CSEP 573
Applications of Artificial Intelligence (AI)

Rajesh Rao (Instructor)
Abe Friesen (TA)

http://www.cs.washington.edu/csep573

Our 2-course meal for this evening

- Part I
  - Goals
  - Logistics
  - What is AI?
  - Examples
  - Challenges

- Part II
  - Agents and environments
  - Rationality
  - PEAS specification
  - Environment types
  - Agent types

CSEP 573 Goals

- To introduce you to a set of key:

    Concepts & Techniques in AI

- Teach you to identify when & how to use

    Heuristic search for problem solving and games

    Logic for knowledge representation and reasoning

    Bayesian inference for reasoning under uncertainty

    Machine learning (for pretty much everything)

CSEP 573 Topics

- Agents & Environments
- Search
- Logic and Knowledge Representation
- Uncertainty and Bayesian Inference
- Machine Learning

CSEP 573 Logistics

- E-mail:

  Rajesh Rao rao@cs

  Abe Friesen afriesen@cs

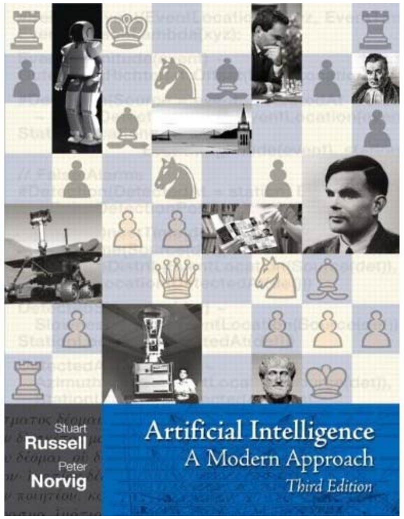

- Required Textbook

  Russell & Norvig's "AIMA3" (2009)

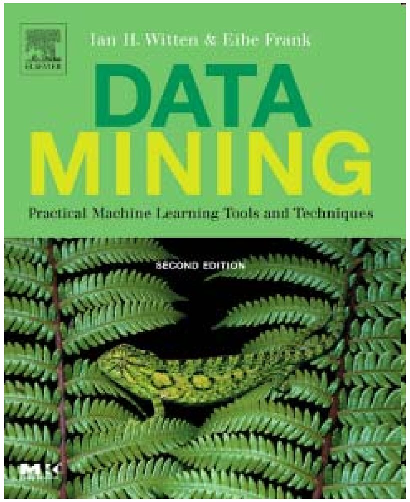

- Recommended Textbook

  Witten & Frank's "Data Mining" (2005)

CSEP 573 Logistics

- Grading:

  4 homework assignments, each 25% of course grade, containing a mix of written and programming problems

- Software tool:

  Some homeworks will use the data mining and machine learning software package Weka:

  http://www.cs.waikato.ac.nz/~ml/weka/index.html

  Documentation online and in the recommended textbook by Witten and Frank (see previous slide)

CSEP 573 Logistics

- 2 University Holidays:

  January 18 and February 15 – No class

- Make-up class:

  Thursday, February 18 6:30-9:20 pm

  Does this work for everyone?

Enough logistics,
let's begin!

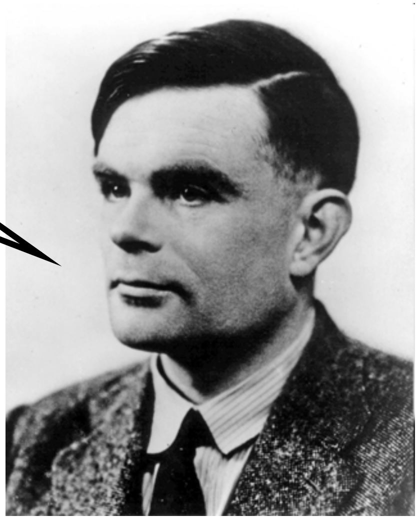

AI as Science

Physics: Where did the ***physical universe*** come from and what laws guide its dynamics?

Biology: How did ***biological life*** evolve and how do living organisms function?

AI: ?????

AI as Science

Physics: Where did the ***physical universe*** come from and what laws guide its dynamics?

Biology: How did ***biological life*** evolve and how do living organisms function?

AI: What is the nature of ***"intelligence"*** and what constitutes intelligent behavior?

AI as Engineering

- How can we make software and robotic devices more powerful, adaptive, and easier to use?
- Examples:
  - Speech recognition
  - Natural language understanding
  - Computer vision and image understanding
  - Intelligent user interfaces
  - Data mining
  - Mobile robots, softbots, humanoids
  - Medical expert systems…

Hardware

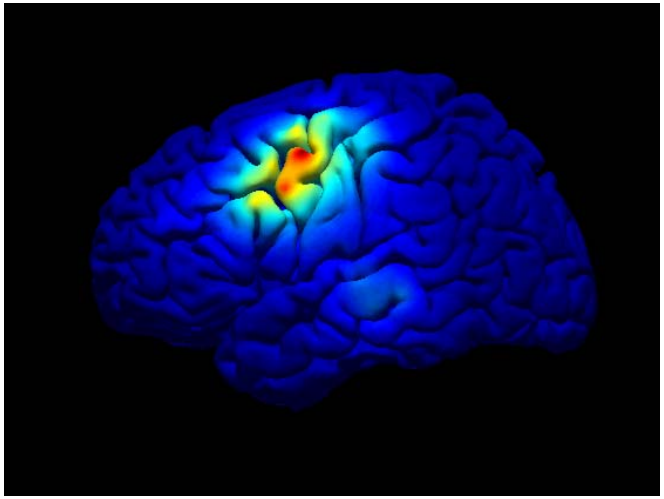

10^{11} neurons
10^{14} synapses
cycle time: 10^{-3} sec

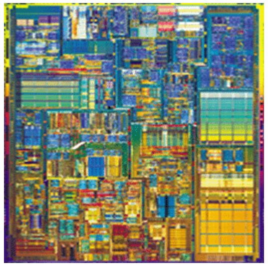

10^{9} transistors (4 CPUs)
10^{11} bits of RAM (12.5 GB)
cycle time: 10^{-9} sec

Computer vs. Brain

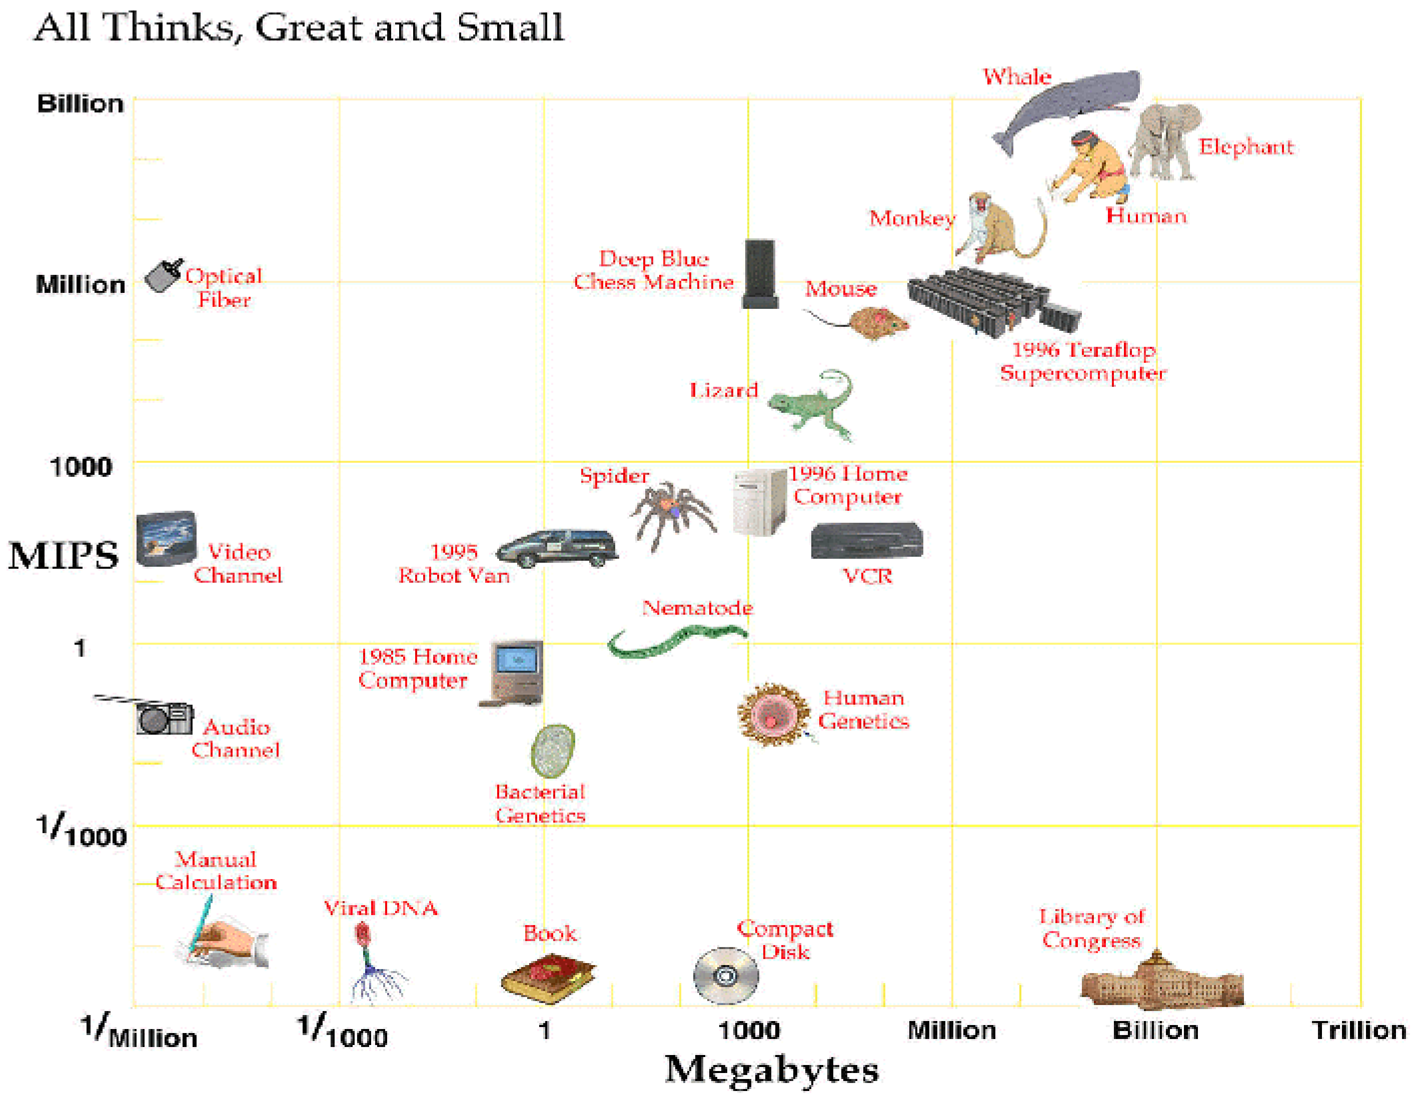

(from Moravec, 1998)

