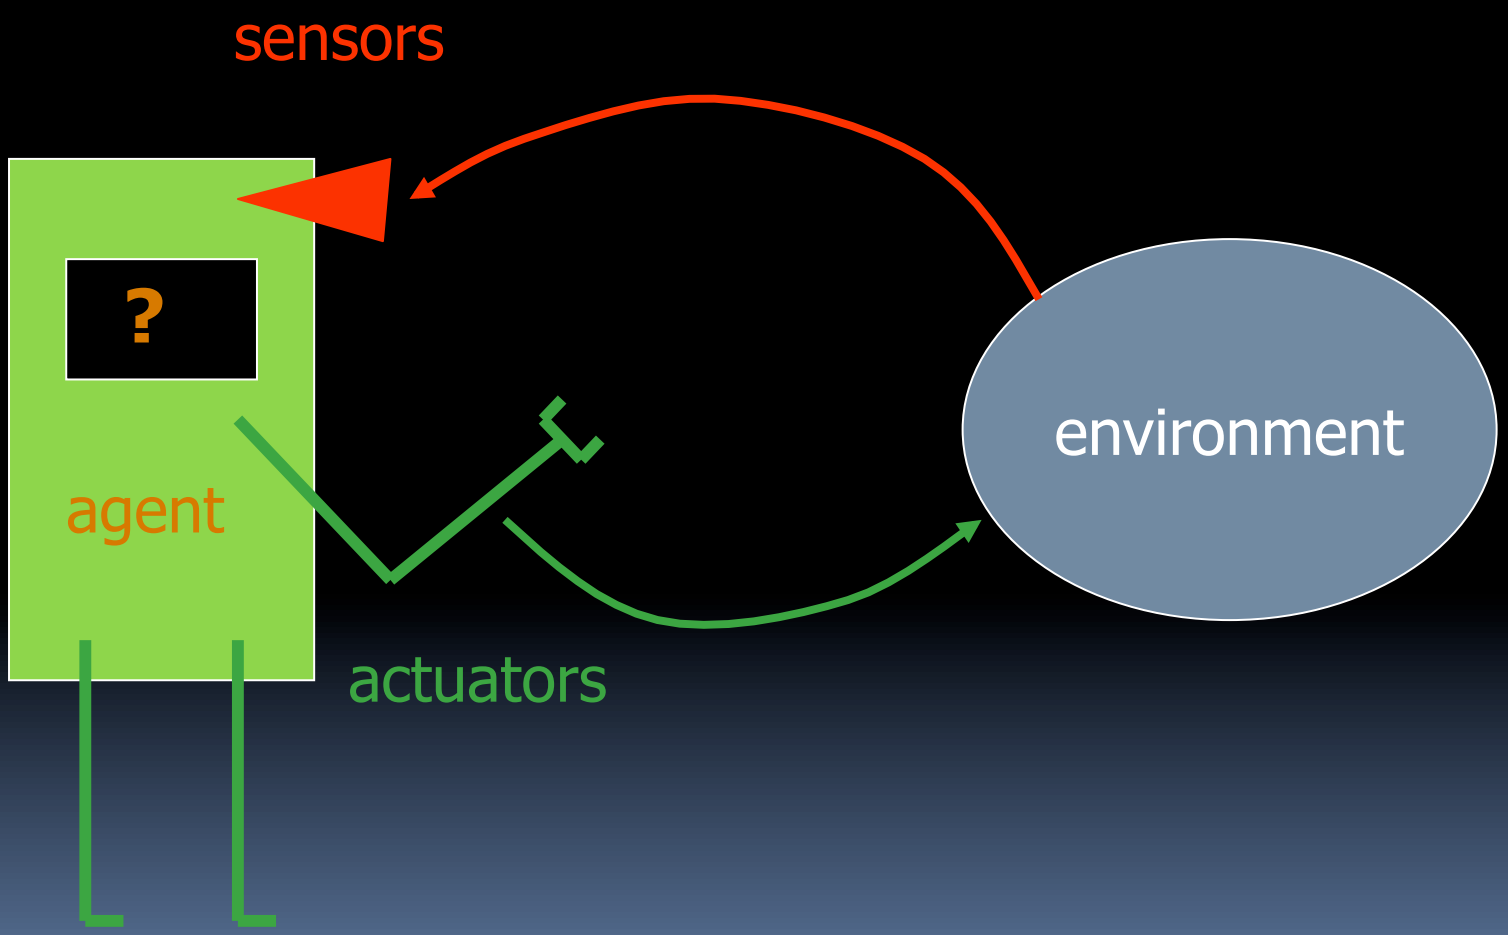

Simple Reflex Agents

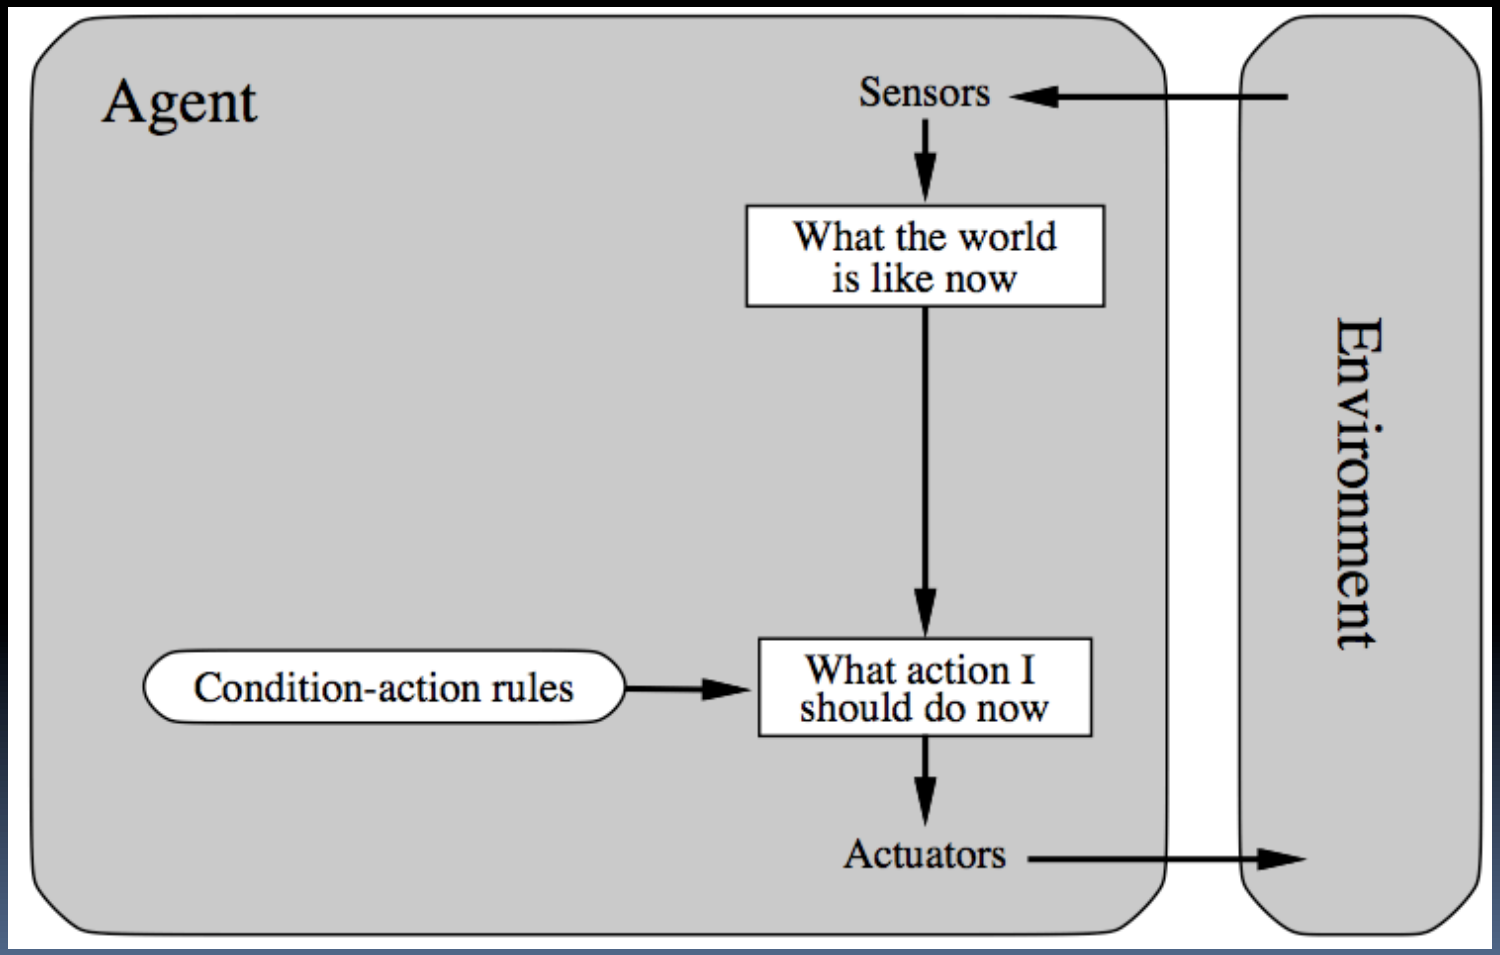

Model-Based Reflex Agents

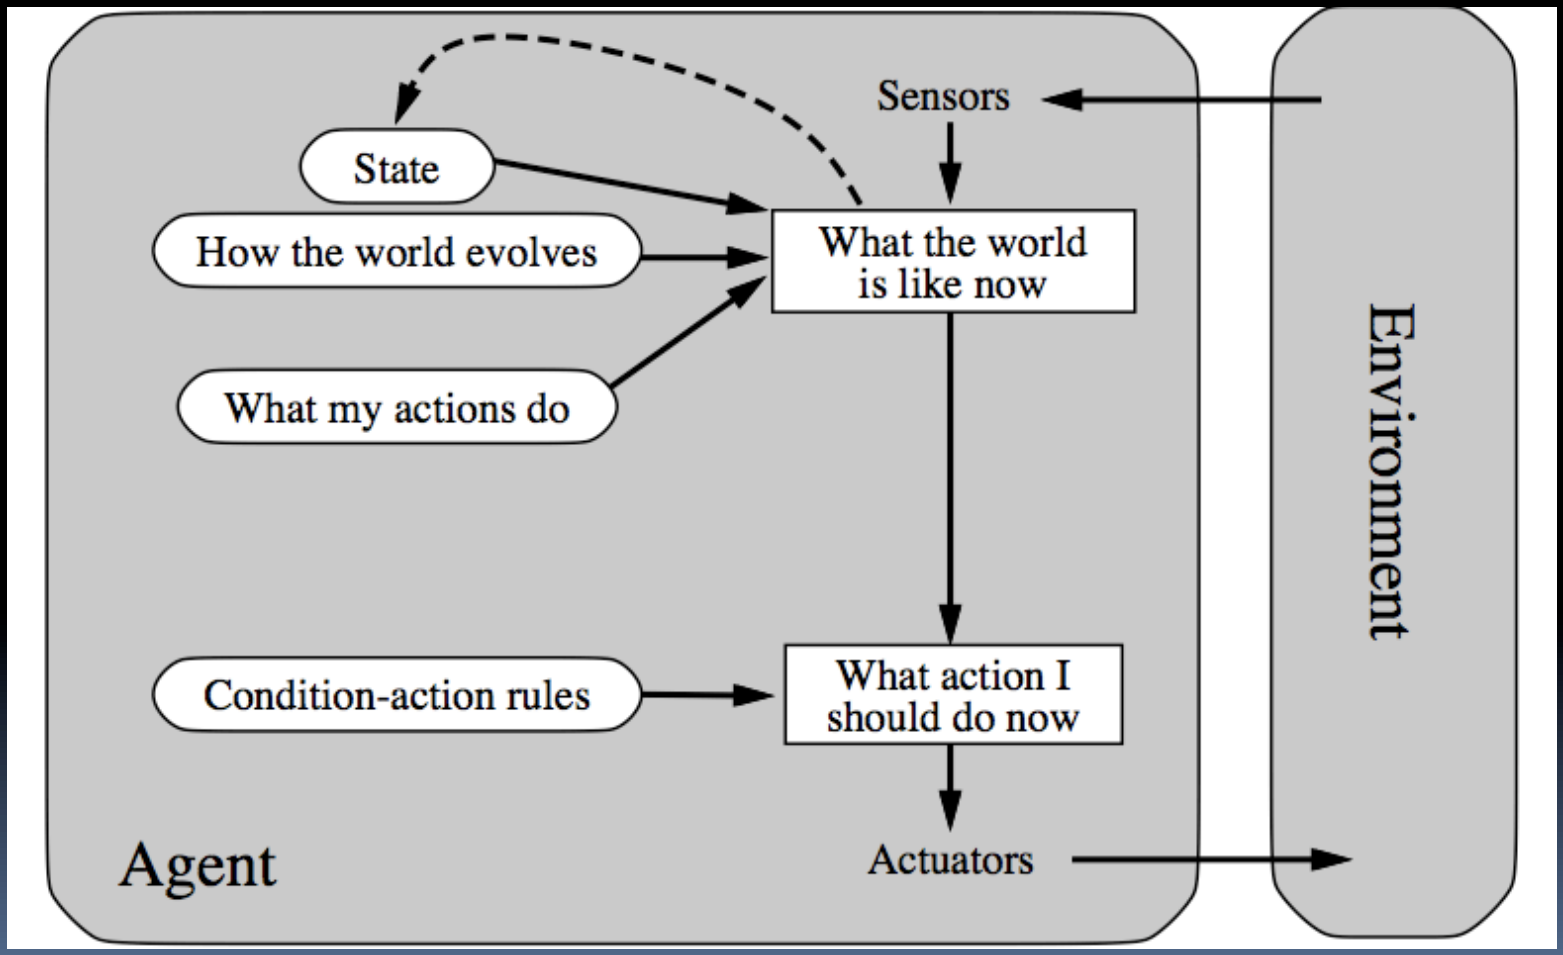

Goal-Based Agents

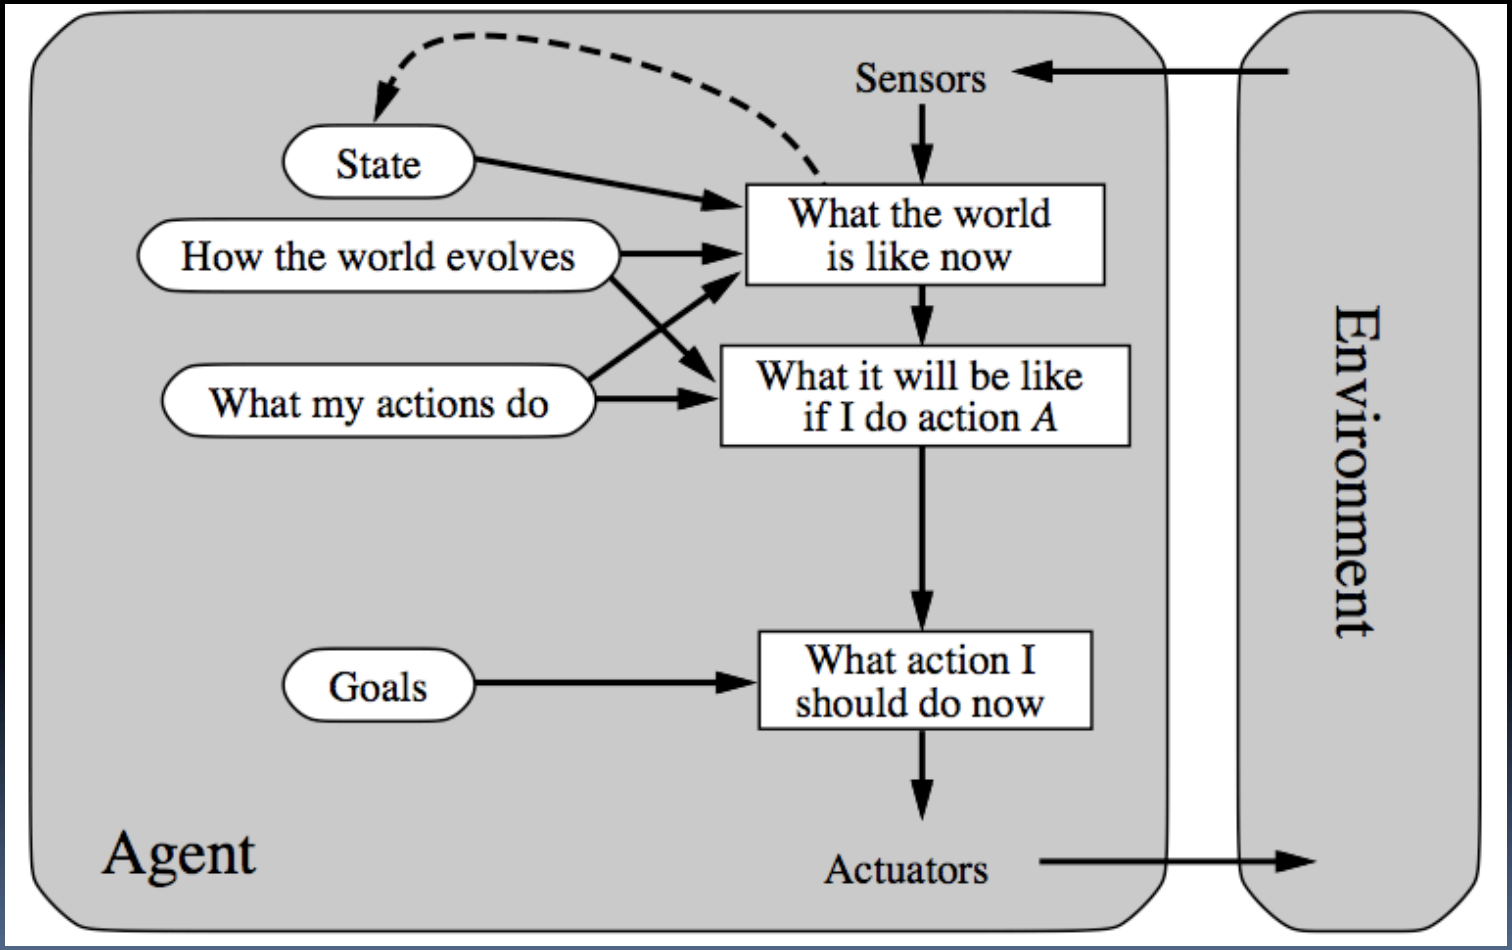

Utility-Based Agents

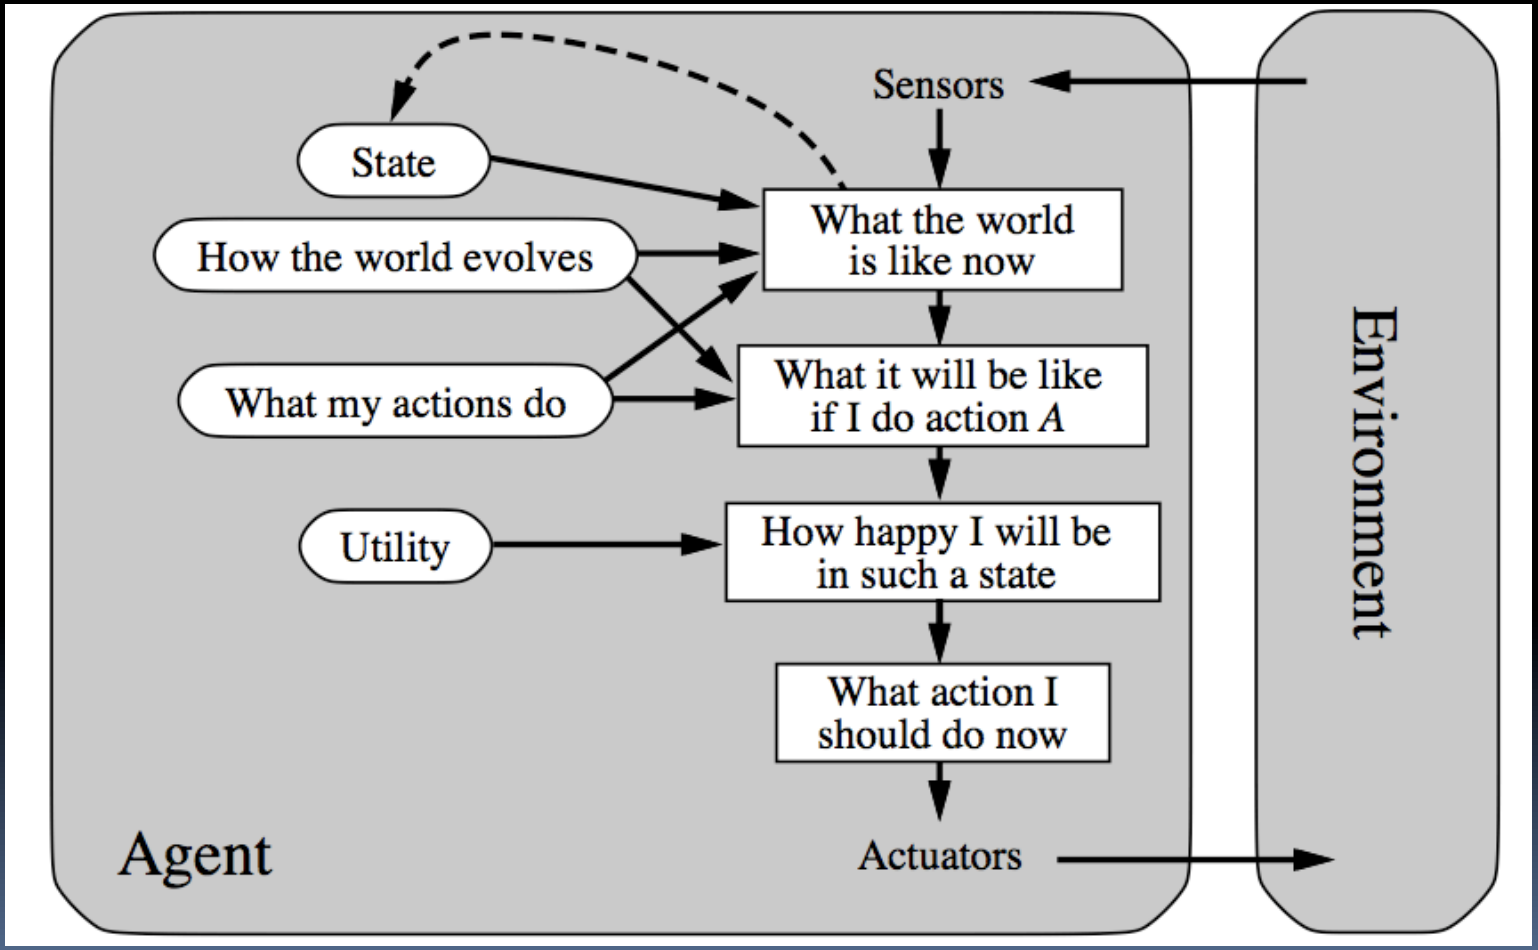

Learning Agents

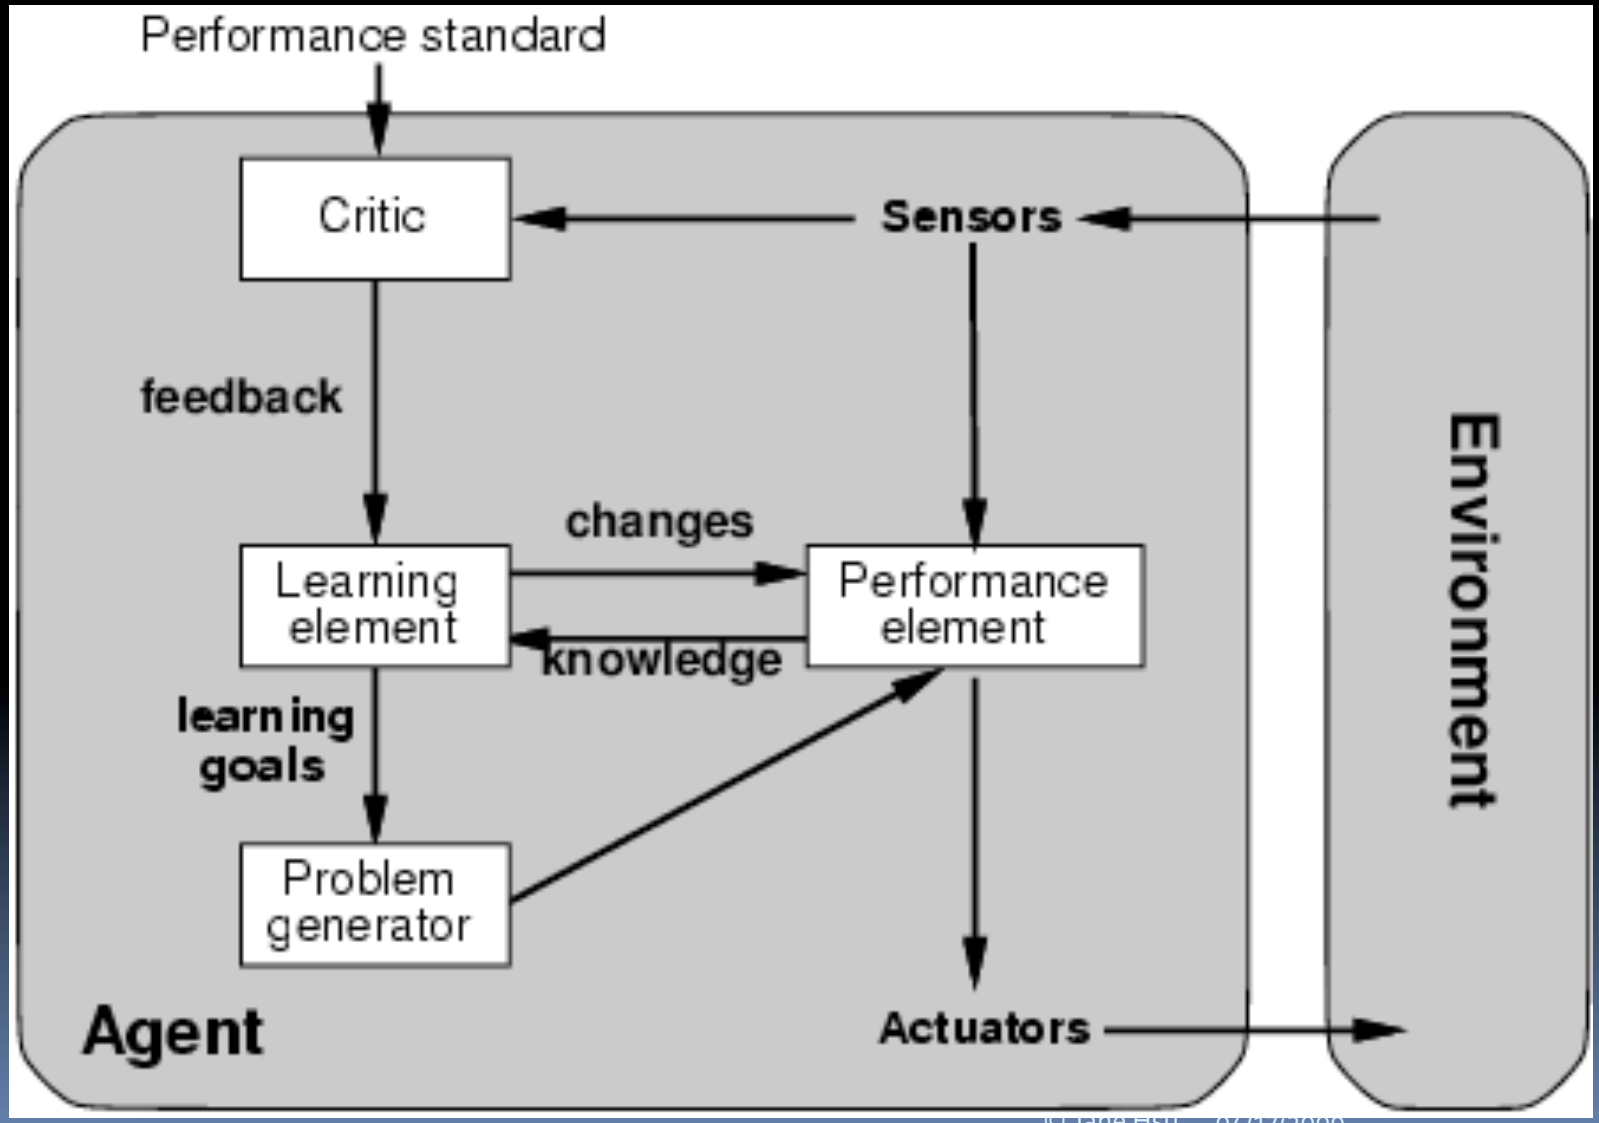

The Physical Symbol System Hypothesis

- *A physical symbol system has the necessary and sufficient means for intelligent action.*
  - "Computer Science as Empirical Inquiry: Symbols and Search" by Allen Newell & Herbert A. Simon, 1975 ACM Turing Award Lecture

AI: The Pioneering Days

- **Shakey the Robot (1966-1972)**
  - the first mobile robot to reason about its actions.
  - Developed by SRI's Artificial Intelligence Center
- Hardware
  - TV camera
  - A triangulating range finder
  - Bump sensors, and was connected to DEC PDP-10 and PDP-15 computers via radio and video links.
- Software
  - perception, world-modeling, and acting.
  - Low-level action routines took care of simple moving, turning, and route planning.
  - Intermediate level actions strung the low level ones together in ways that robustly accomplished more complex tasks.
  - The highest level programs could make and execute plans to achieve goals given it by a user.
  - The system also generalized and saved these plans for possible future use.
- Shakey currently resides in the Computer History Museum in Mountain View, CA.
- In 2004, Shakey was selected for induction to the Robot Hall of Fame at Carnegie Mellon University.

Historical Achievements

- Proverb solves crossword puzzles better than most humans
- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- Stanley drove 132 miles to win the Grand Challenge.