Evolution of Computers

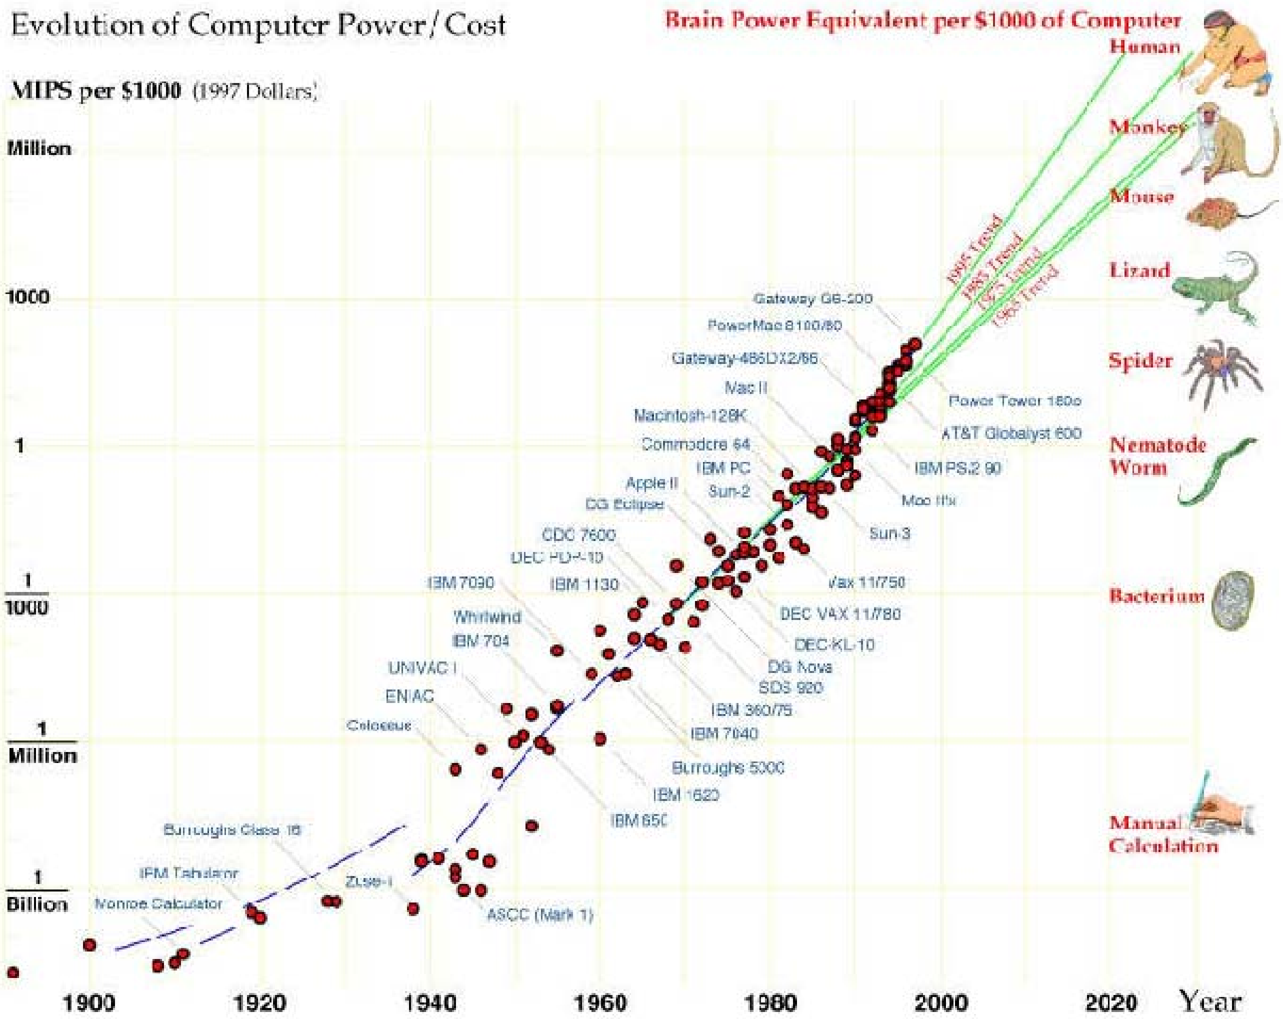

(from Moravec, 1998)

Projection

- In near future (~2020) computers will

  become cheap enough and have enough processing power and memory capacity to *match the general intellectual performance of the human brain*

- But…what "software" does the human brain run?

  Very much an open question

Defining AI Systems

	human-like	rational
thought	**Systems that think like humans**	**Systems that think rationally**
behavior	**Systems that act like humans**	**Systems that act rationally**

History of AI: Foundations

- Logic: rules of rational thought

  Aristotle (384-322 BC) – syllogisms

  Boole (1815-1864) – propositional logic

  Frege (1848-1925) – first-order logic

  Hilbert (1962-1943) – “Hilbert’s Program”

  Gödel (1906-1978) – incompleteness

  Turing (1912-1954) – computability, Turing test

  Cook (1971) – NP completeness

History of AI: Foundations

- Probability & Game Theory

  Cardano (1501-1576) – probabilities (*Liber de Ludo Aleae*)

  Bernoulli (1654-1705) – random variables

  Bayes (1702-1761) – belief update

  von Neumann (1944) – game theory

  Richard Bellman (1957) – Markov decision processes

Early AI

- Neural networks

  McCulloch & Pitts (1943) – simple neural nets

  Rosenblatt (1962) – perceptron learning

- Symbolic processing

  Dartmouth AI conference (1956)

  Newell & Simon – logic theorist

  John McCarthy – symbolic knowledge representation

  Arthur Samuel – Checkers program

Battle for the Soul of AI

- Minsky & Papert (1969) – *Perceptrons*
  - Single-layer networks cannot learn XOR
  - Argued against neural nets in general
- Backpropagation
  - Invented in 1969 and again in 1974
  - Hardware too slow, until rediscovered in 1985
- Research funding for neural nets disappears
- Rise of rule-based expert systems

Knowledge is Power

- Expert systems (1969-1980)
  - Dendral – molecular chemistry
  - Mycin – infectious disease
  - R1 – computer configuration
- AI Boom (1975-1985)
  - LISP machines – single user workstations
  - Japan's 5th Generation Project – massive parallel computing

AI Winter

- Expert systems oversold
  - Fragile
  - Hard to build, maintain
- AI Winter (1985-1990)
- Science went on... looking for
  - Principles for robust reasoning
  - Principles for learning

AI Now

- Probabilistic graphical models

  Pearl (1988) – Bayesian networks
- Machine learning

  Quinlan (1993) – decision trees (C4.5)

  Vapnik (1992) – Support vector machines (SVMs)

  Schapire (1996) – Boosting

  Neal (1996) – Gaussian processes
- Recent progress:

  Probabilistic relational models, deep networks, active learning, structured prediction, etc.

AI Now: Applications

- Countless AI systems in day to day use
  - Industrial robotics
  - Data mining on the web
  - Speech recognition
  - Security: Face & Iris recognition
  - Stock market prediction
  - Space exploration
  - Computational biology
  - Hardware verification
  - Credit card fraud detection
  - Surveillance and threat assessment
  - Military applications (bomb-defusing robots, drones)
  - Etc.

Notable Examples: Chess (Deep Blue, 1997)

Deep blue wins 2-1-3 (wins-losses-draws)

"I could feel –
I could smell –
a new kind of
intelligence
across the
table"

-Gary
Kasparov

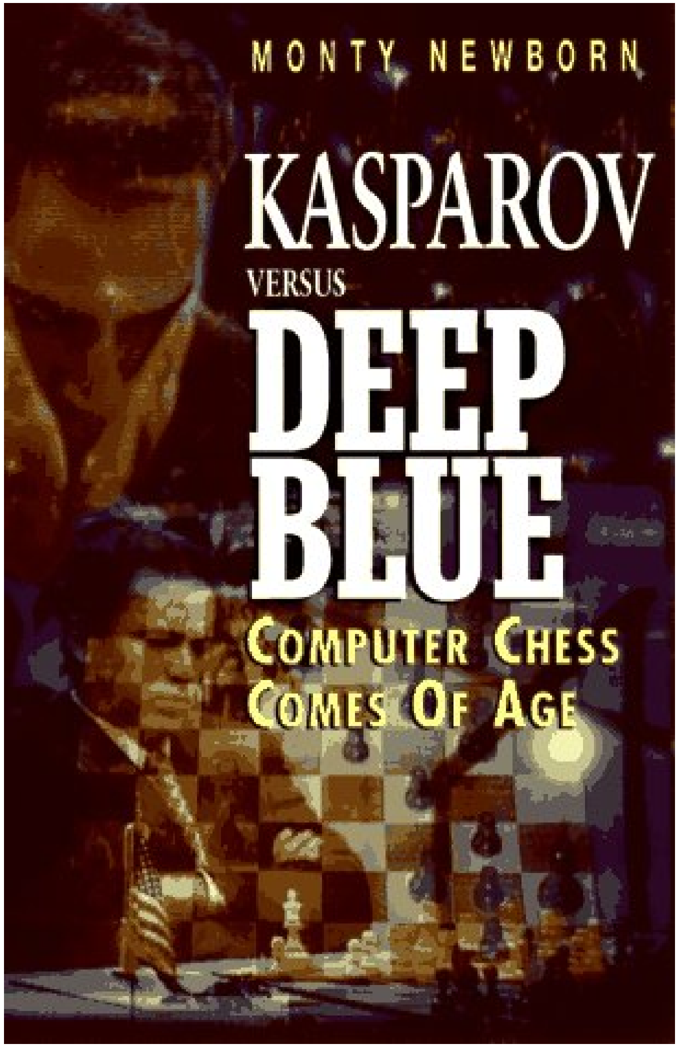

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

Speech Recognition

Navigation Systems

Automated call centers

Natural Language Understanding

- Speech Recognition
  - "word spotting" feasible today
  - continuous speech – inching closer
- WWW Information Extraction
  - E.g., KnowItAll project
- Machine Translation / Understanding
  - ***The spirit is willing but the flesh is weak.* (English)**
  - ***The vodka is good but the meat is rotten.* (Russian)**
  - (i.e., very much a work in progress…)