PROBLEM SOLVING AS SEARCH

Crossword Puzzle

Placement Route

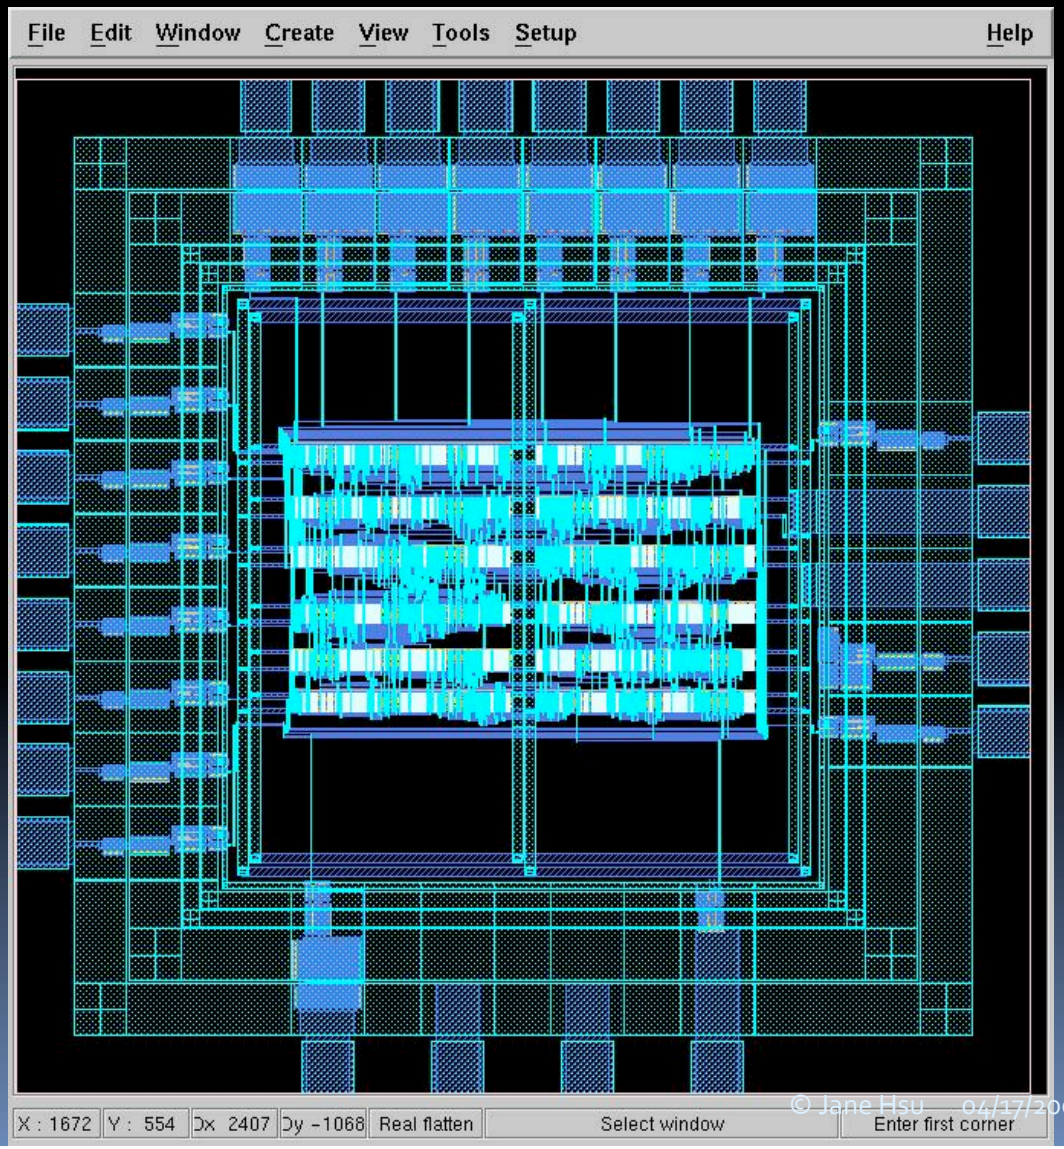

Romania with Step Costs

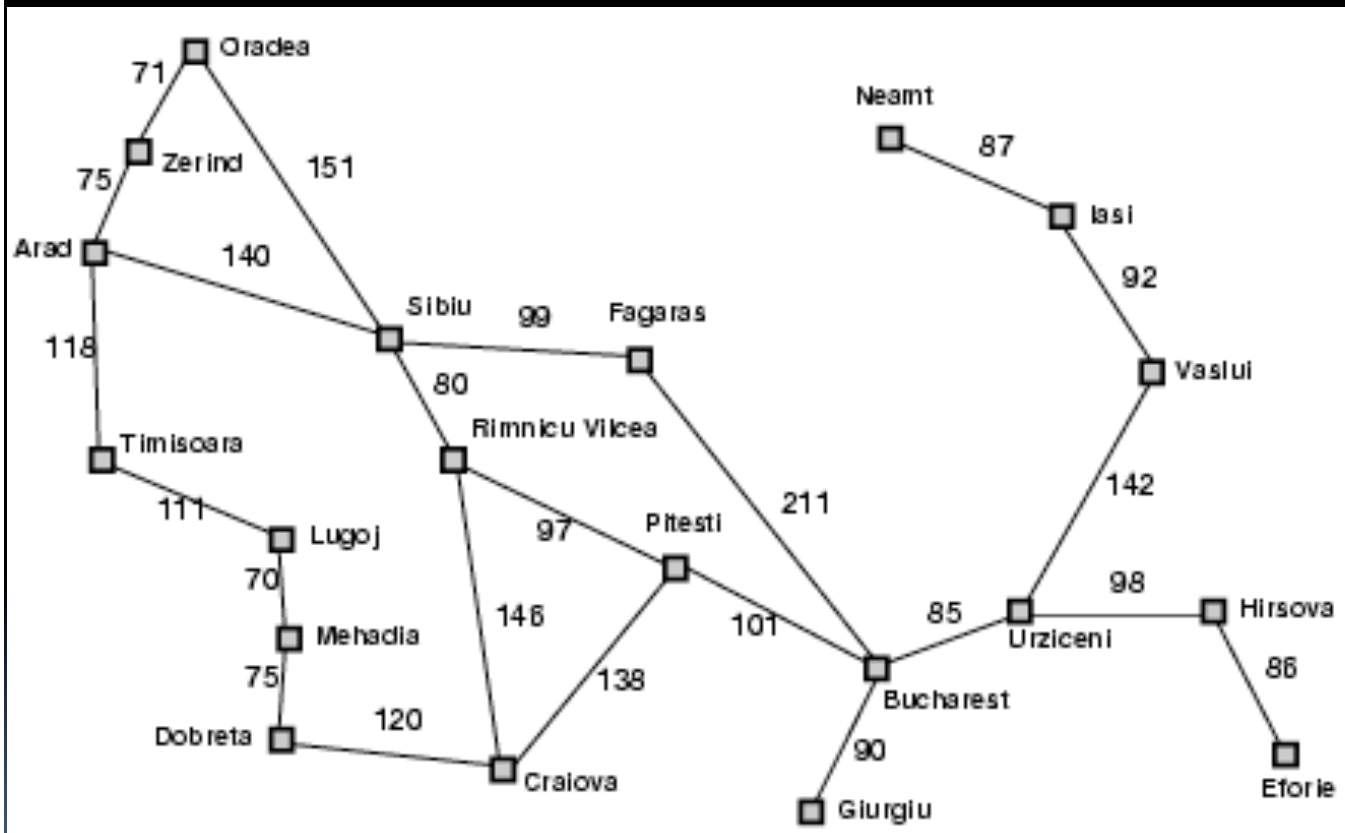

Straight–line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	176
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	10
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Informed Search Algorithms

- Heuristic Search
  - Greedy best-first search
  - A* search
- Local search algorithms
  - Hill-climbing search
  - Simulated annealing search
  - Local beam search
  - Genetic algorithms

A* Search

- Idea: avoid expanding paths that are already expensive
- Evaluation function $f(n) = g(n) + h(n)$
  - $g(n)$ = cost so far to reach n
  - $h(n)$ = estimated cost from n to goal
  - $f(n)$ = estimated total cost of path through n to goal

A* Search Example

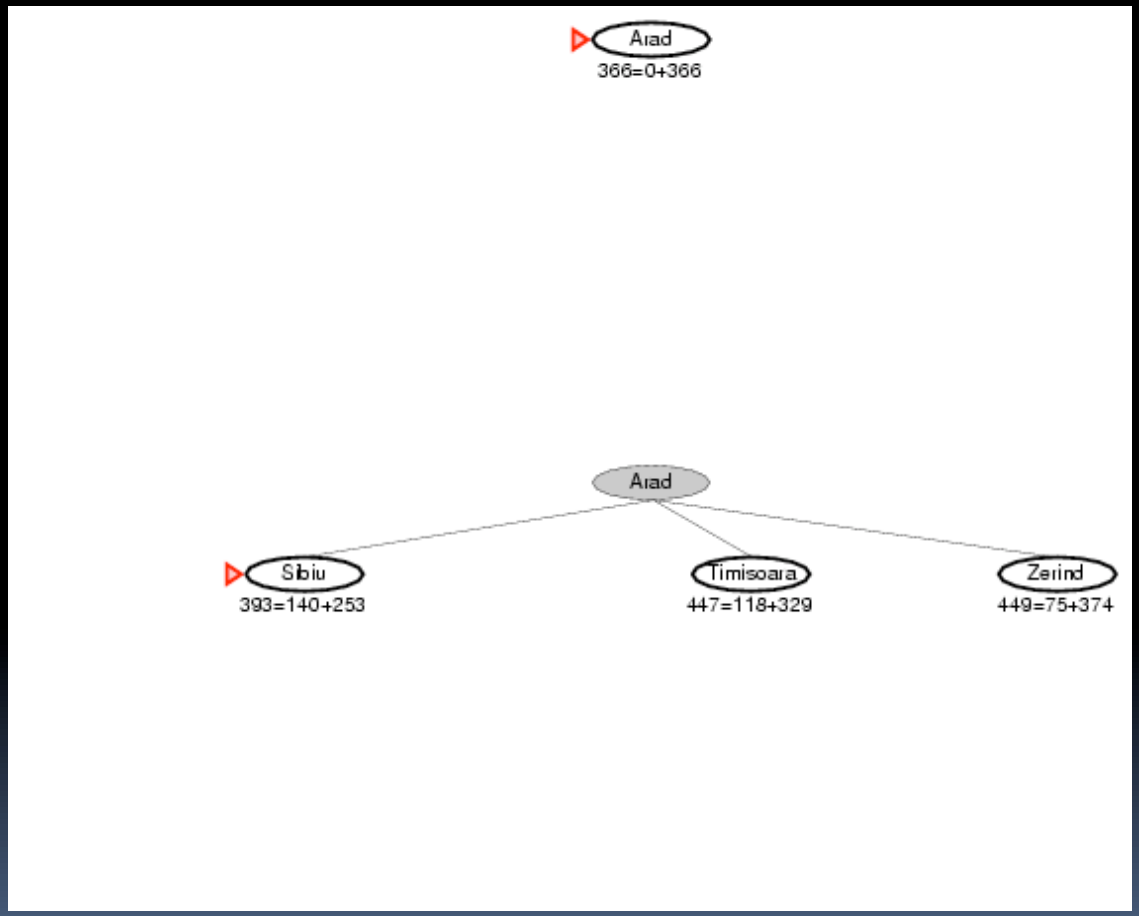

A* Search Example

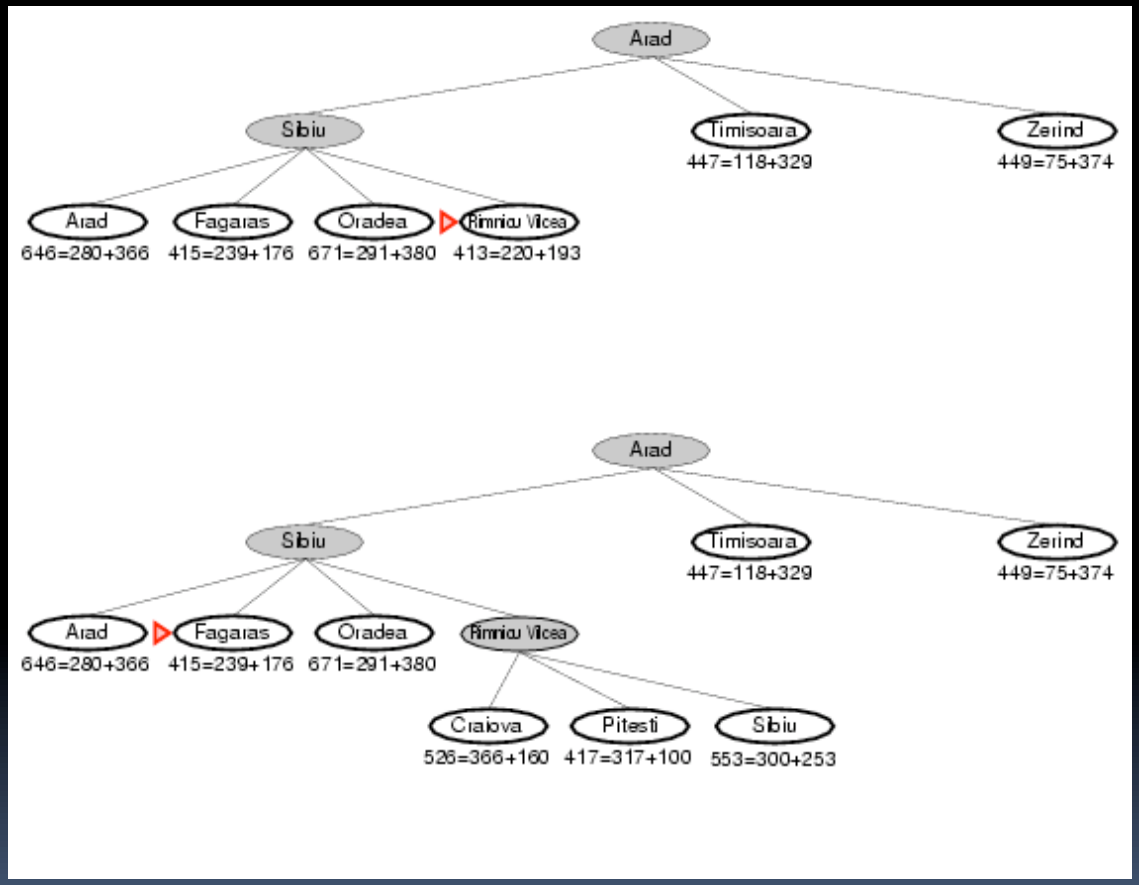

A* Search Example

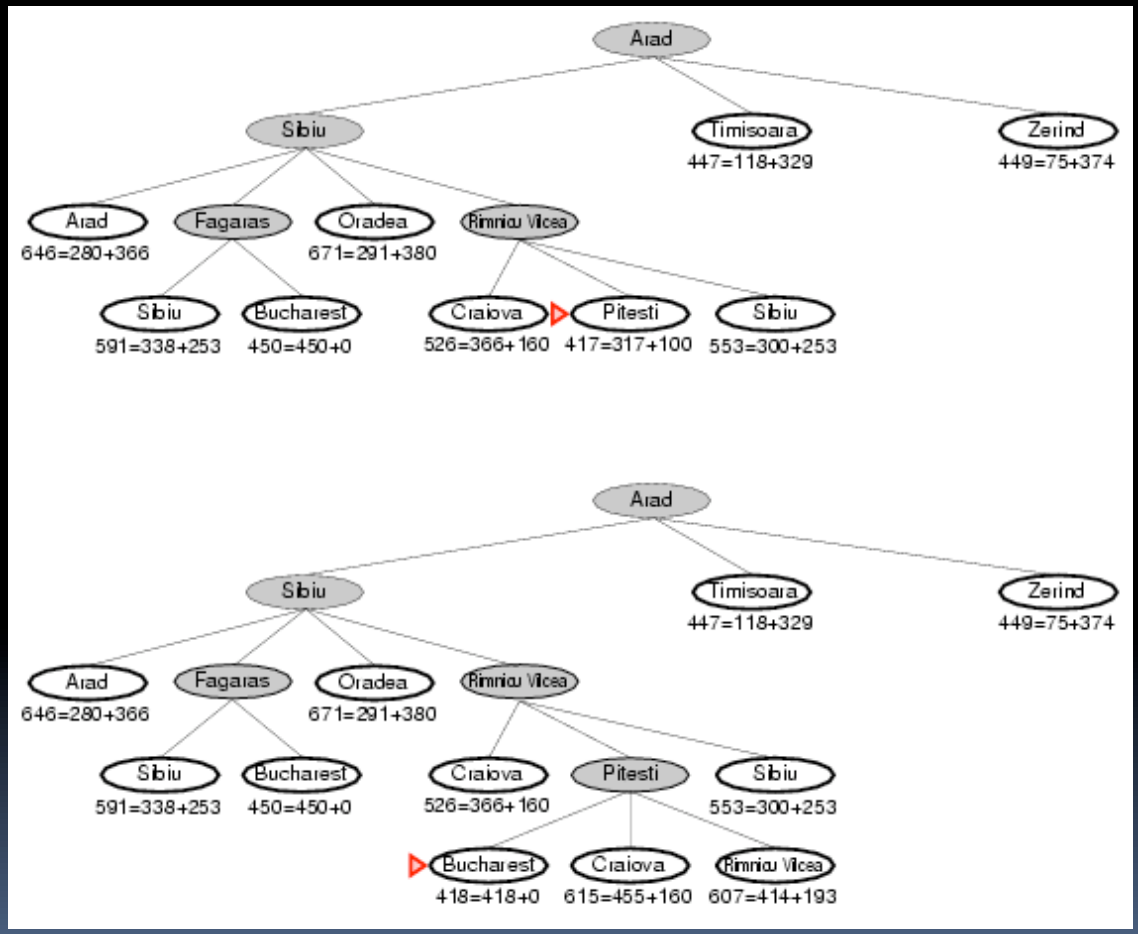

GAME-PLAYING COMPUTERS

Computer Chess

- Kasparov vs. Deep Blue
- The Match
  - May 3~11, 1997
  - Deep Blue won in 6 games

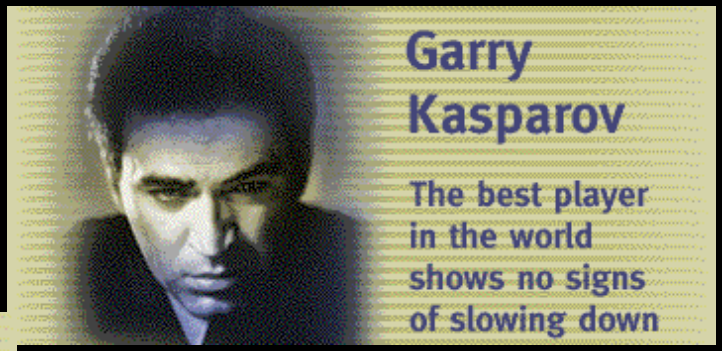

- Contrast in styles

End Game

Differences

1. Chess positions per second: up to 200,000,000 vs. 3
2. Amount of chess knowledge vs. amount of calculation ability
3. Sense of feeling and intuition
4. Guidance of five IBM research scientists and one international grandmaster vs. personal coach Yuri Dokhoian
5. Learning and adaptation

Differences (continued)

- Human frailties: forgetfulness, distraction, intimidation, fatigue, boredom and loss of concentration.
- Task-specific
- Changes by development team vs. self-modification
- evaluating its opponent's weaknesses
- Exhaustive vs. selective search through the possible positions

No Triumph for AI

- In 1957, the AI pioneer Herbert Simon predicted that a machine would be chess champion of the world within 10 years.
- He was off by three decades.
- More importantly, however, his prediction of how computers would solve chess proved to be entirely wrong — to artificial intelligence's enduring chagrin.

Games as Search Problems

- Games are idealization of worlds in which
  - the world state is fully accessible
  - the (small number of) actions are well-defined
  - uncertainty exists due to moves of the opponent, and the complexity of games
- A game can be defined as a search problem:
  - initial state
  - successor function (next moves or board situations)
  - terminal states
  - utility function (chance of win)

Game Tree

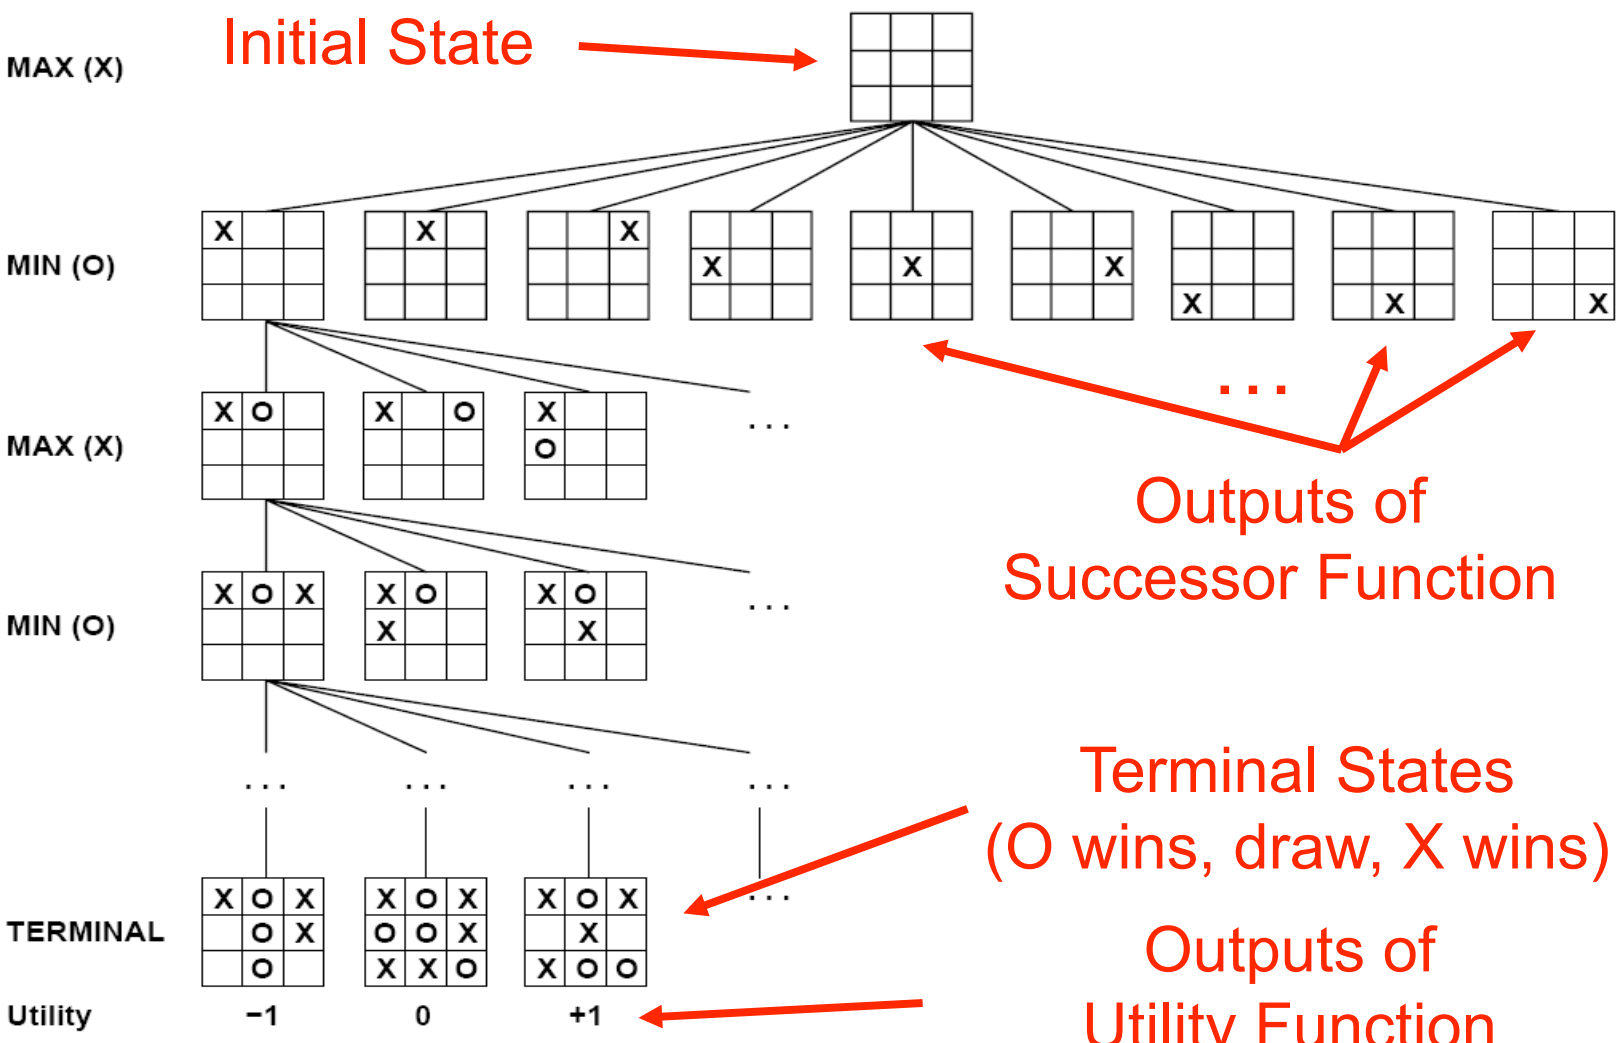
Initial State
MAX (X)
MIN (O)
MAX (X)
MIN (O)
TERMINAL
Utility
−1
0
+1
Outputs of
Successor Function
Terminal States
(O wins, draw, X wins)
Outputs of
Utility Function

When Game Tree Is Huge...

In a typical chess game, the game tree is huge such that exploration is limited within a given depth.

- Average branching factor: 35
- Average moves by a player: 50 (100 plies)
- Average size of a game tree: 35^{100}

  If a leaf node is not in a terminal state, an efficiently computable evaluation function is used to approximate its utility.