Museum Tour-Guide Robots

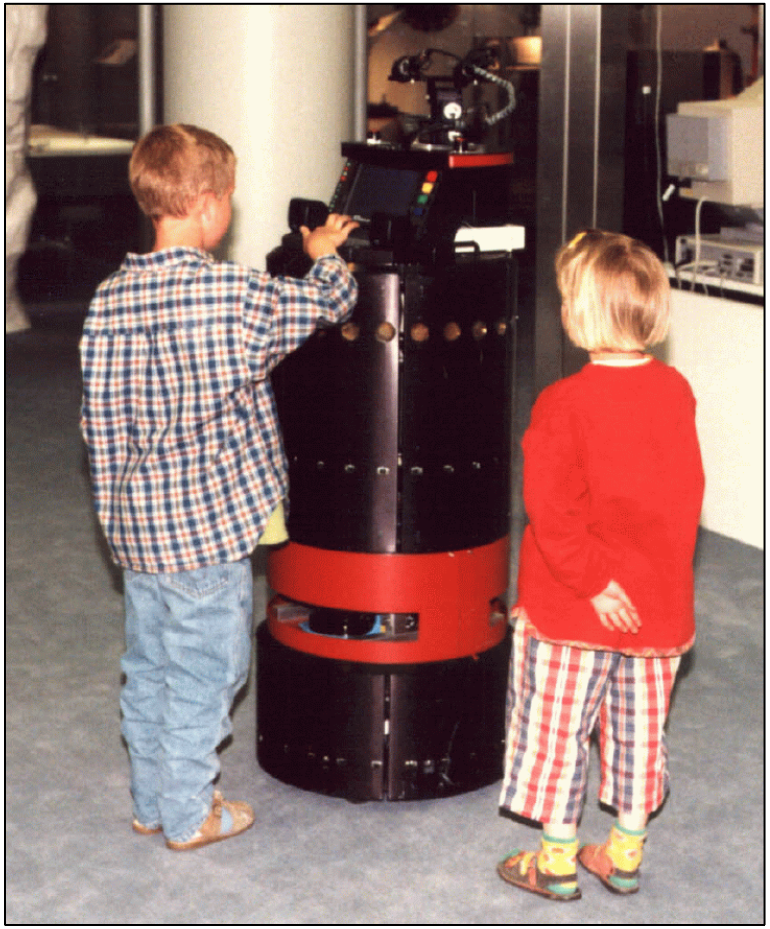

Rhino, 1997

Minerva, 1998

Mars Rovers (2003-now)

Europa Mission ~ 2018?