MiniMax Algorithm

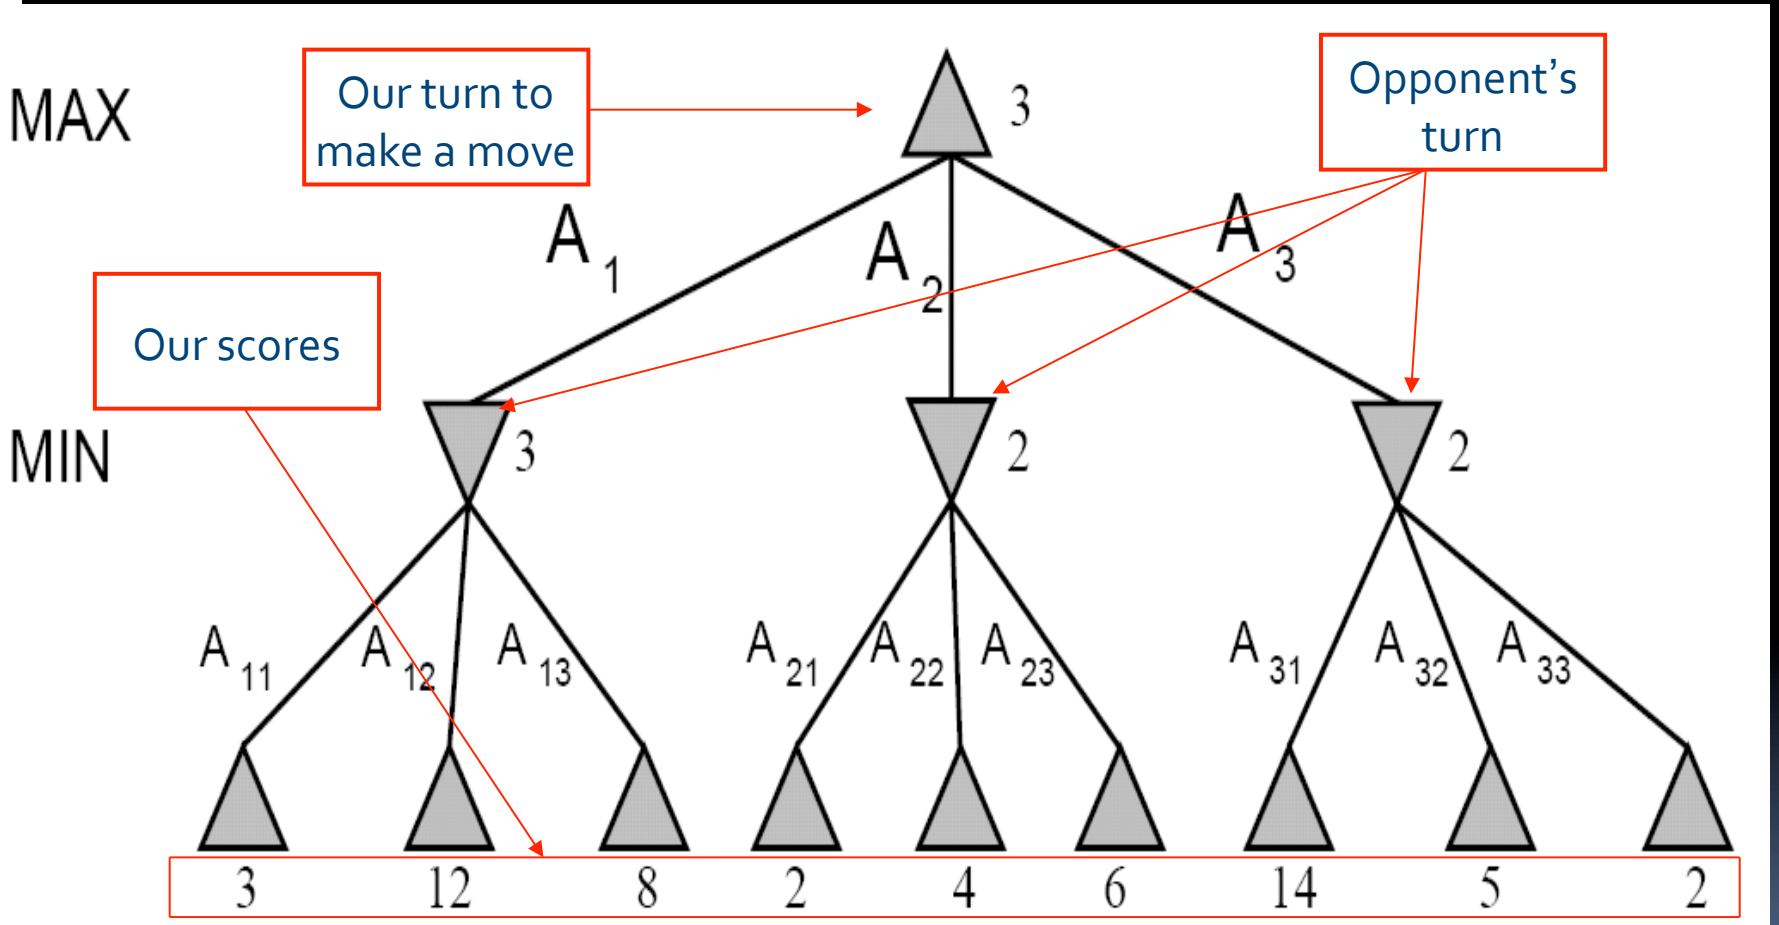

Opponent Minimizes Our Score

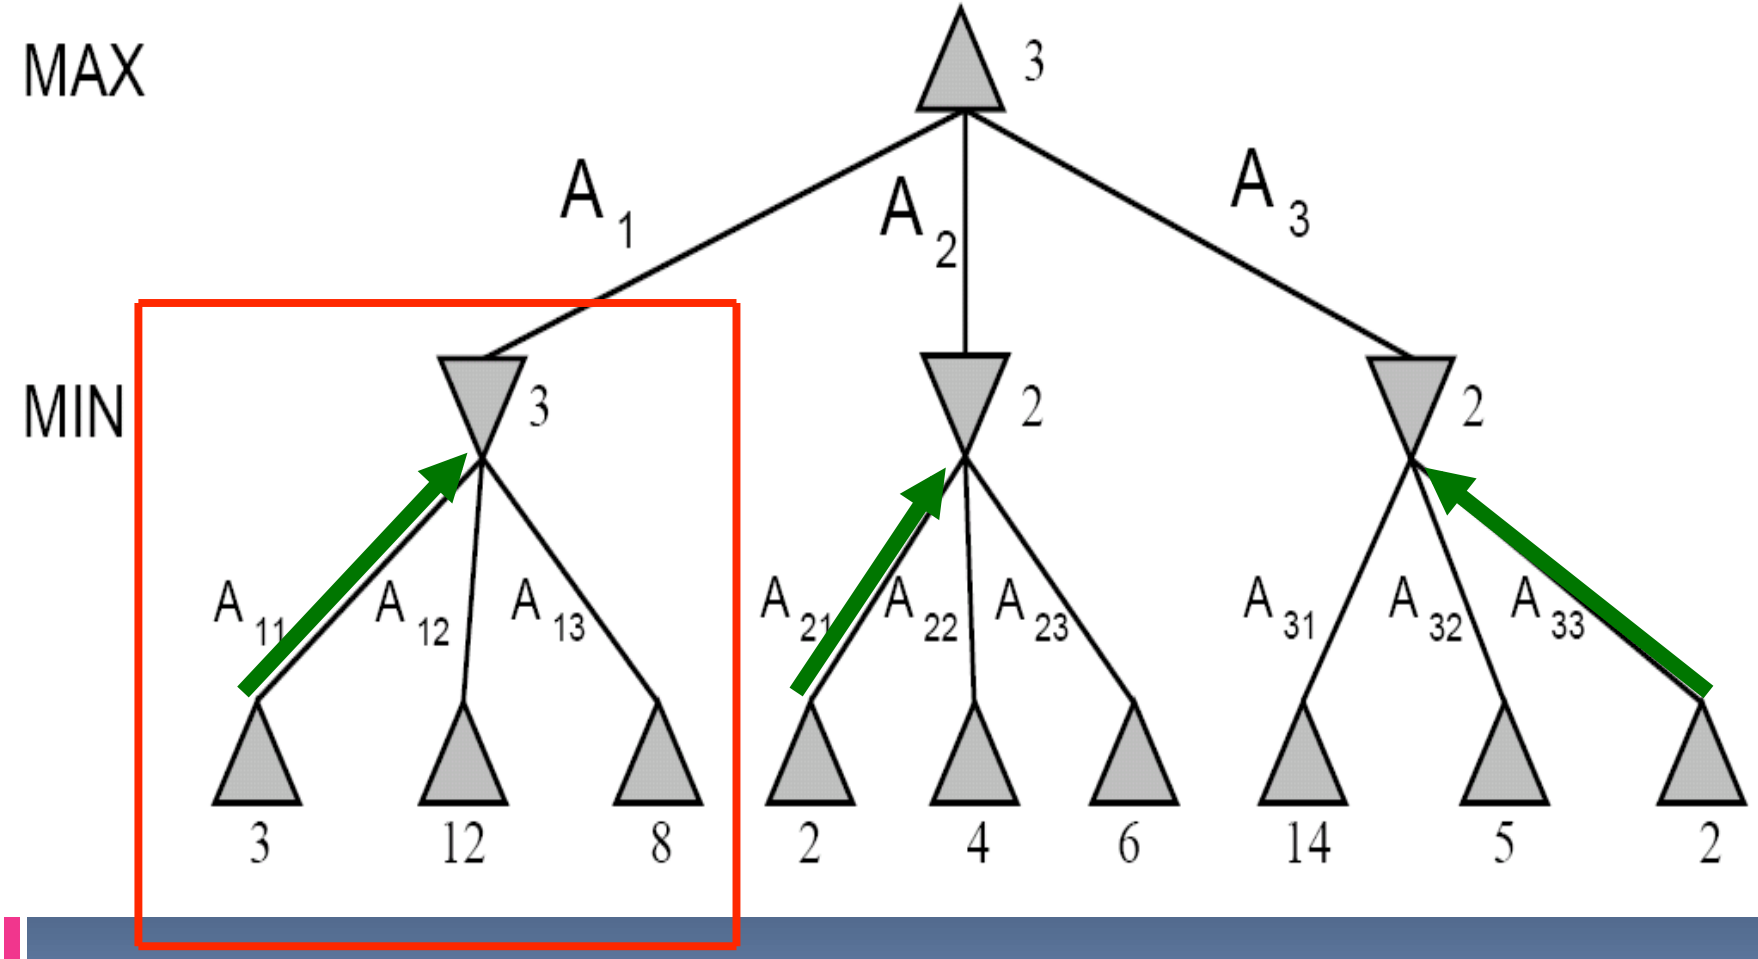

We Maximize Our Score

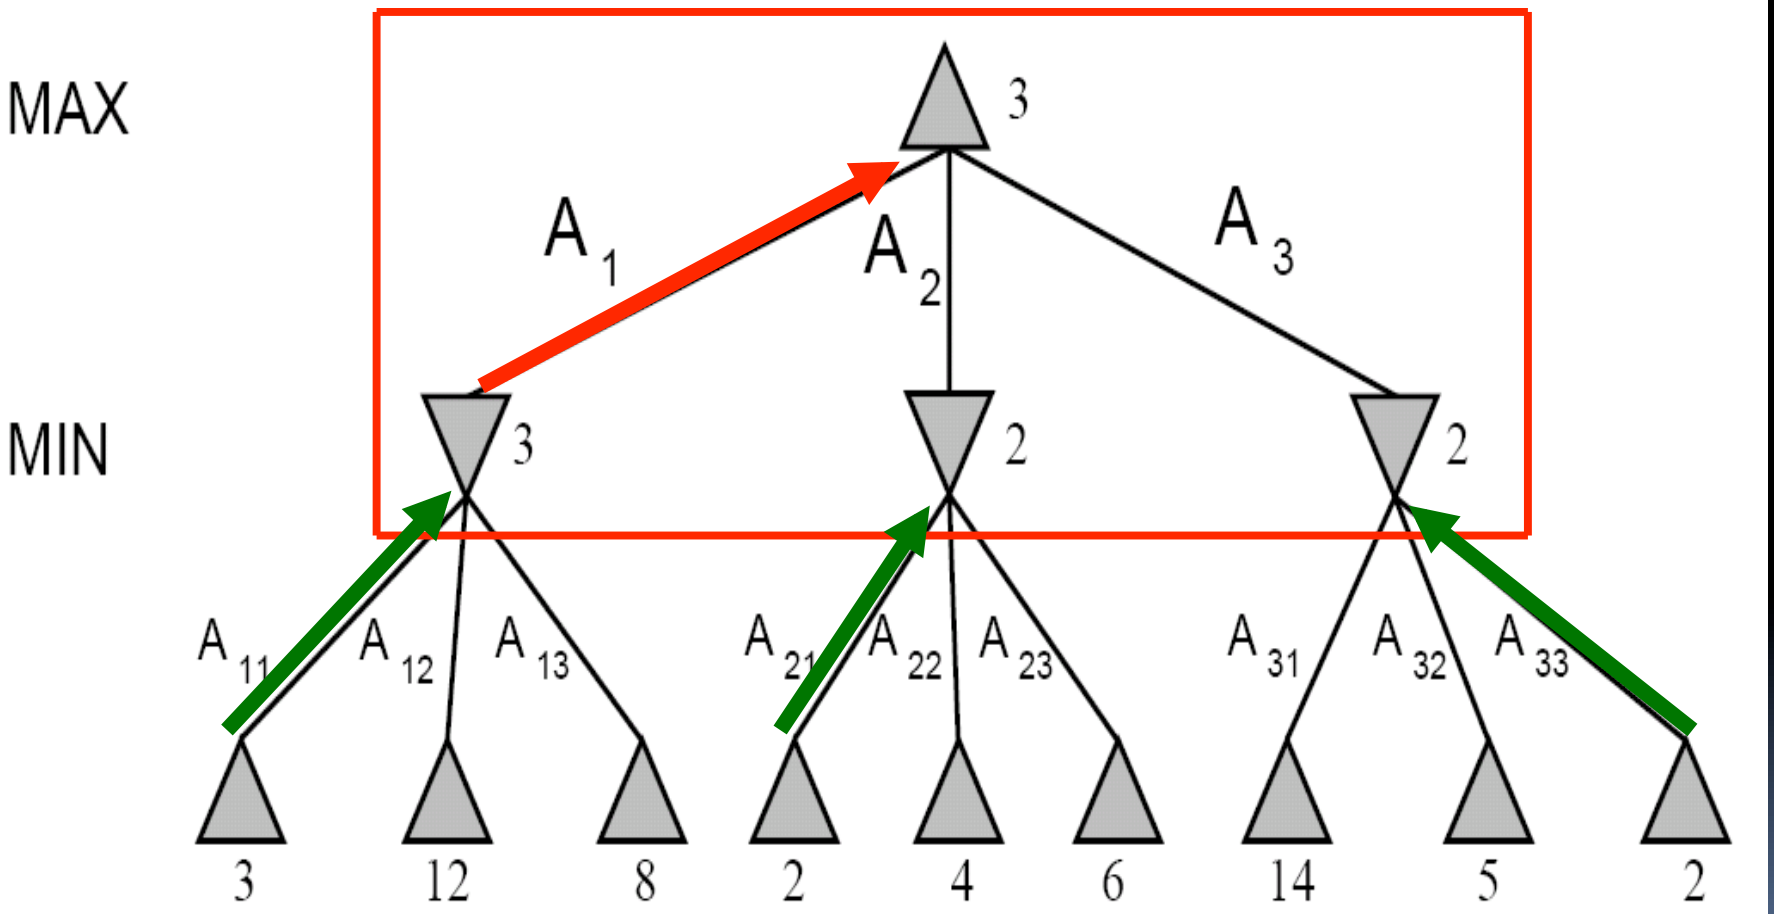

Minimax Algorithm

```
function MINIMAX-DECISION(state) returns an action
    inputs: state, current state in game
    return the a in ACTIONS(state) maximizing MIN-VALUE(RESULT(a, state))

function MAX-VALUE(state) returns a utility value
    if TERMINAL-TEST(state) then return UTILITY(state)
    v ← −∞
    for a, s in SUCCESSORS(state) do v ← MAX(v, MIN-VALUE(s))
    return v

function MIN-VALUE(state) returns a utility value
    if TERMINAL-TEST(state) then return UTILITY(state)
    v ← ∞
    for a, s in SUCCESSORS(state) do v ← MIN(v, MAX-VALUE(s))
    return v
```

Alpha-Beta Pruning

- Goal: to speed up MiniMax (MaxMax) algorithm, such that it
  - traverses fewer nodes in a game tree, and
  - returns a solution with the same score as MiniMax.

What kind of node (move) can be omitted?