Humanoid Robots

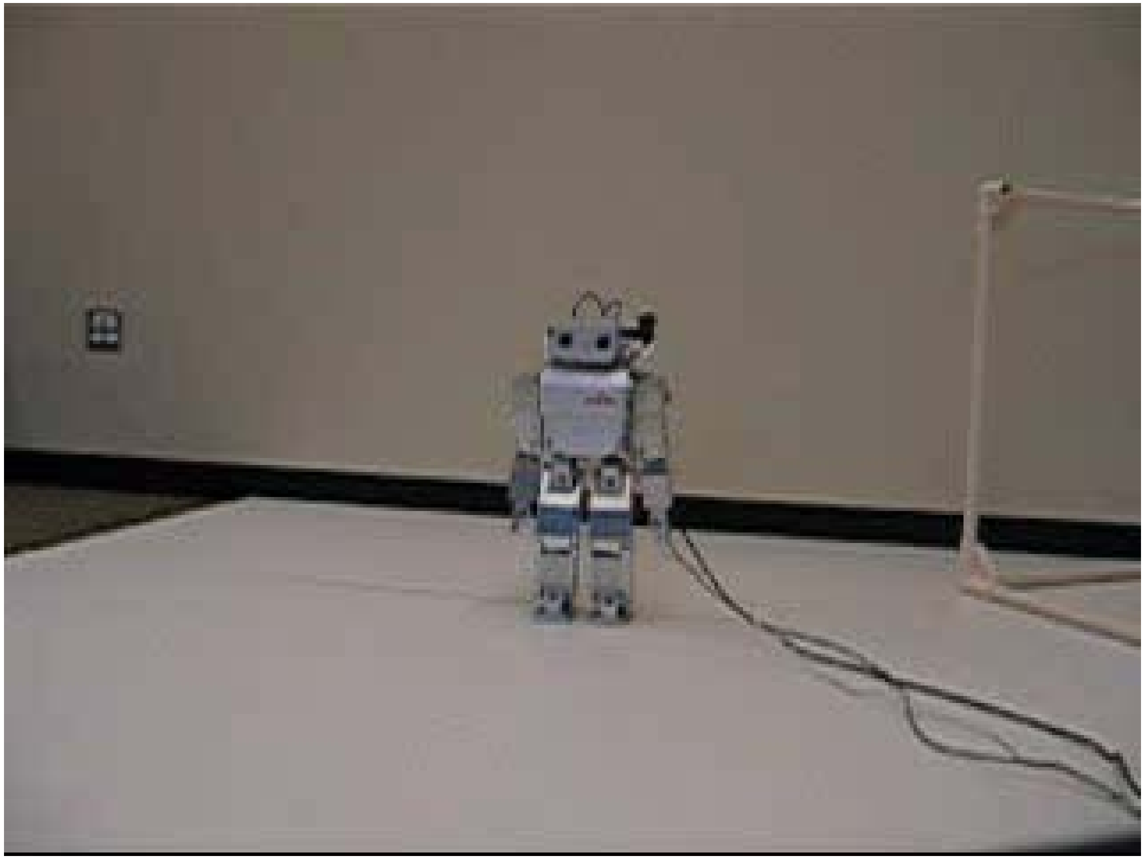

Humanoid robot “Mo” in UW CSE’s Neural Systems Lab

Robots that Learn

Before Learning

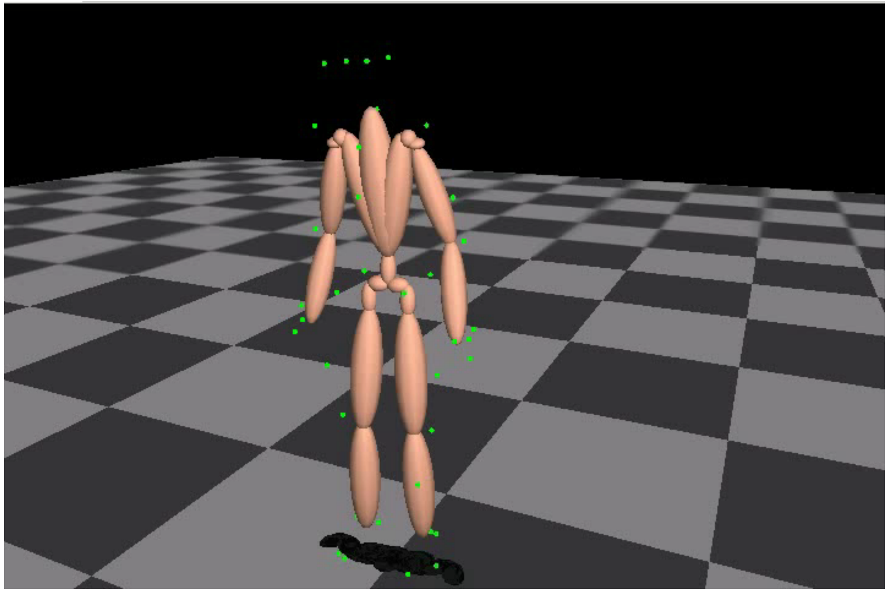

Human Motion Capture

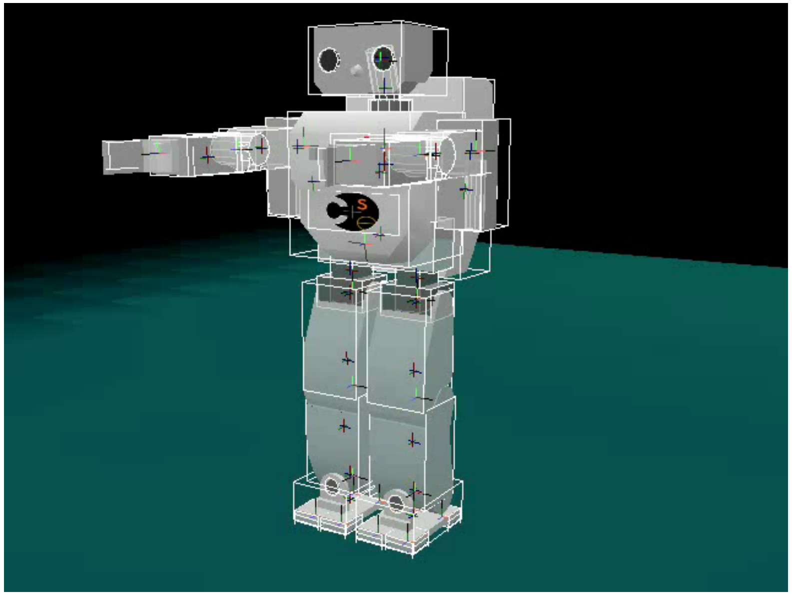

Attempted Imitation

Robots that Learn

After Learning

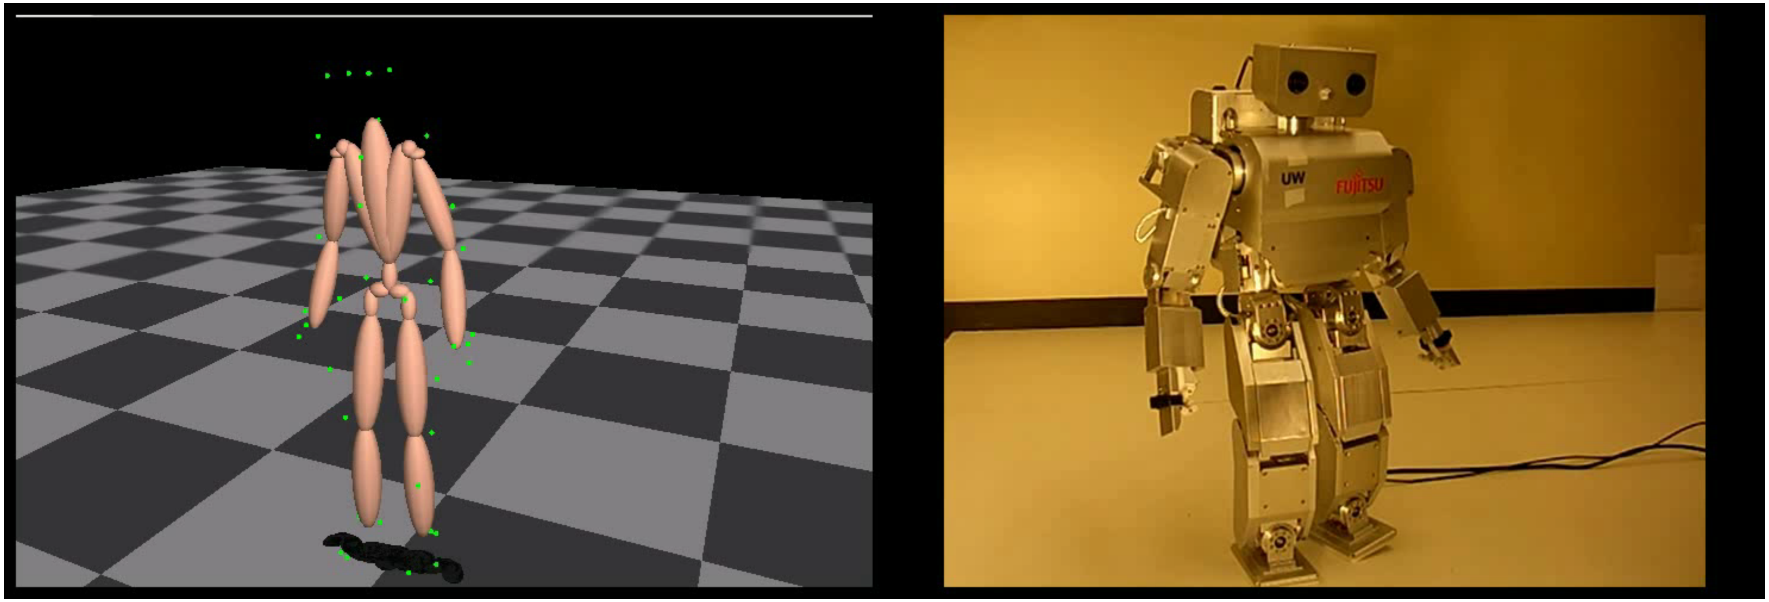

Chess Playing vs. Robots

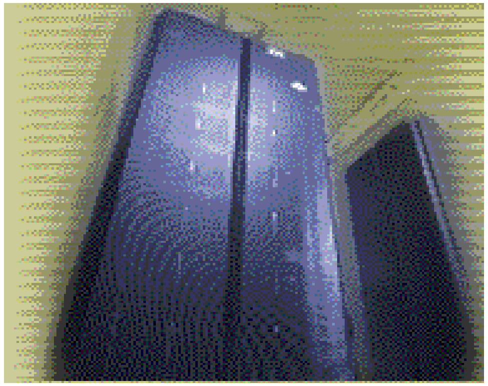

Deep Blue

Robot

- Static
- Deterministic
- Turn-based

- Dynamic
- Stochastic
- Real-time

Robotic Prosthetics

Brain-Computer Interfaces

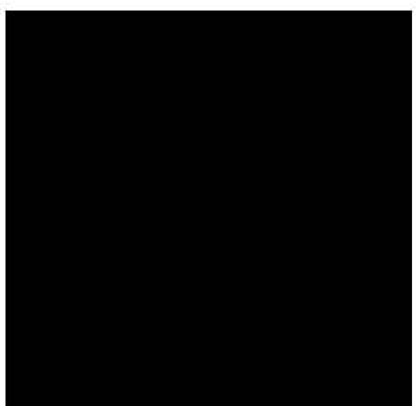

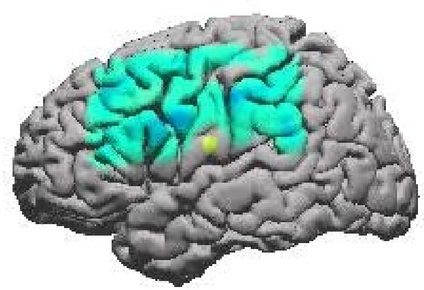

Limitations of AI Systems Today

- Today's successful AI systems
  - operate in well-defined domains
  - employ narrow, specialized hard-wired knowledge
- *Needed:* Ability to
  - Operate in complex, open-ended dynamic worlds
    - E.g., Your kitchen vs. GM factory floor
  - Adapt to unforeseen circumstances
  - Learn from new experiences
- In this class, we will explore some potentially useful techniques for tackling these problems

5 Minute Break…

Next:

Agents & Environments (Chapter 2 in AIMA)

Outline

- Agents and environments
- Rationality
- PEAS specification
- Environment types
- Agent types

Agents

- An agent is any entity that can perceive its environment through sensors and act upon that environment through actuators

- Human agent:

  Sensors: Eyes, ears, and other organs

  Actuators: Hands, legs, mouth, etc.

- Robotic agent:

  Sensors: Cameras, laser range finders, etc.

  Actuators: Motorized limbs, wheels, etc.

Types of Agents

- Immobots (Immobile Robots)
  - Intelligent buildings
  - Intelligent forests

- Softbots
  - Jango (early softbot for shopping)
  - Microsoft Clippy
  - Askjeeves.com (now Ask.com)
  - Expert Systems
    - Cardiologist

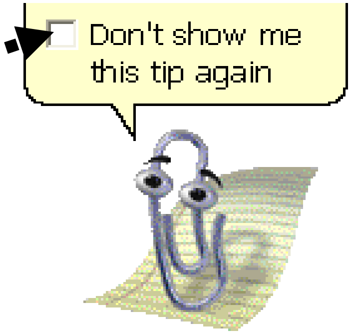

Intelligent Agents

- Have sensors and actuators (effectors)
- Implement mapping from percept sequence to actions

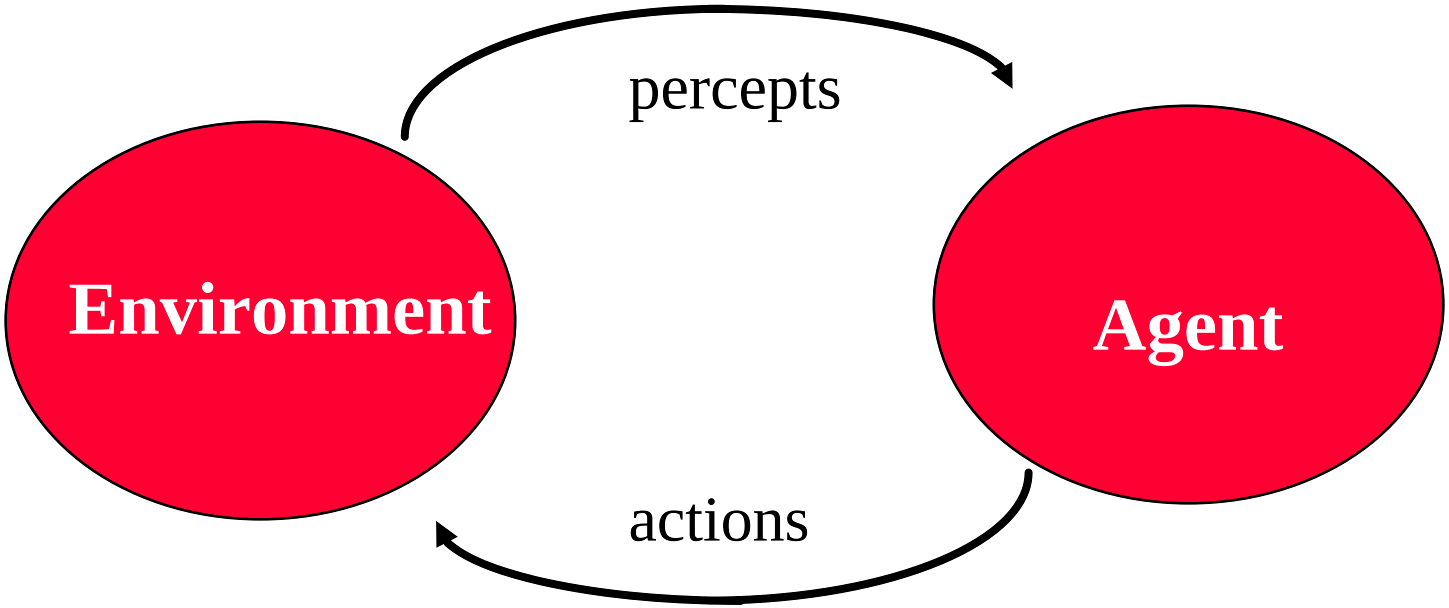

- Maximize a Performance Measure

Performance Measures

- Performance measure = An objective criterion for success of an agent's behavior
- E.g., vacuum cleaner agent

  performance measure: amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Rational Agent

“For each possible percept sequence, does whatever action is expected to maximize its performance measure on the basis of evidence perceived so far and built-in knowledge."

- Rationality vs. omniscience
  - Rationality maximizes *expected* performance
  - Omniscience maximizes *actual* performance (but impossible to achieve in reality)
- Rational agents need to use information gathering actions and learning

Autonomy

A rational agent is autonomous if it can learn to compensate for partial or incorrect prior knowledge

Why is this important?

Task Environments

- The “task environment” for an agent is comprised of PEAS (Performance measure, Environment, Actuators, Sensors)
- E.g., Consider the task of designing an automated taxi driver:

    Performance measure = ?

    Environment = ?

    Actuators = ?

    Sensors = ?