Intuition: Good and Bad M

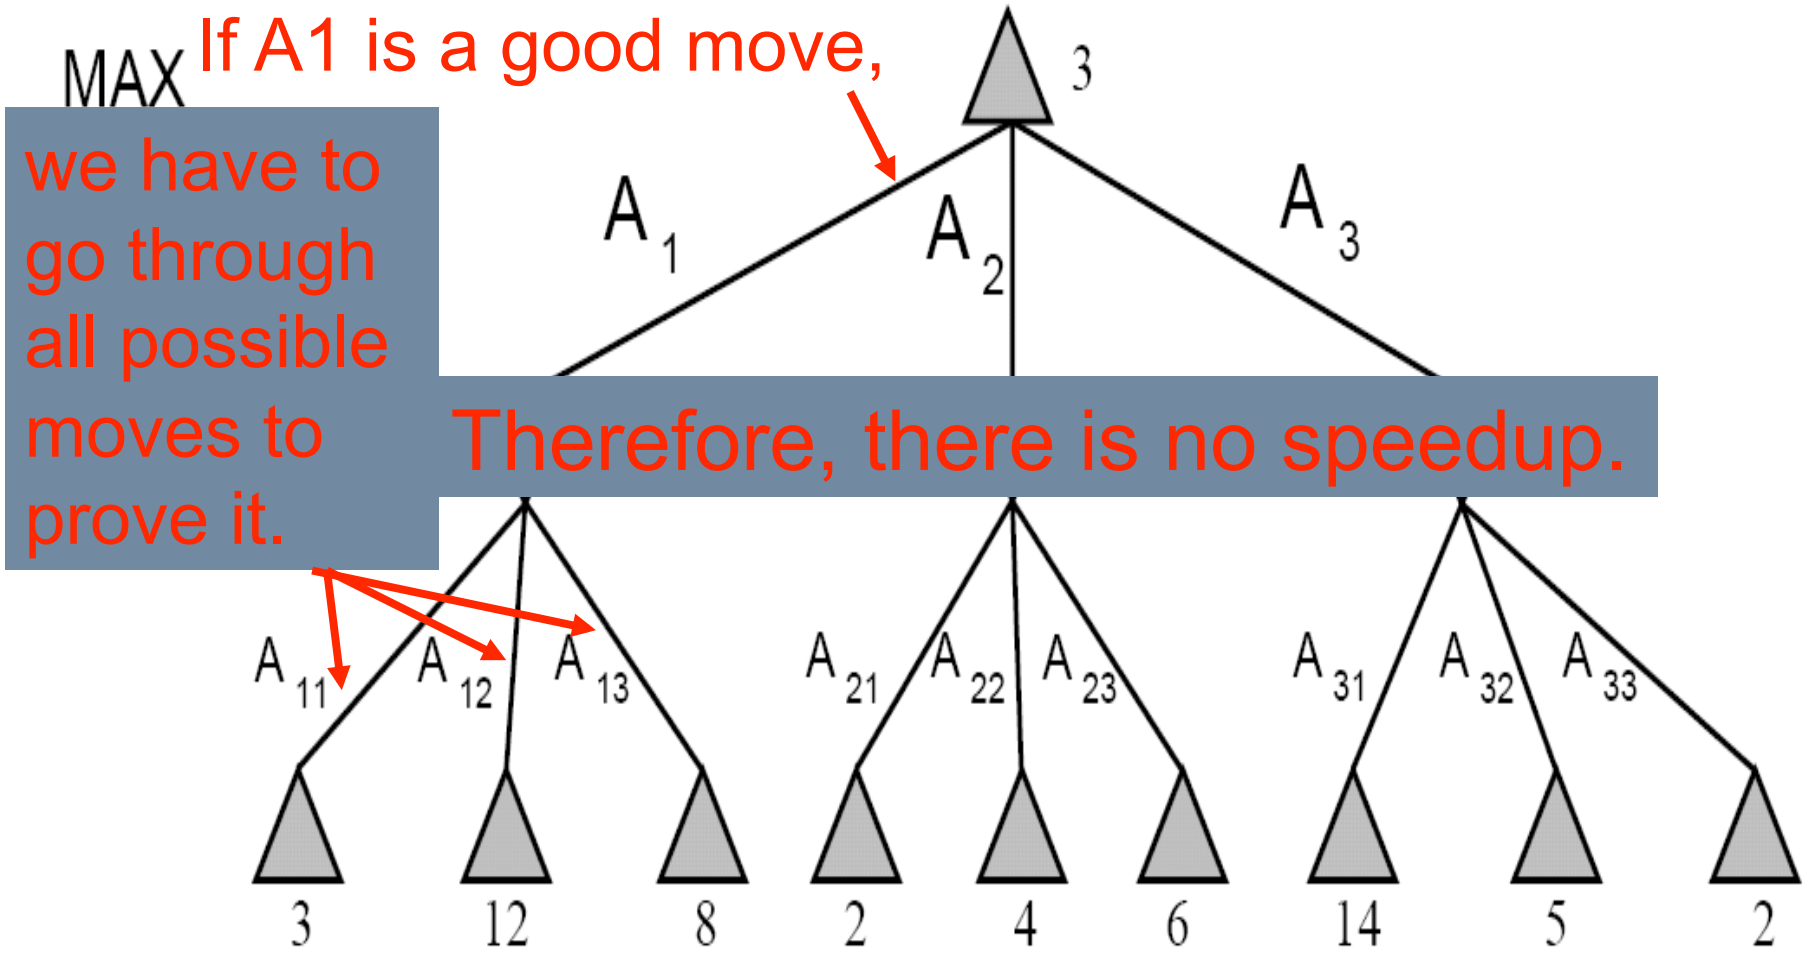

Intuition: Good and Bad M

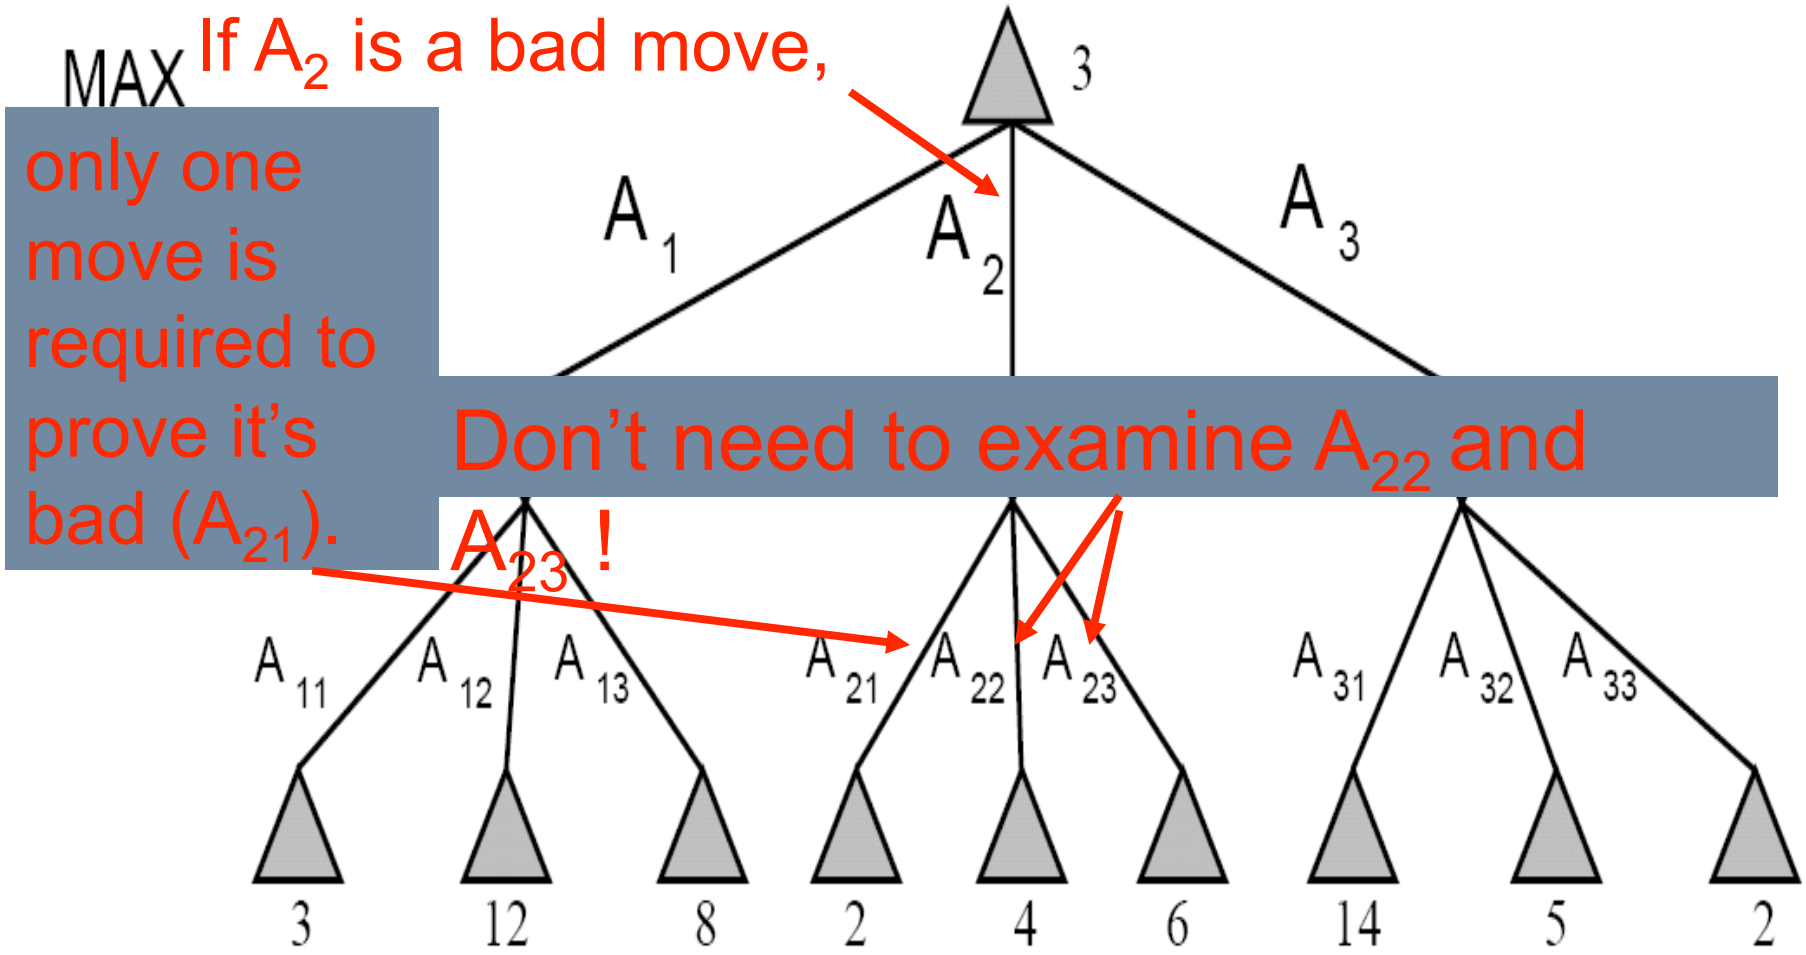

Pruning Example

1.Suppose that A_1 has been traversed.

2.A_2 is considered bad if its score <= 3.

3.That is, A_2 is considered bad if opponent's score >= -3.

4. A_{21} = -2 >= -3, which shows that A_2 is a bad move.

5. -3 is returned. (A_{22} and A_{23} are not traversed.)

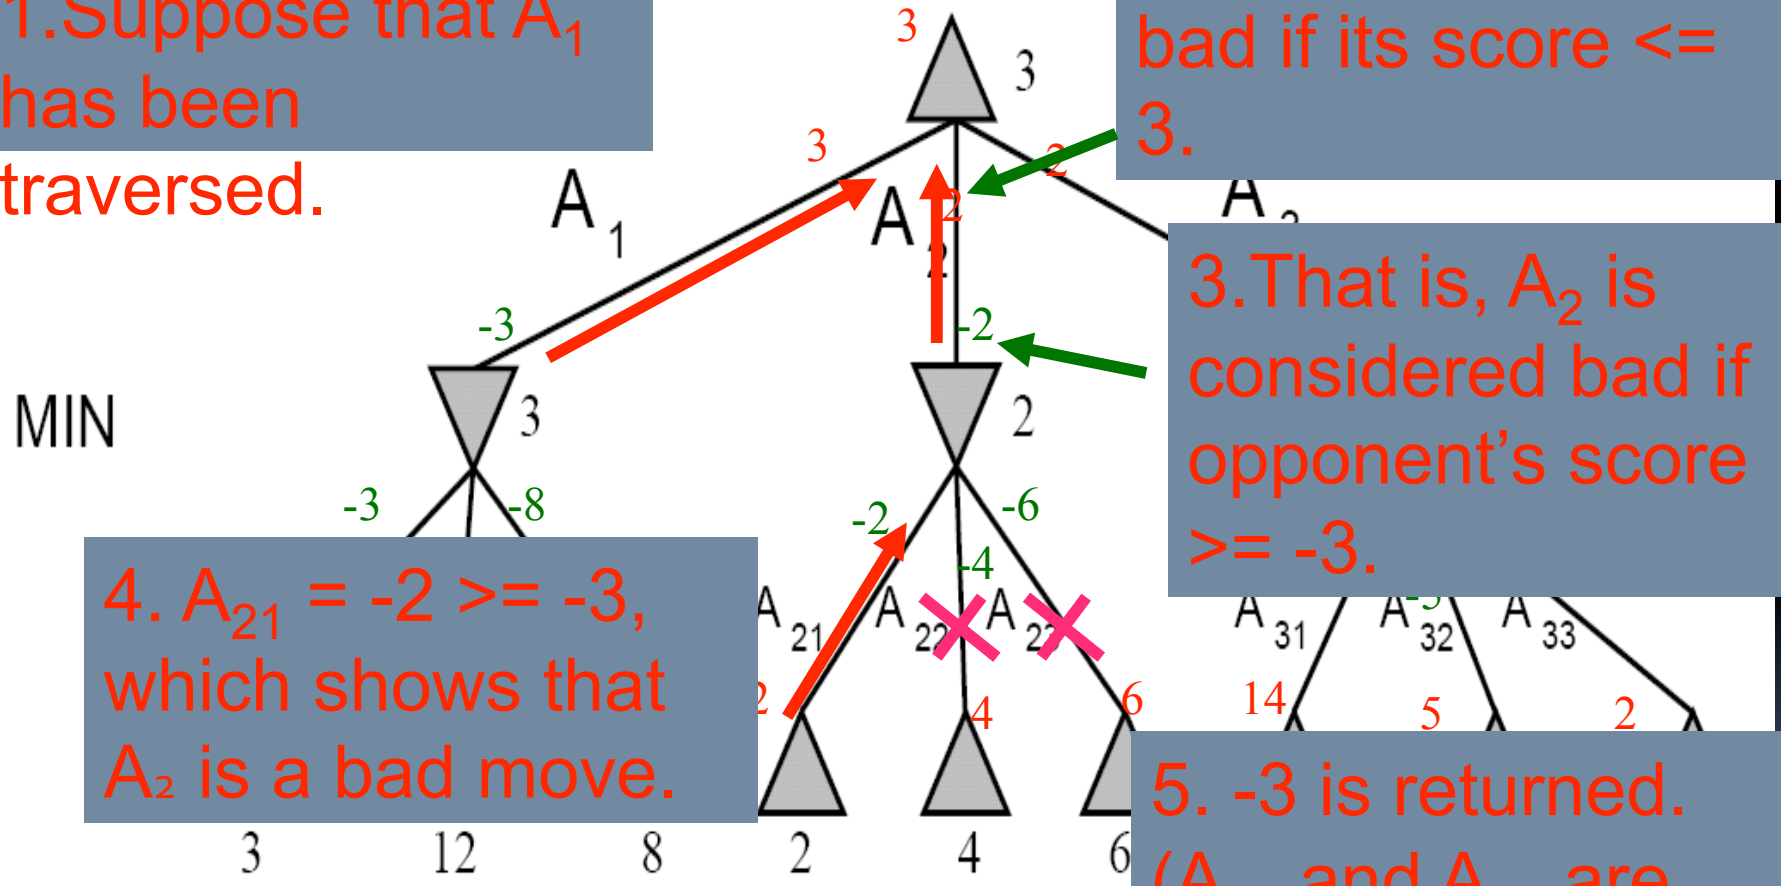

KNOWLEDGE REPRESENTATION AND ACQUISITION

Semantic Labeling Task

- Input: Collection of Photos
- Output: Photo metadata
- Requirement: High efficiency and accuracy

Face Recognition

- riya – a hybrid approach
  - Add – Train – Share
- Training data collection
  - Coverage
  - Correctness
- Challenges on illumination, pose, and incomplete information.

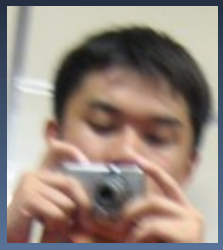

Human Computation

- Luis von Ahn, CMU
  - The ESP Game (2004)
  - Peekaboom (2006)

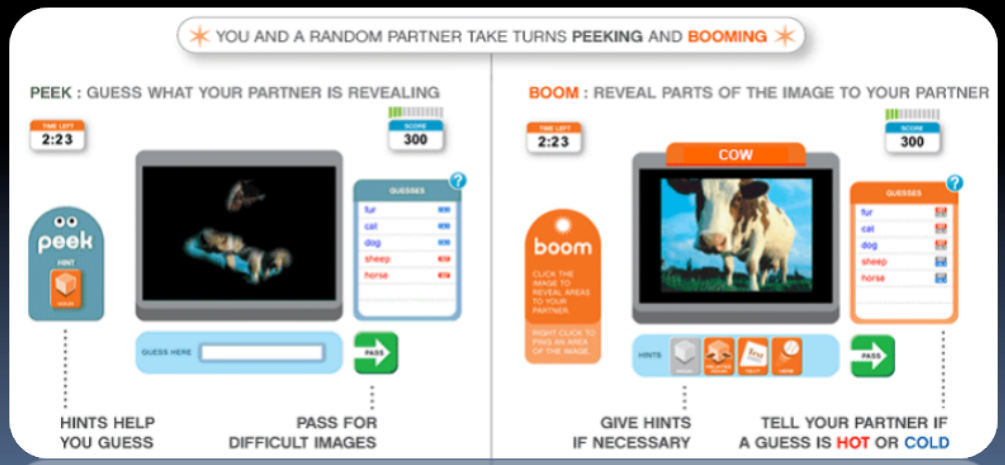

PhotoSlap: A Productive Game

- Motivation: to make it fun for the people involved
- Games as productivity tools
- People play games while producing useful information simultaneously