PEAS

- PEAS for Automated taxi driver

- Performance measure:
  Safe, fast, legal, comfortable trip, maximize profits
- Environment:
  Roads, other traffic, pedestrians, customers
- Actuators:
  Steering wheel, accelerator, brake, signal, horn
- Sensors:
  Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS

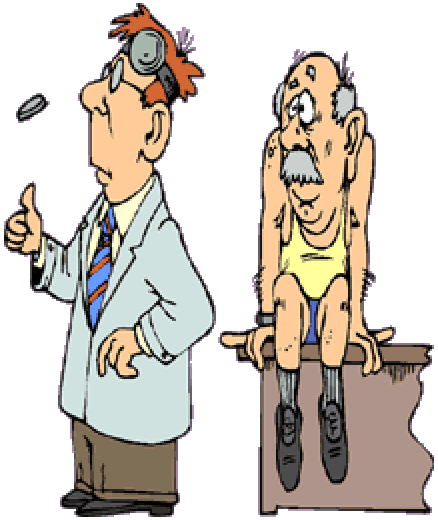

- PEAS for Medical diagnosis system

- Performance measure:

  Healthy patient, minimize costs, lawsuits
- Environment:

  Patient, hospital, staff
- Actuators:

  Screen display (questions, tests, diagnoses, treatments, referrals)
- Sensors:

  Keyboard (entry of symptoms, findings, patient's answers)

Properties of Environments

- Observability: full *vs.* partial
  Sensors detect all aspects of state of environment relevant to choice of action?
- Deterministic *vs.* stochastic
  Next state completely determined by current state and action?
- Episodic *vs.* sequential
  Current action independent of previous actions?
- Static *vs.* dynamic
  Can environment change over time?
- Discrete *vs.* continuous
  State of environment, time, percepts, and actions discrete or continuous-valued?
- Single *vs.* multiagent

Properties of Environments

- Observability: full *vs.* partial
- Deterministic *vs.* stochastic
- Episodic *vs.* sequential
- Static *vs.* dynamic
- Discrete *vs.* continuous
- Single *vs.* multiagent

- **Crossword puzzle**
- **Chess**
- **Poker**
- **Coffee delivery mobile robot**

Agent Functions and Agent Programs

- An agent's behavior can be *described* by an agent function mapping percept sequences to actions taken by the agent
- An *implementation* of an agent function running on the agent architecture (e.g., a robot) is called an agent program
- Our goal: Develop concise agent programs for implementing rational agents

Example

Vacuum-cleaner world

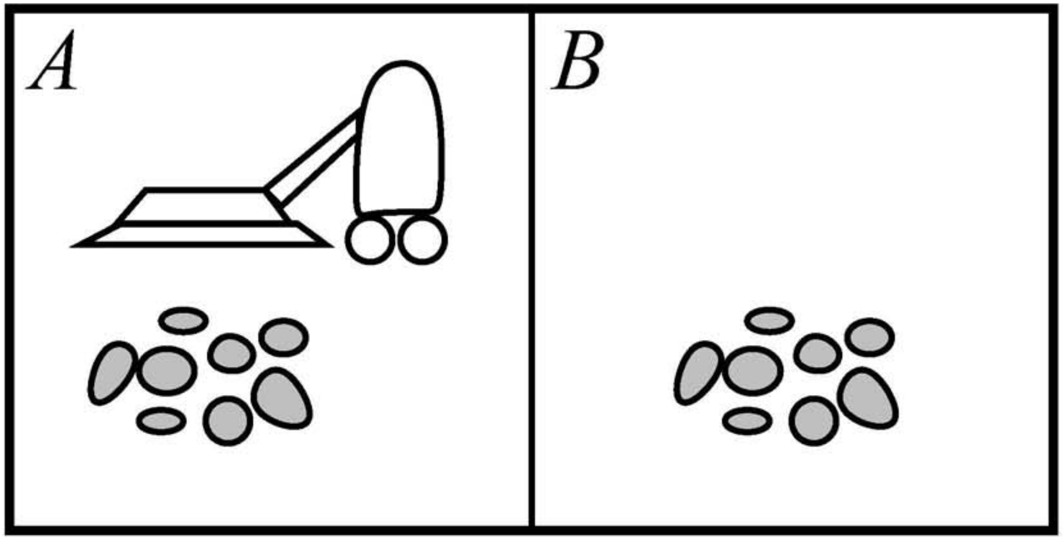

Percepts: location and contents, e.g., $[A, Dirty]$

Actions: $Left$, $Right$, $Suck$, $NoOp$

Example: vacuum world state space graph

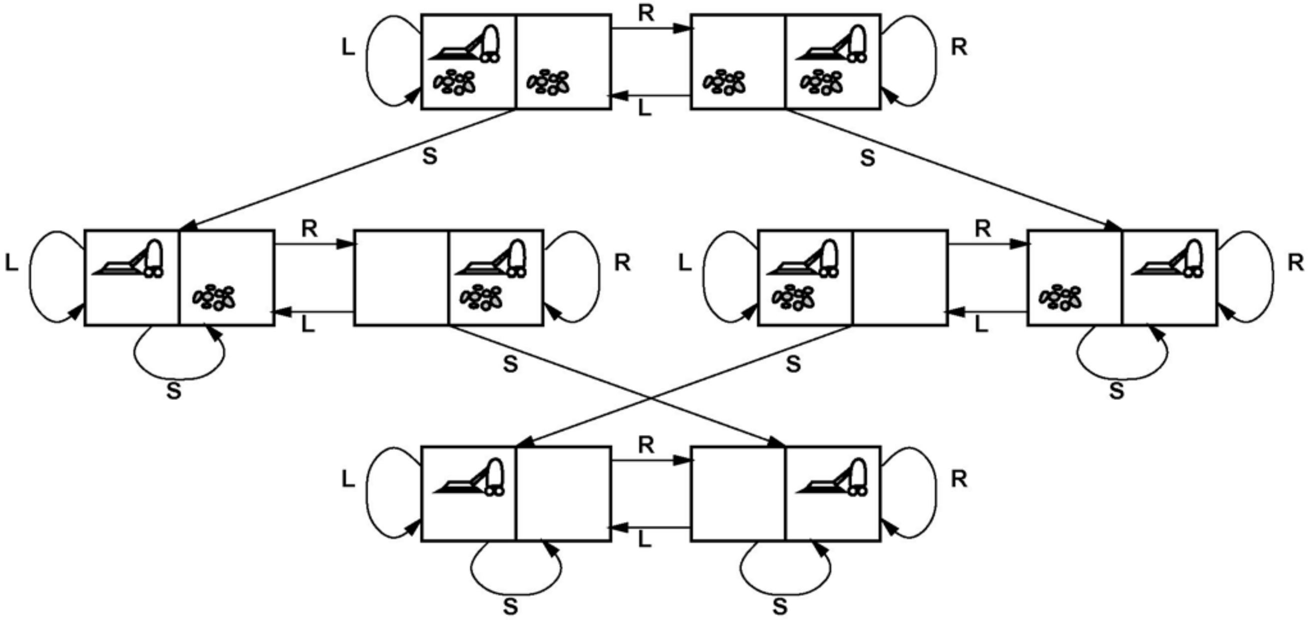

How should the agent be designed if…

- It has location and dirt sensors, but no internal state?
- It has no sensors, but knows the starting state?
- It has no sensors, and does not know the starting state?

Implementing Rational Agents

- Table lookup based on percept sequences
  - Infeasible
- Agent programs:
  - Simple reflex agents
  - Agents with memory
    - Reflex agent with internal state
    - Goal-based agents
    - Utility-based agents