Photos

Multi-Player Online Game

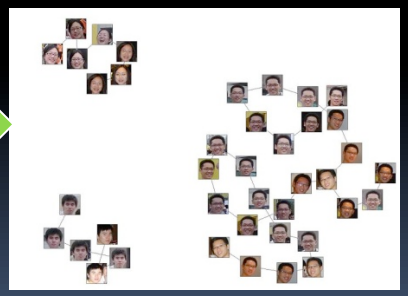

Clusters with
Semantic Relationship

Photoslap: Play to Annotate

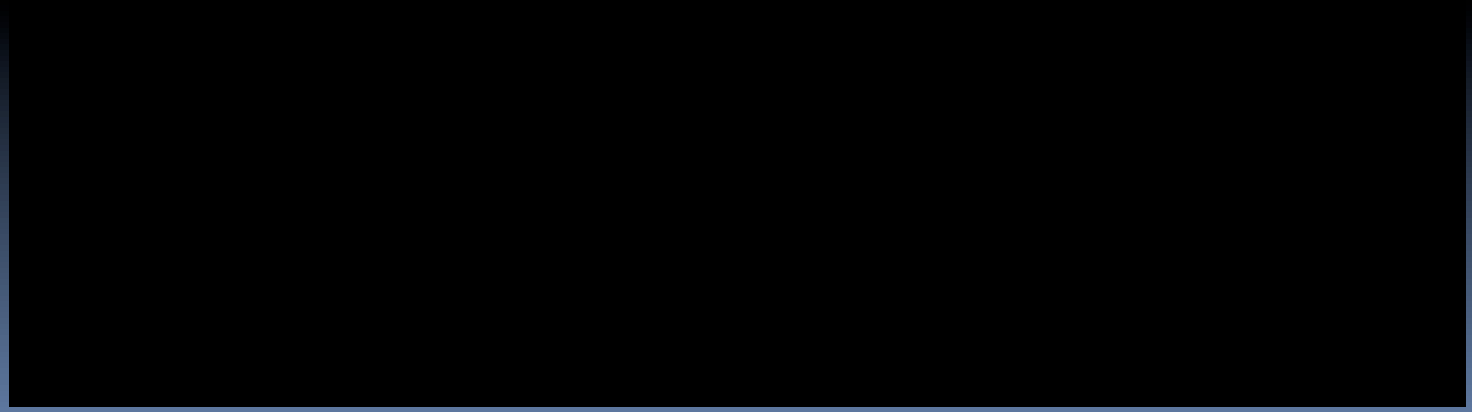

Landmark Tagging

- Data: keyword search from the web
- Partially labeled ground truth

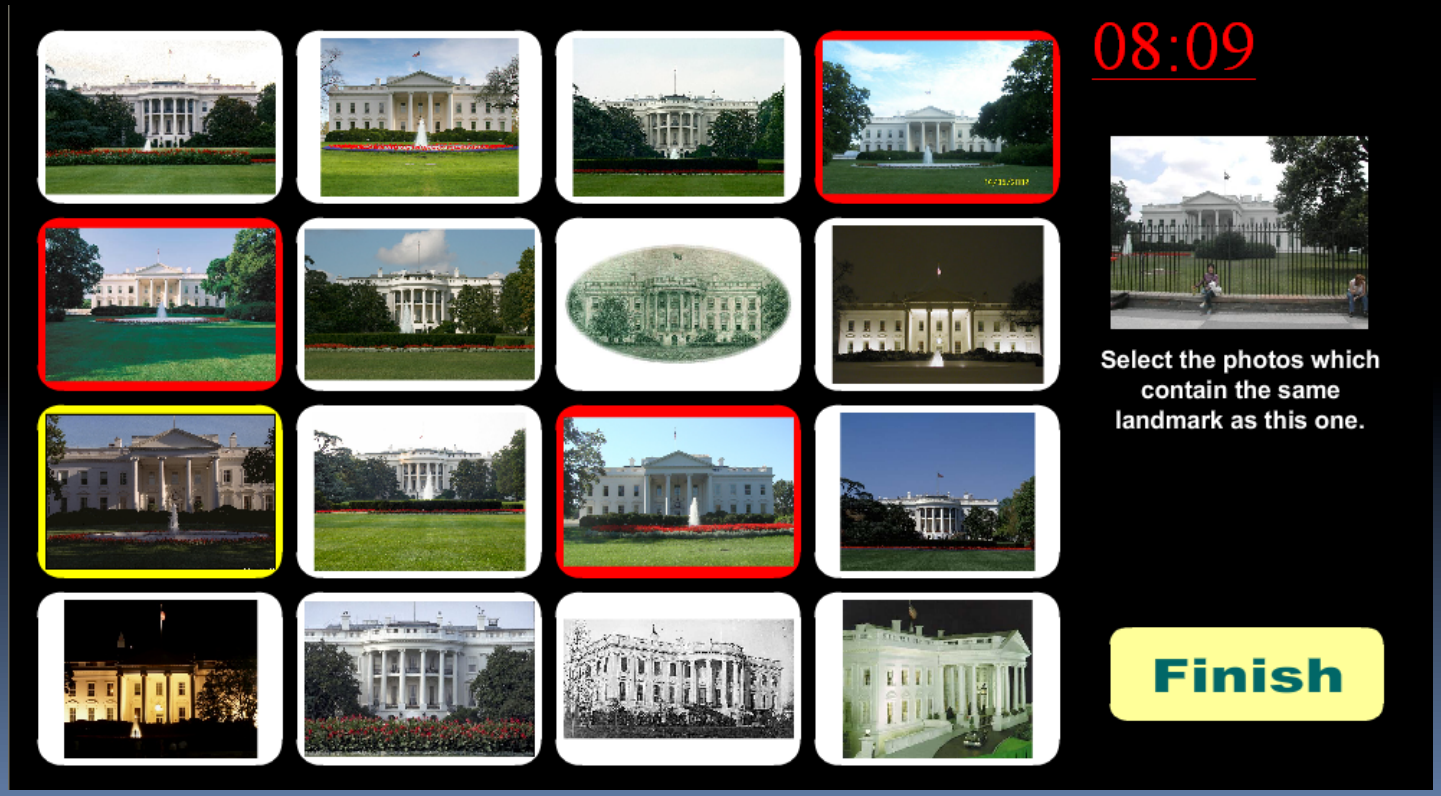

Open Mind Common Sense [MIT]

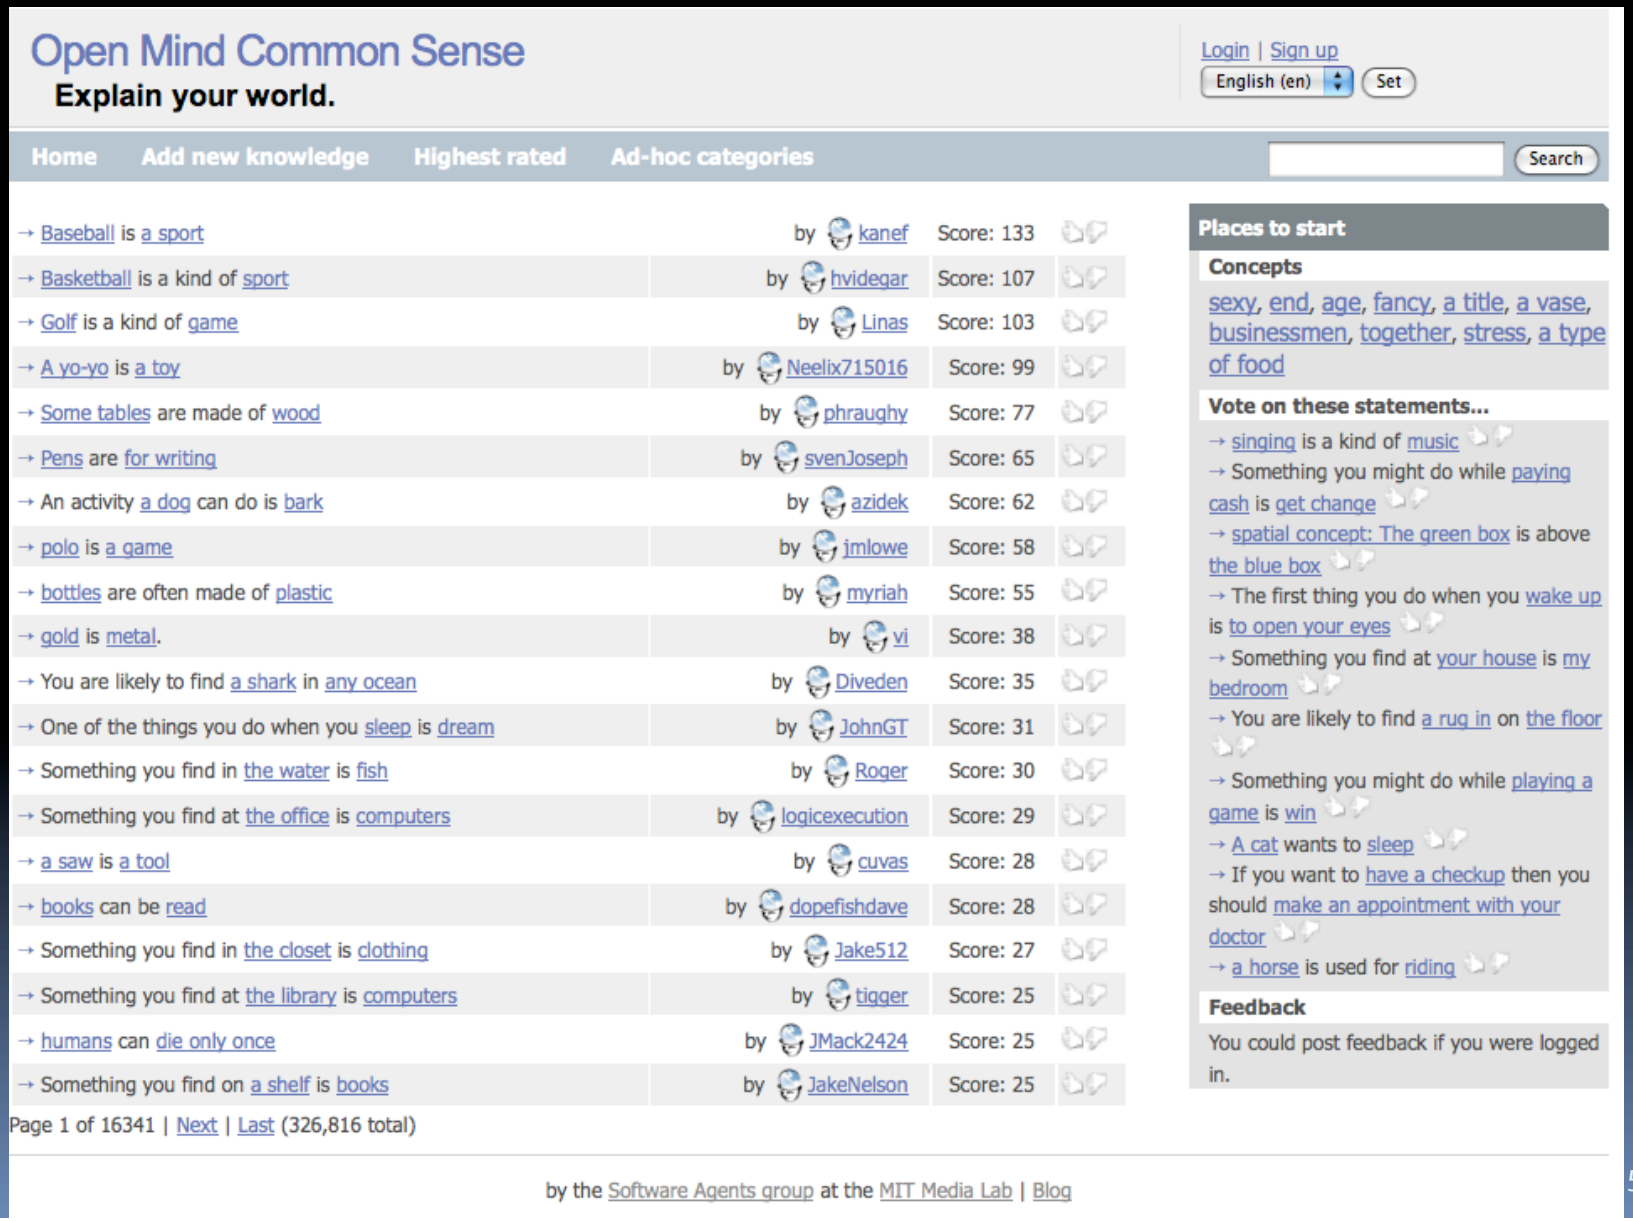

Open Mind Common Sense

Explain your world.

Login | Sign up

English (en) Set

Home | Add new knowledge | Highest rated | Ad-hoc categories | Search

→ Baseball is a sport	by kanef	Score: 133
→ Basketball is a kind of sport	by hvidegar	Score: 107
→ Golf is a kind of game	by Linas	Score: 103
→ A yo-yo is a toy	by Neelix715016	Score: 99
→ Some tables are made of wood	by phraughy	Score: 77
→ Pens are for writing	by svenJoseph	Score: 65
→ An activity a dog can do is bark	by azidek	Score: 62
→ polo is a game	by jmlowe	Score: 58
→ bottles are often made of plastic	by myriah	Score: 55
→ gold is metal.	by vi	Score: 38
→ You are likely to find a shark in any ocean	by Diveden	Score: 35
→ One of the things you do when you sleep is dream	by JohnGT	Score: 31
→ Something you find in the water is fish	by Roger	Score: 30
→ Something you find at the office is computers	by logicexecution	Score: 29
→ a saw is a tool	by cuvas	Score: 28
→ books can be read	by dopefishdave	Score: 28
→ Something you find in the closet is clothing	by Jake512	Score: 27
→ Something you find at the library is computers	by tigger	Score: 25
→ humans can die only once	by JMack2424	Score: 25
→ Something you find on a shelf is books	by JakeNelson	Score: 25

Page 1 of 16341 | Next | Last (326,816 total)

Places to start

Concepts

sexy, end, age, fancy, a title, a vase, businessmen, together, stress, a type of food

Vote on these statements...

→ singing is a kind of music

→ Something you might do while paying cash is get change

→ spatial concept: The green box is above the blue box

→ The first thing you do when you wake up is to open your eyes

→ Something you find at your house is my bedroom

→ You are likely to find a rug in on the floor

→ Something you might do while playing a game is win

→ A cat wants to sleep

→ If you want to have a checkup then you should make an appointment with your doctor

→ a horse is used for riding

Feedback

You could post feedback if you were logged in.

by the Software Agents group at the MIT Media Lab | Blog

ConceptNet Representation

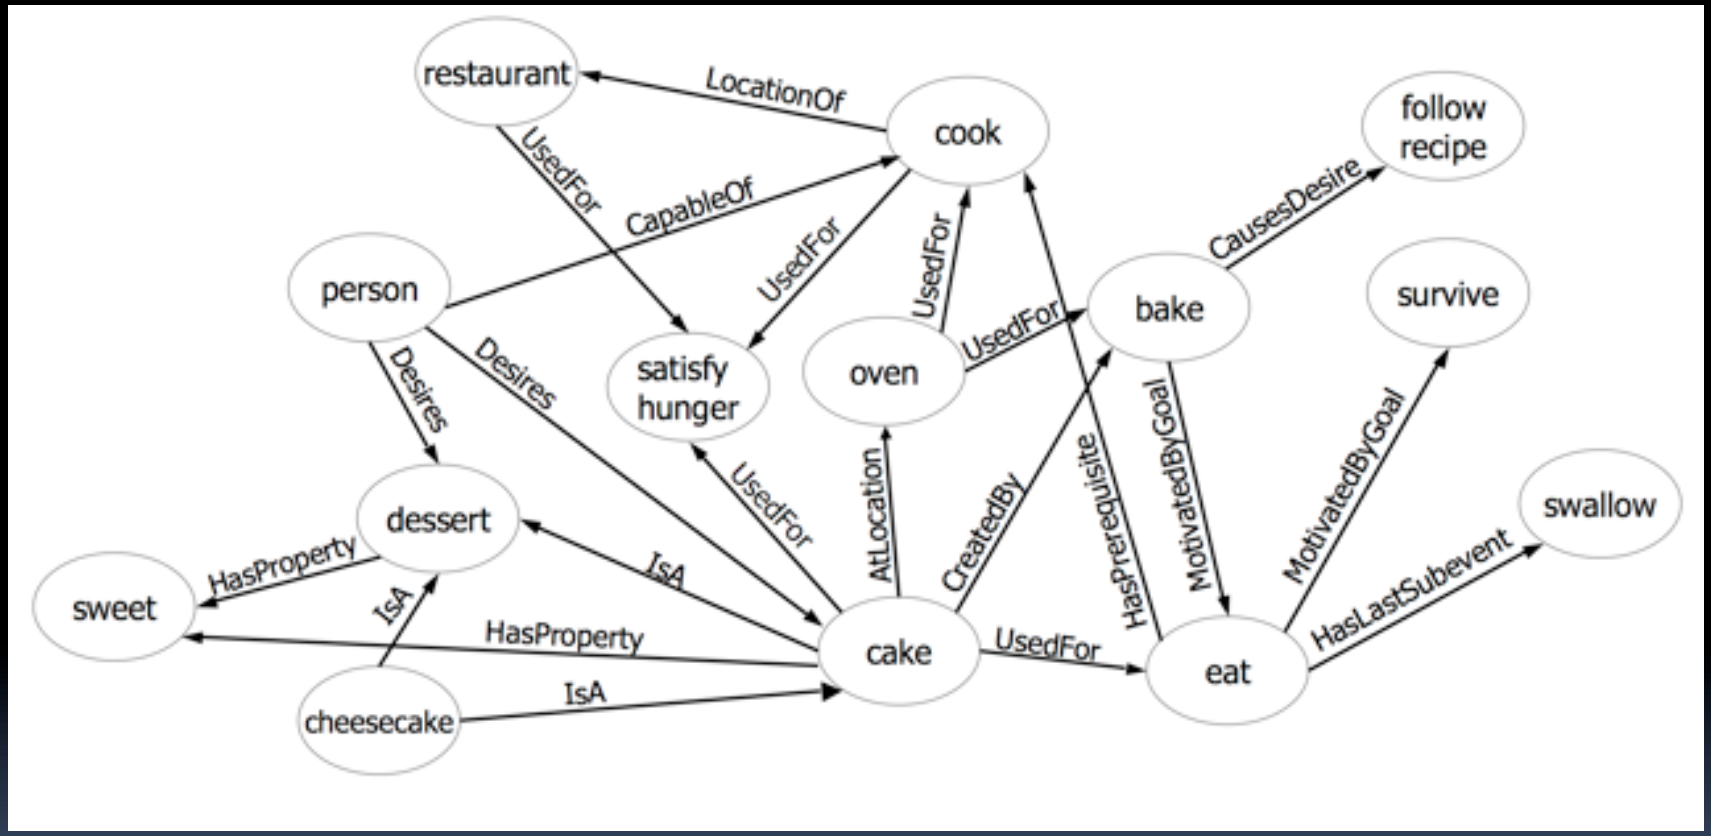

The Rapport Game on Facebook

1.

2.

Pet Game on PTT

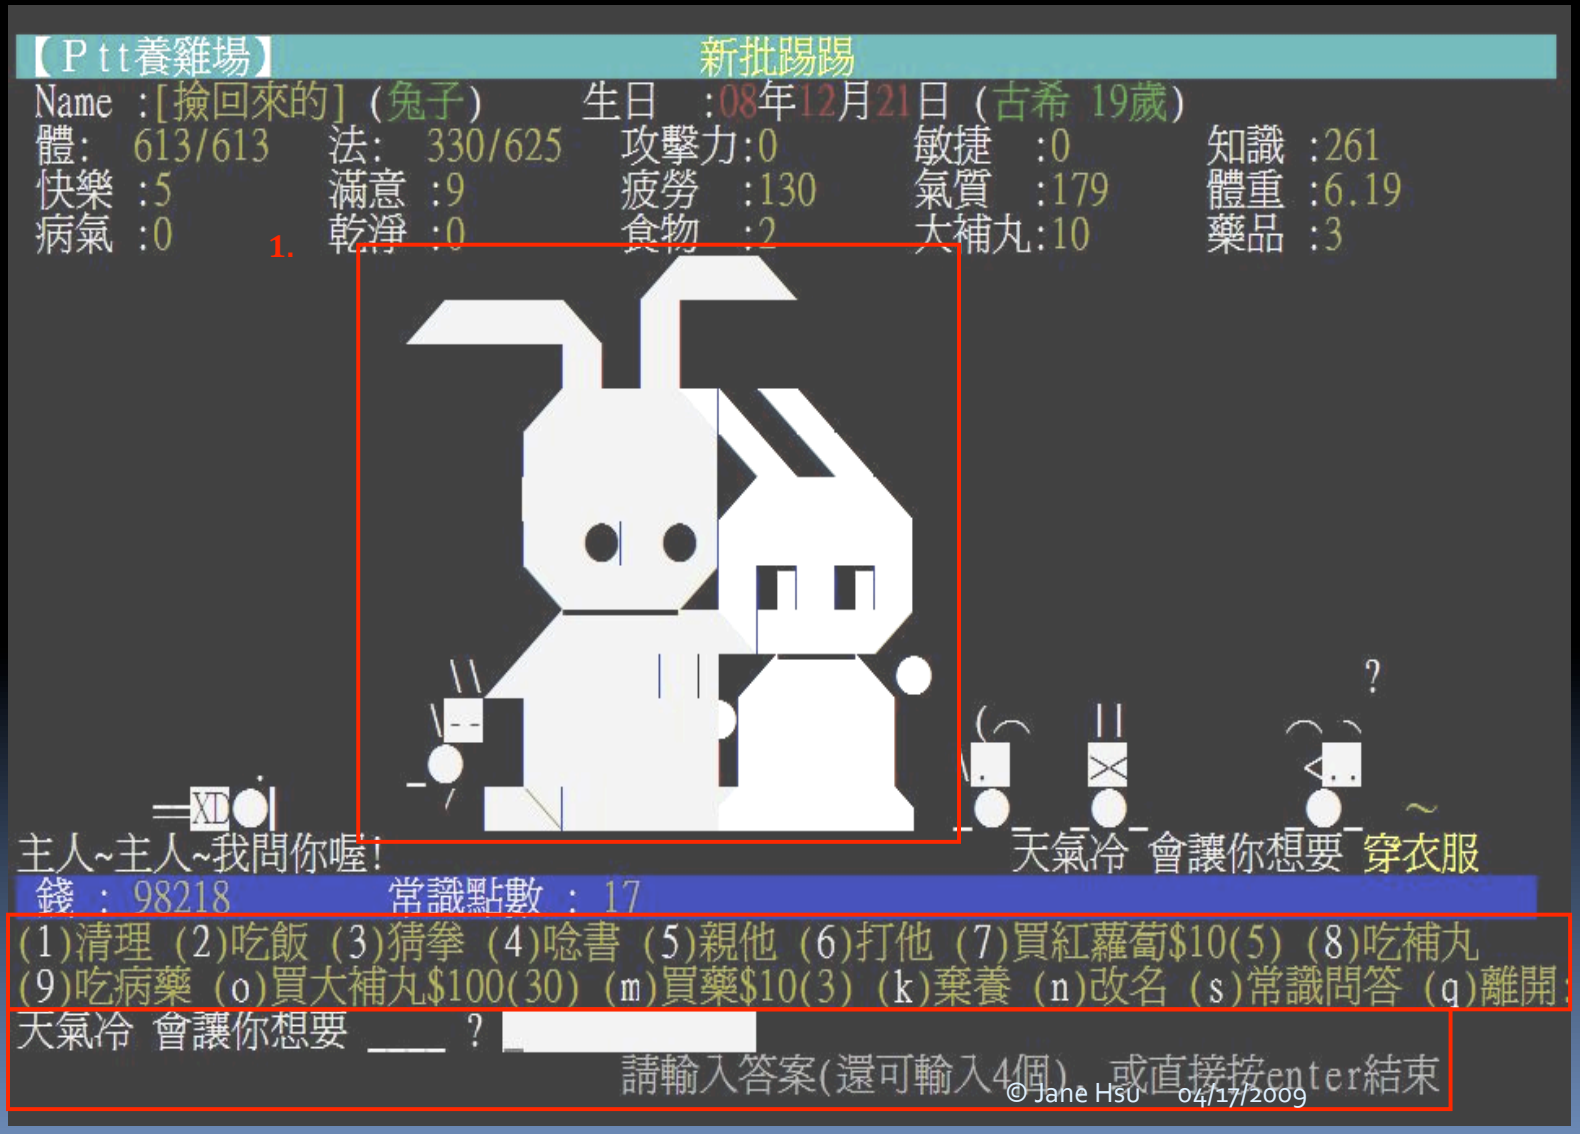

Commonsense Knowledge

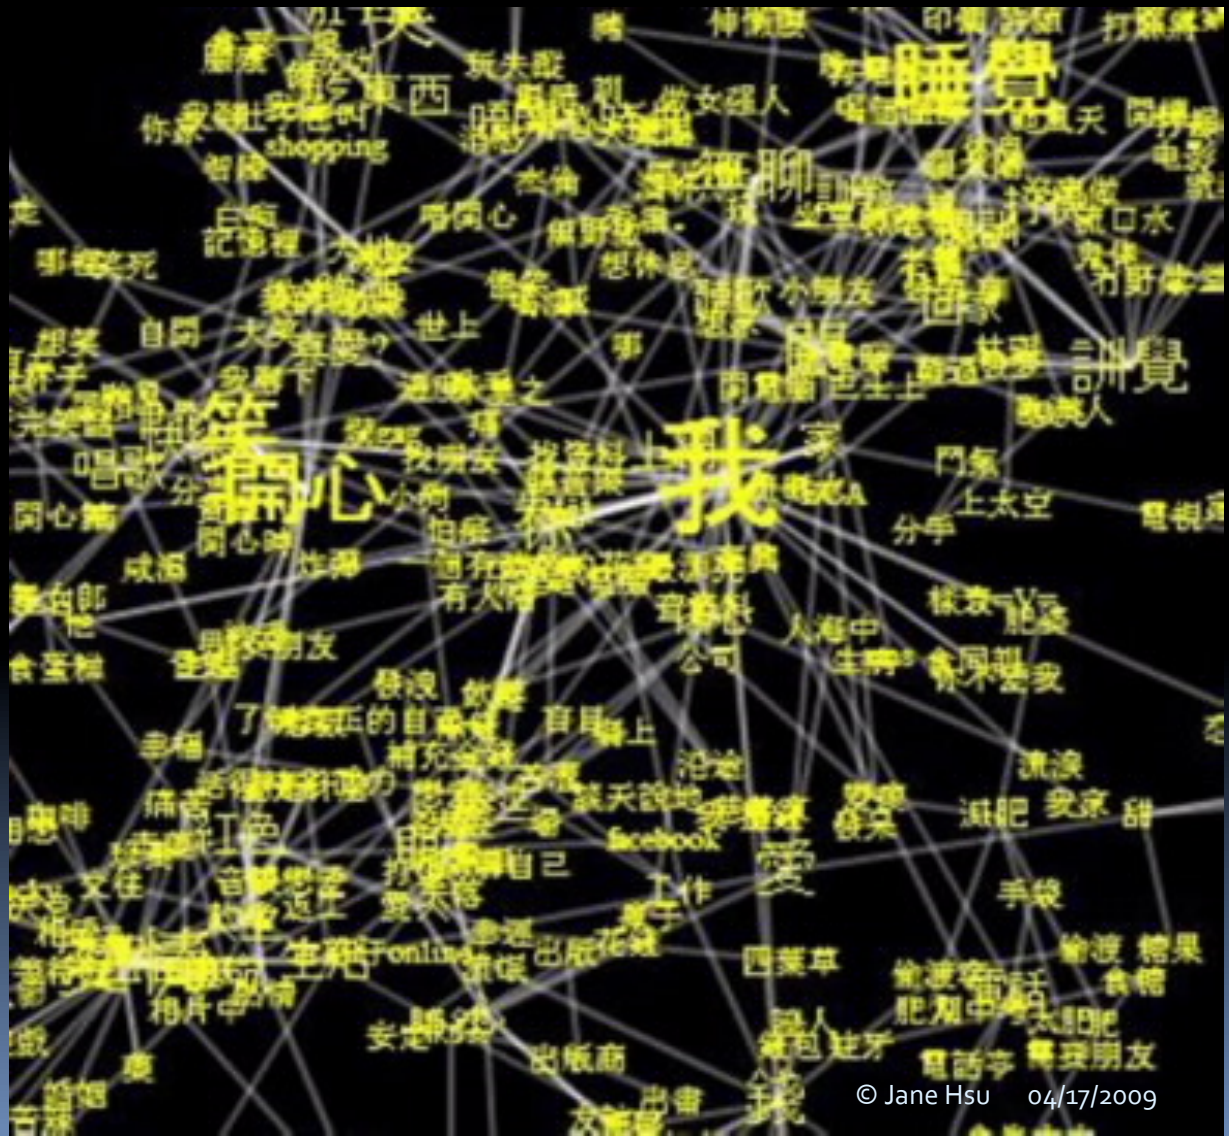

Coming Up

- Conversational Agents
  - Natural Language
  - Speech processing
- Vision
- Integrated AI
- Autonomous vehicles
- Emotional agents
- AI in art and other applications

HOMEWORK ASSIGNMENT

DUE: 05/01/2009

Connect-4

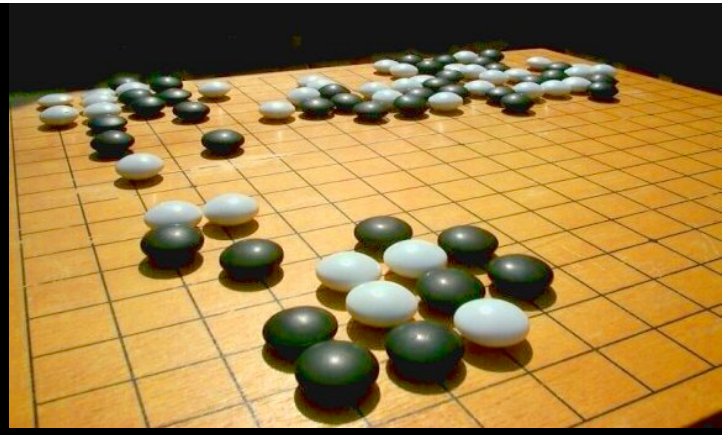

- Players and stones: There are two players. The first player, called Black here, holds a set of black stones, like Go or Go-Moku games. The second player, called White here, holds a set of white stones.
- Game boards: 9x9 Go boards.
- Game moves: Black plays first and puts only one black stone on one unoccupied intersection (or called grid). Subsequently, Black and White alternately put two of their own stones on two unoccupied grids.
- Game winning: The one who first gets four or more consecutive stones (horizontally, vertically or diagonally) of her/his own wins. When all squares on the board are placed without connecting four, the game draws.

Bonus: Commonsense Voting

- Due: 2009/4/17 23:59

1. 到 www.chickenschool.cc/vote.php
2. 選註冊新帳號:
   - 帳號: 你的學號, e.g. b94902000
   - 邀你做實驗的人: 選 ”前瞻資訊科技課程”
3. 開始投票. 請對每一句”常識”判斷是否合理. (若一時無法判斷, 可按”skip”跳過)
4. 投完票後將會顯示你的投票結果, 投滿30張票即可☺

Bonus: Semantic Labeling Game

http://daisy.csie.org/~leeaa/intro_page/kkd.html

- Due: 2009/4/18 5:00pm

KissKissDamn Introduction

This is a word-matching game. One pair of couple and one cursor play 15 rounds.

Cursor:

Make curses in 7 seconds at first to prevent the matching from couples.
Wait and see couples being cursed.
Win if couples have no matching.

Couple:

Avoid curses and match with another couple in 40 seconds.
If any word is curse, couples lost 5 seconds.
Win if any matching occurs in couples.

	Cursor	Couple
Cursor win	+200	-100
Couple win	-200	+100

START