Simple Reflex Agents

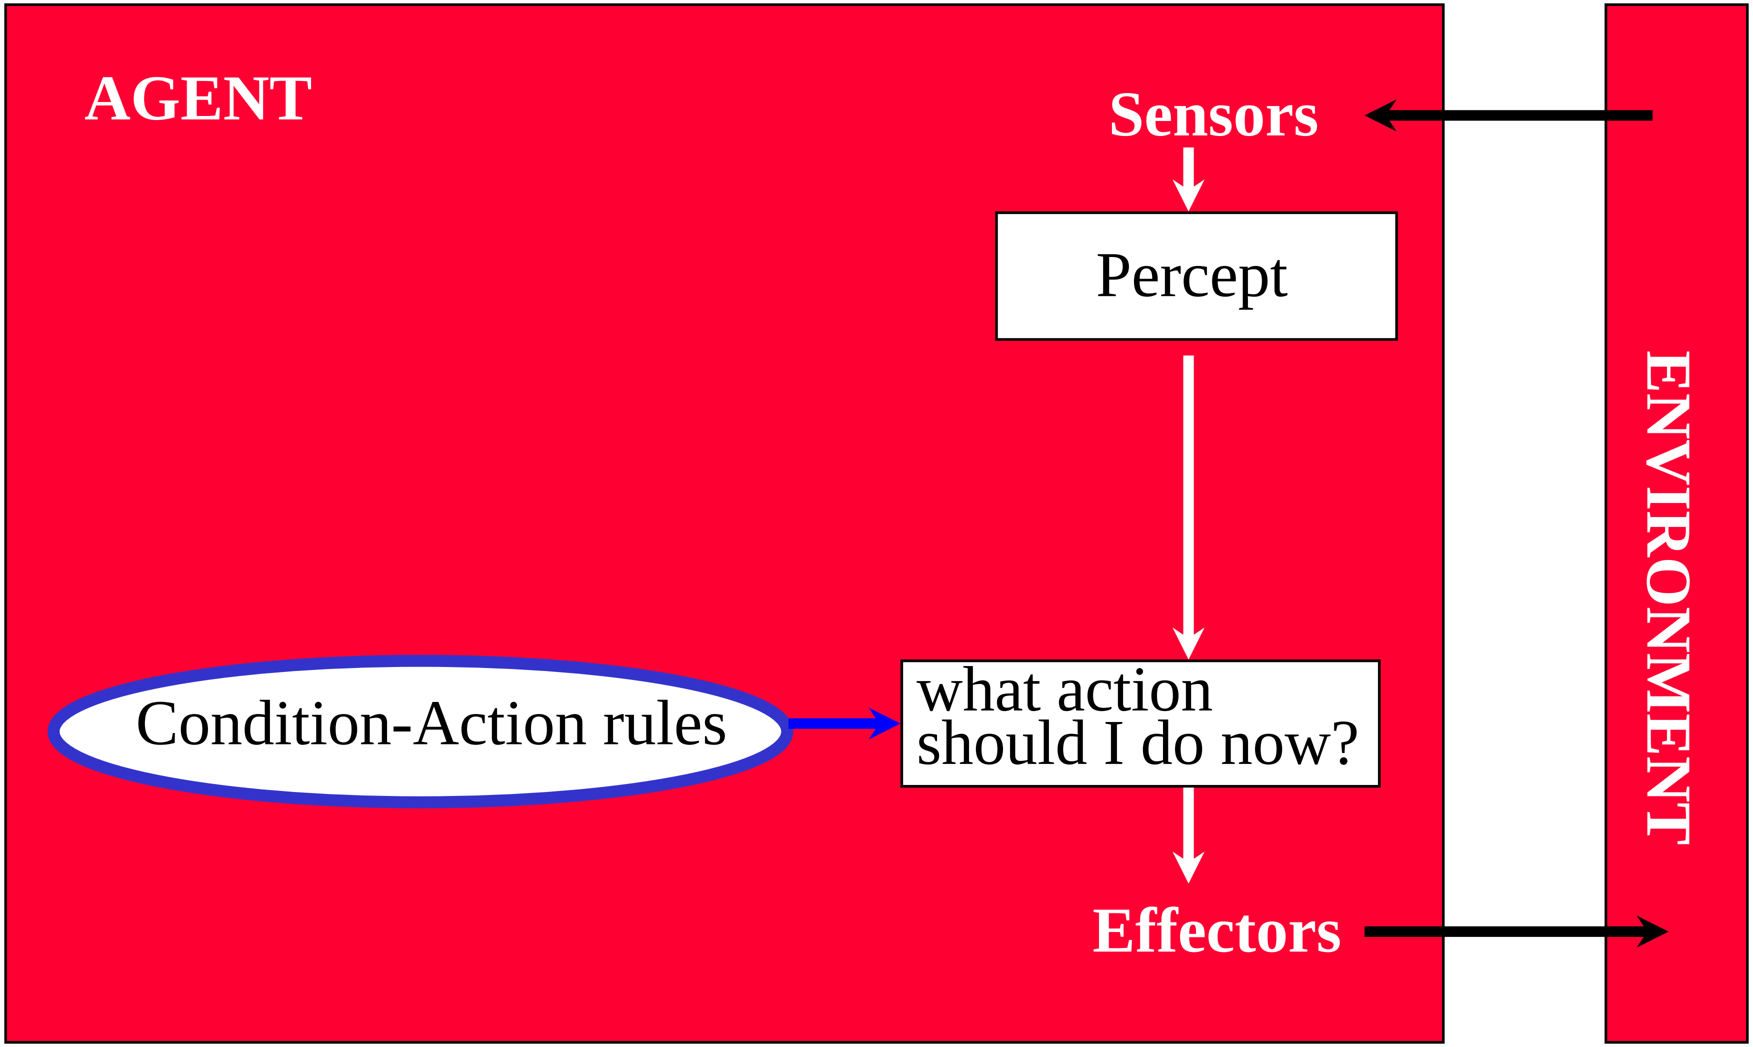

Simple Reflex Agents

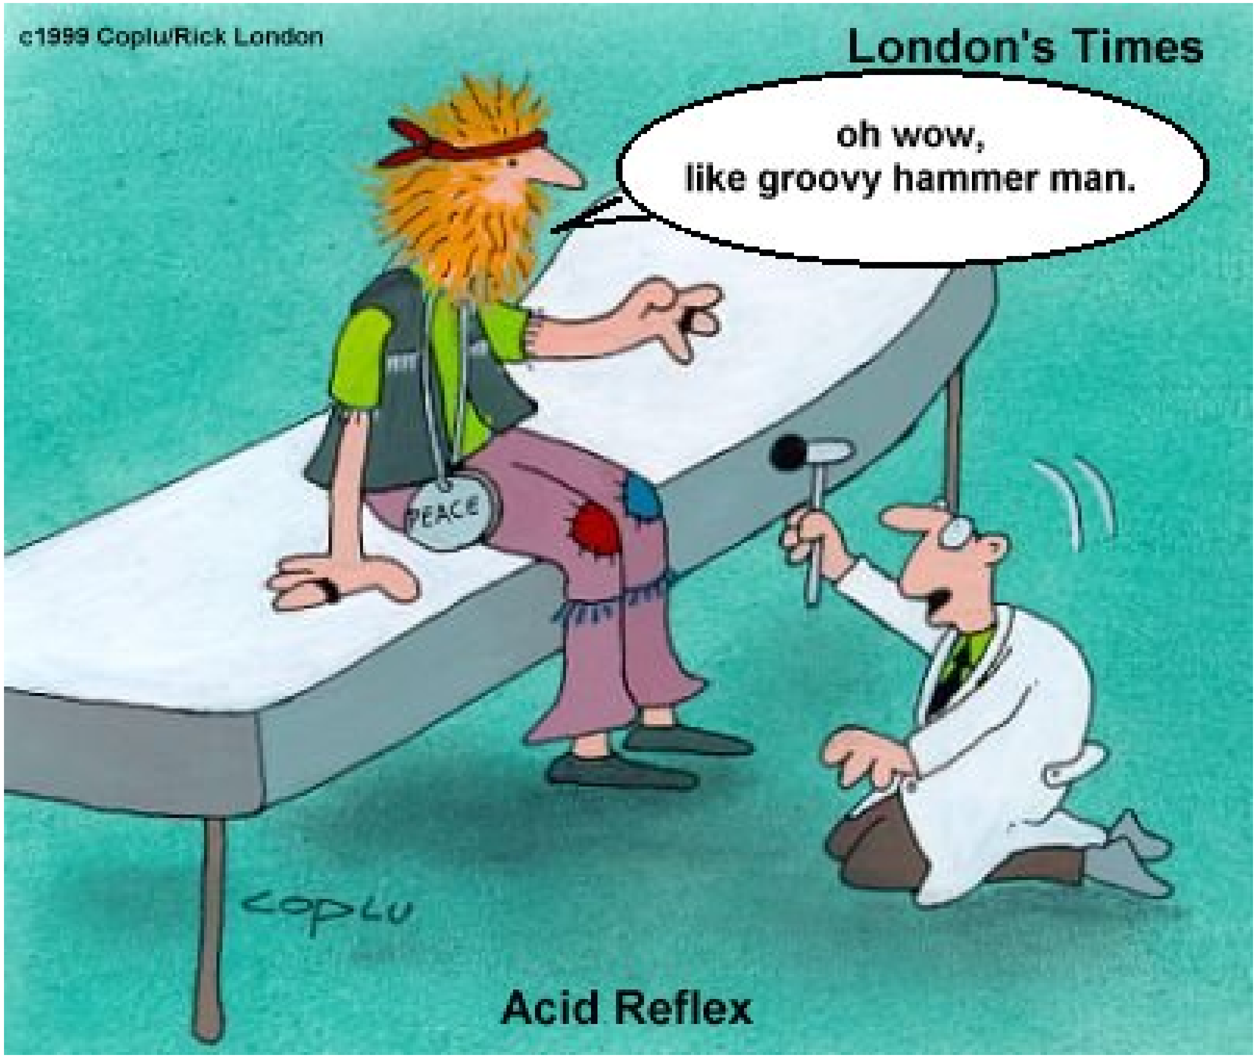

Reflex Agent with Internal State

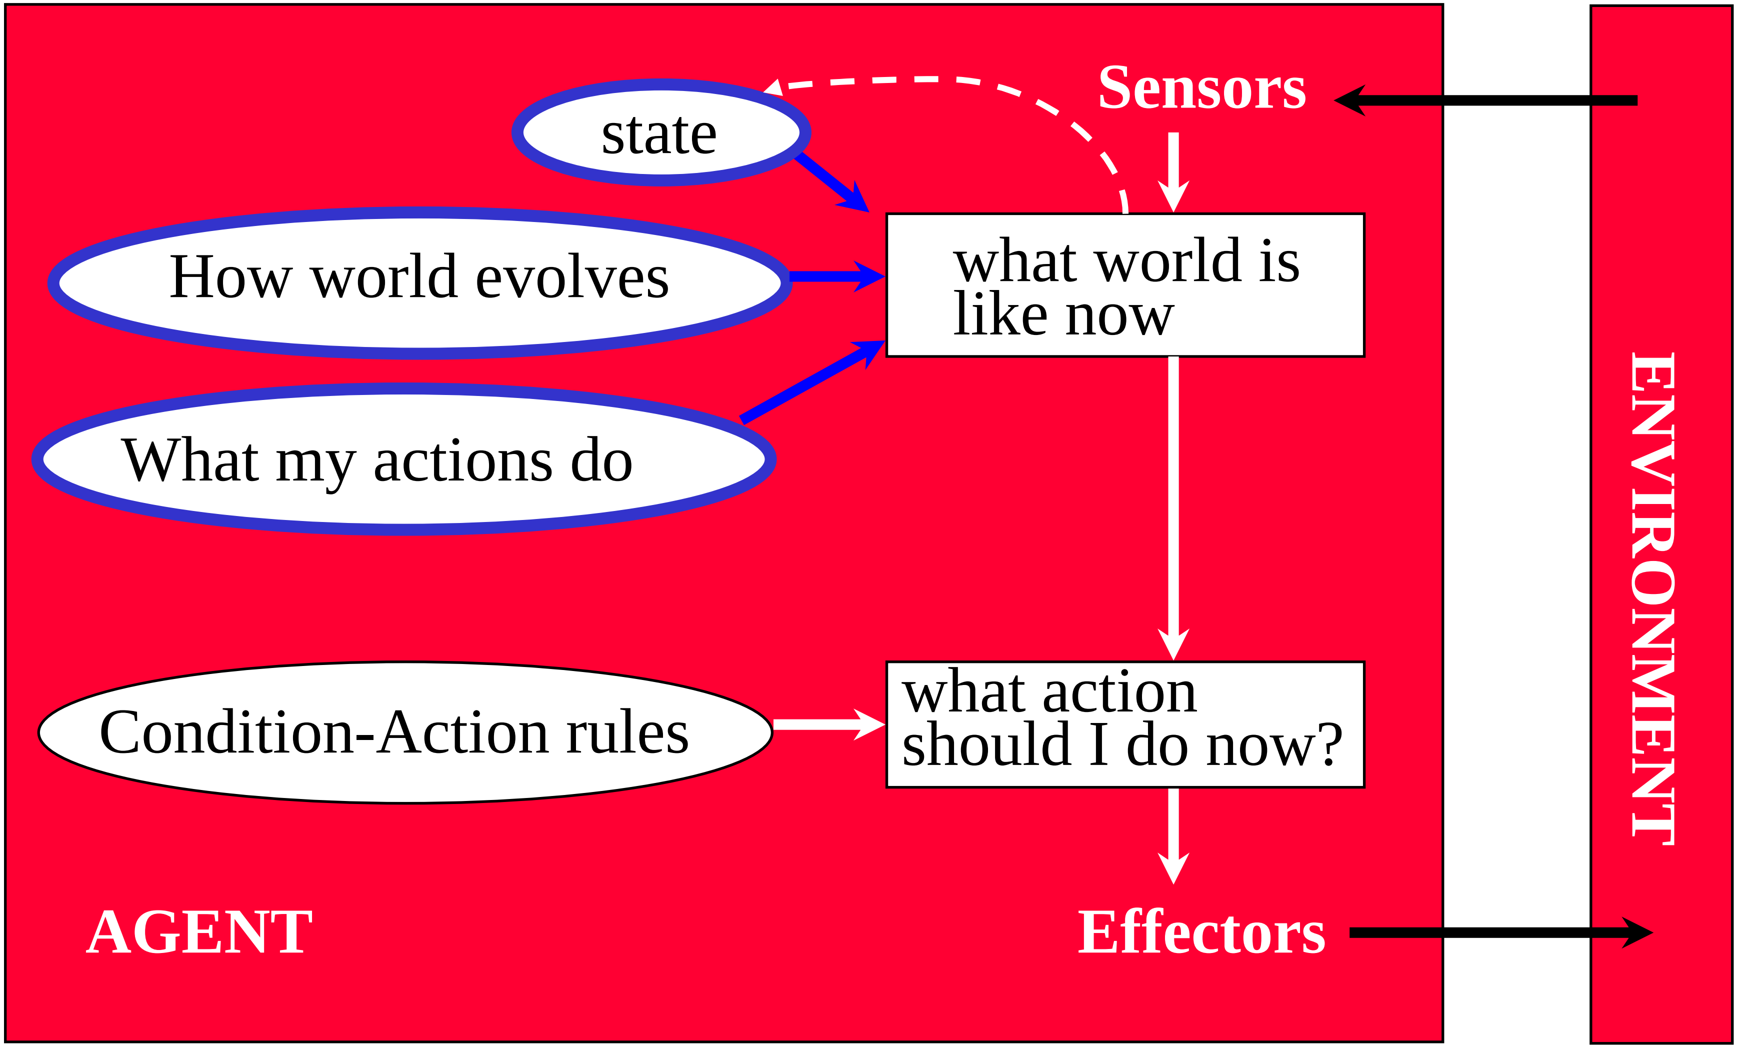

Goal-Based (Planning) Agents

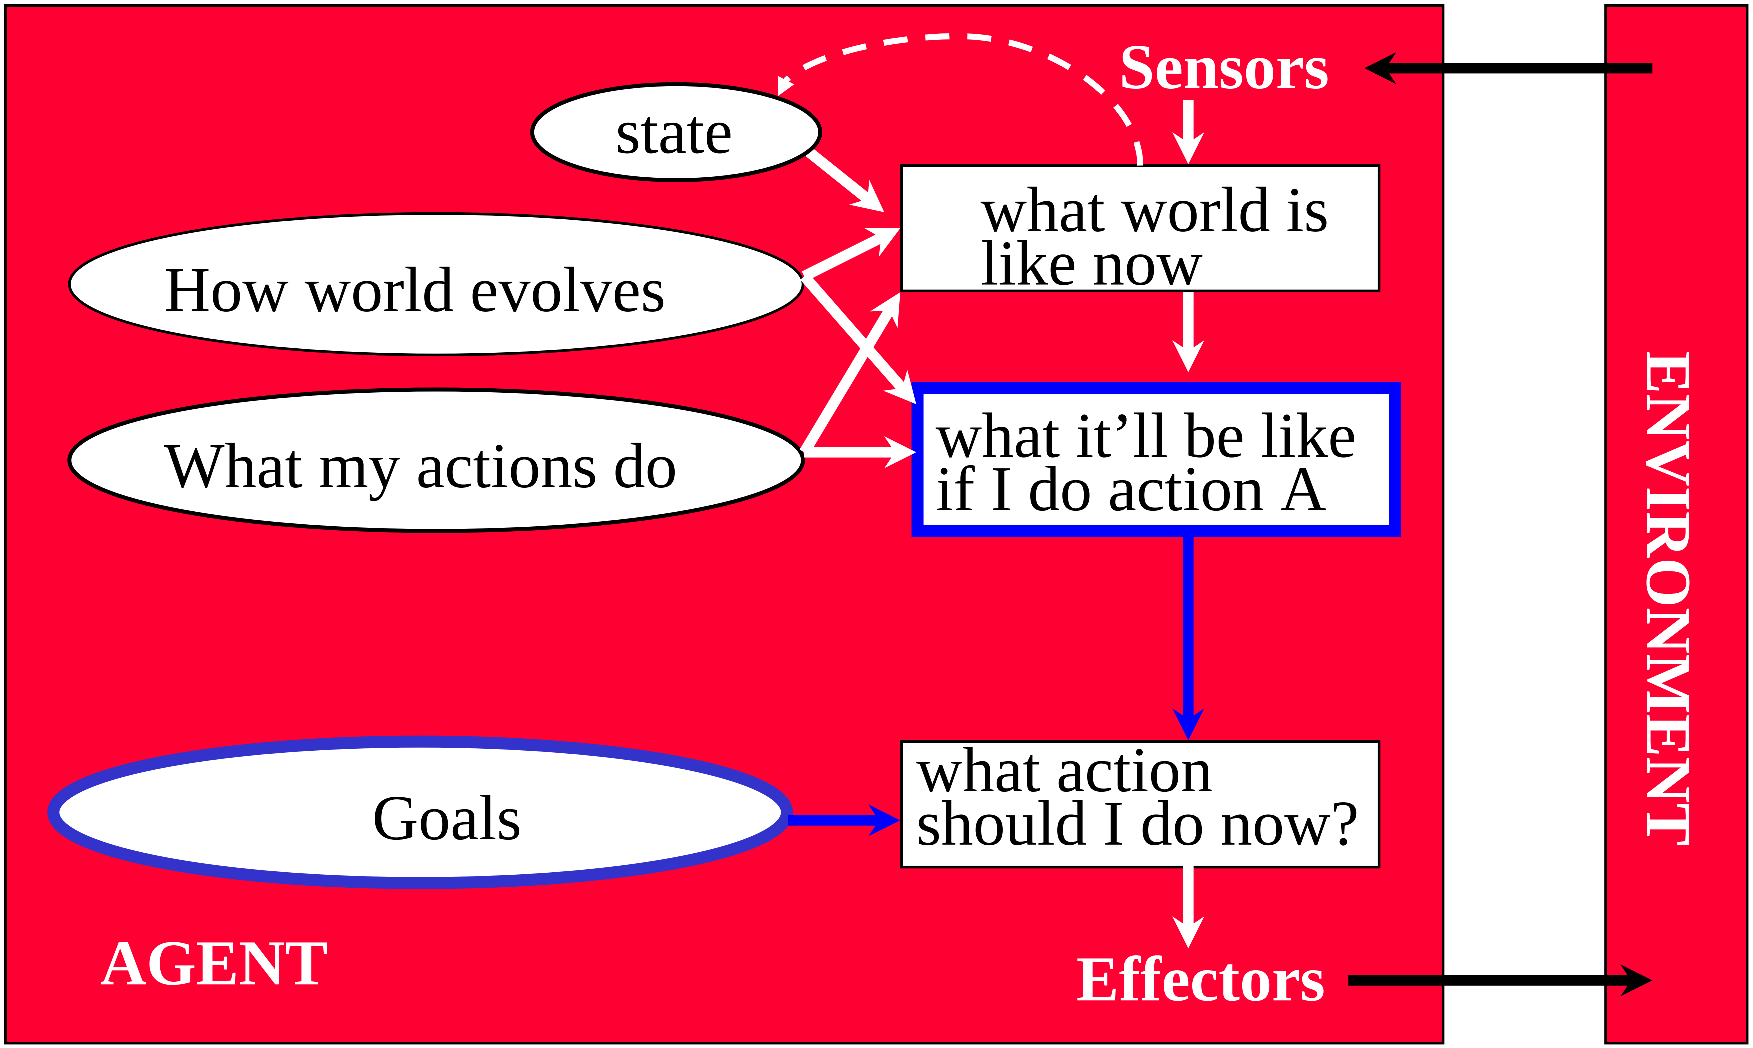

Utility-Based Agents

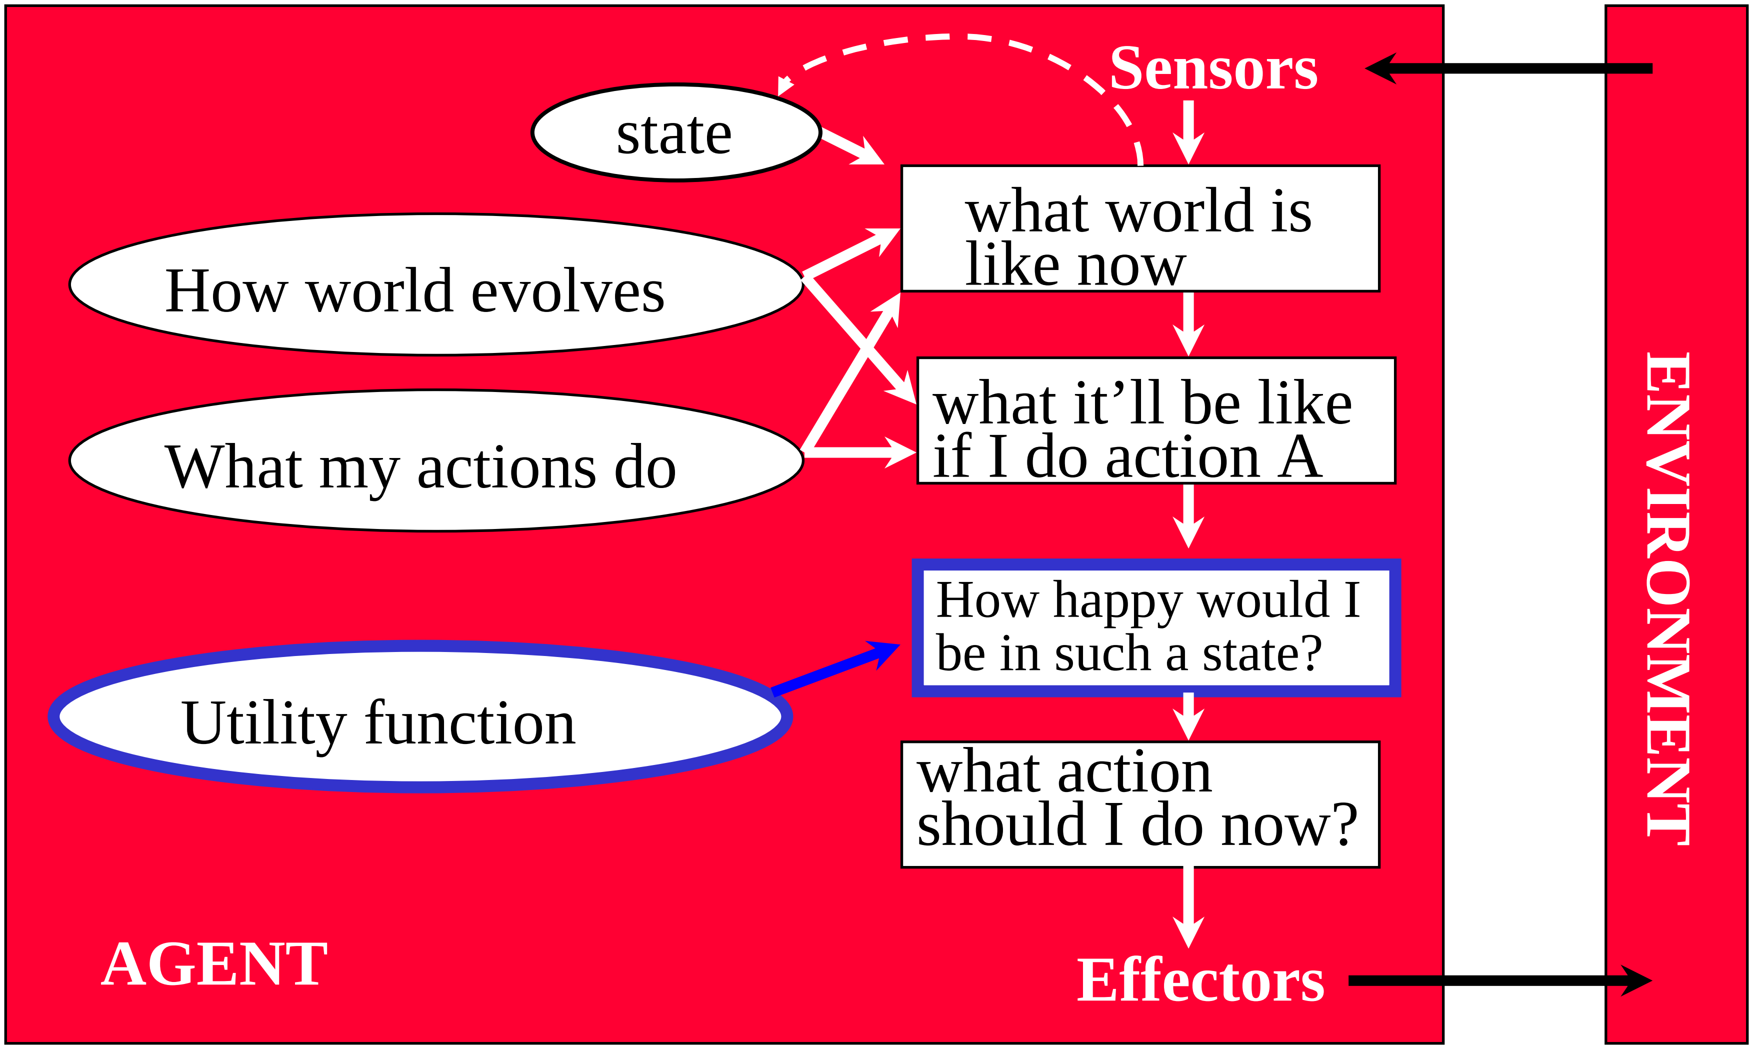

Learning Agents

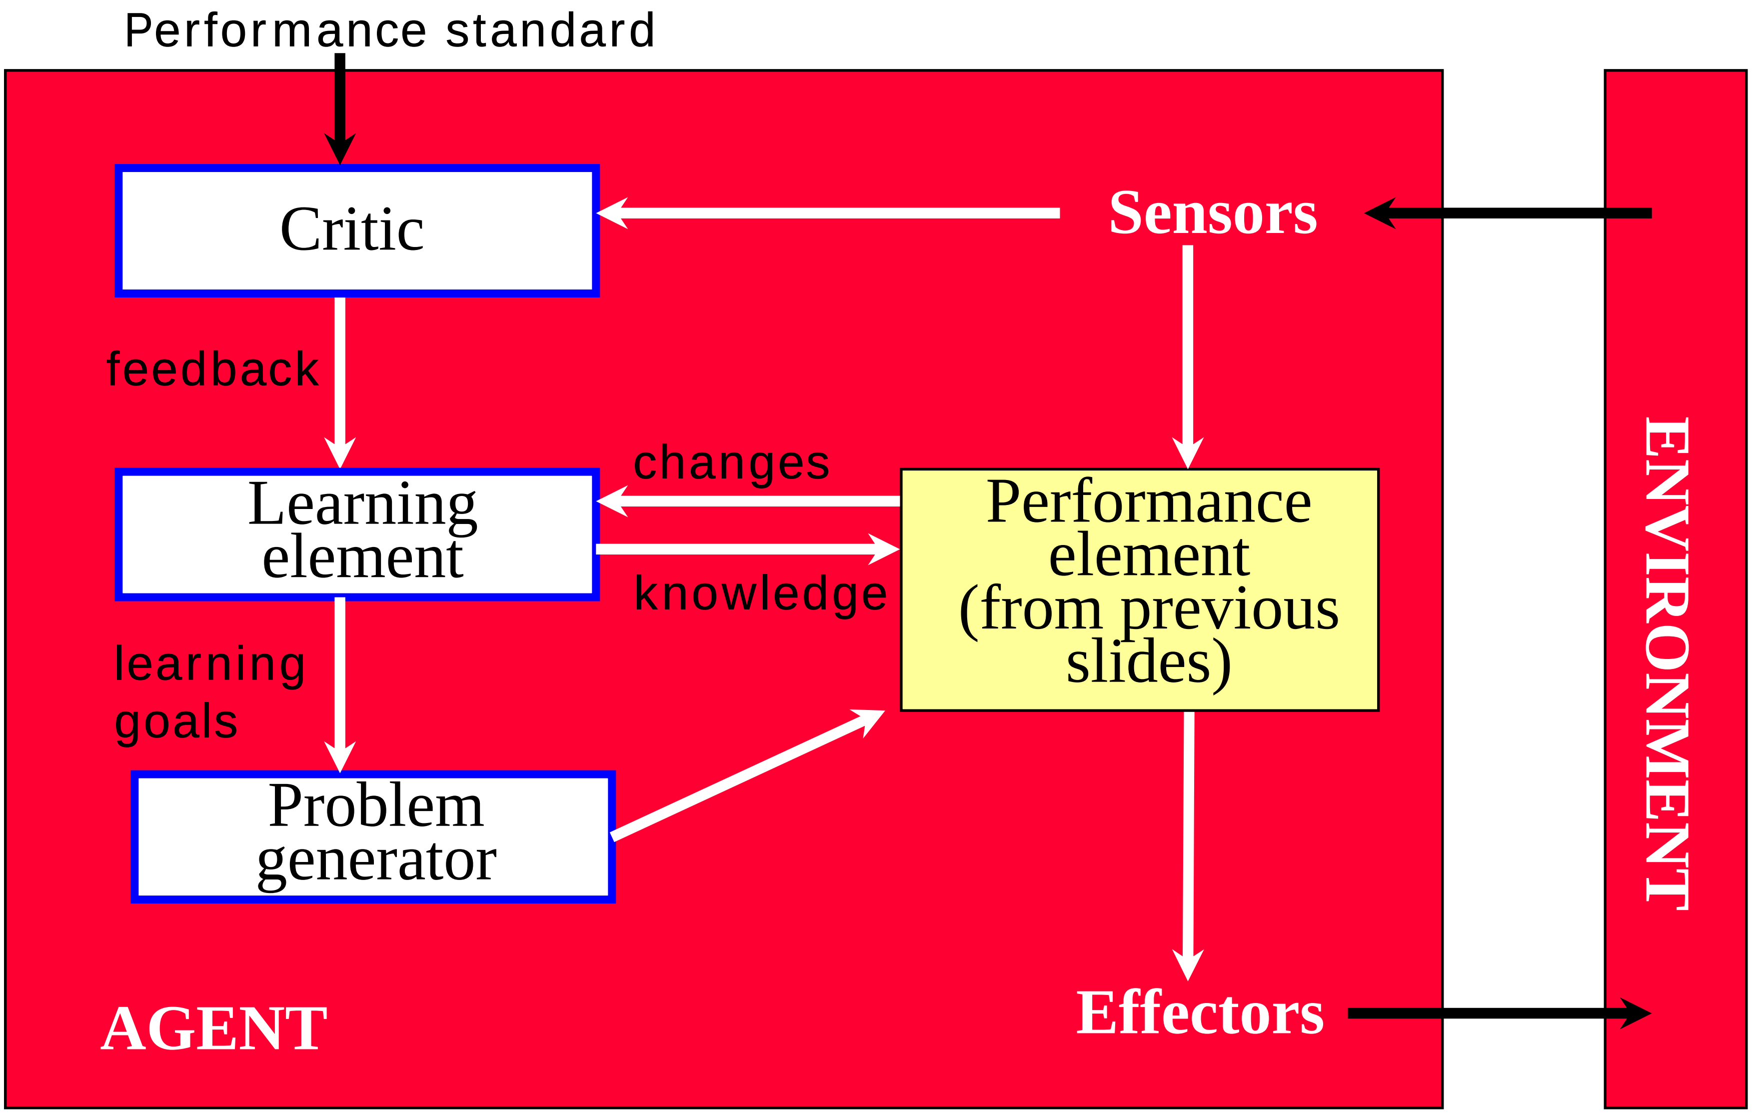

While driving, what's the best policy?

- Always stop at a stop sign
- Never stop at a stop sign
- Look around for other cars and stop only if you see one approaching
- Look around for a cop and stop only if you see one

- **What kind of agent are you?**
  - reflex, goal-based, utility-based?

Best policy not applicable

(http://www.gonomad.com/traveltalesfromindia/archives/2007_09_01_archive.html)

For You To Do

- Browse CSEP 573 course web page
- Get on class mailing list
- Read Chapters 3-5 in AIMA text
- HW #1 to be assigned next week (watch course website)