CS 460: Artificial Intelligence

- <u>Instructor:</u> Prof. Laurent Itti, itti@pollux.usc.edu
- <u>Lectures:</u> Th 5:00-7:50, WPH-B27
- <u>Office hours:</u> Mon 2:00 – 4:00 pm, HNB-30A, and by appointment

- **<u>Course web page:</u> http://iLab.usc.edu/classes/2004cs460**
  - Up to date information
  - Lecture notes
  - Relevant dates, links, etc.

- <u>Course material:</u>
  - [AIMA] **Artificial Intelligence: A Modern Approach, by Stuart Russell and Peter Norvig.**

CS 460: Artificial Intelligence

- Course overview: foundations of symbolic intelligent systems. Agents, search, problem solving, logic, representation, reasoning, symbolic programming, and robotics.

- Prerequisites: CS 455x, i.e., programming principles, discrete mathematics for computing, software design and software engineering concepts. Some knowledge of C/C++ for some programming assignments.

- **Grading:** 30% for midterm +
  30% for final +
  40% for mandatory homeworks/assignments

Practical issues

- **Class mailing list:**
  will be setup on the backboard system at learn.usc.edu

- **Submissions:** See class web page under Assignments

Why study AI?

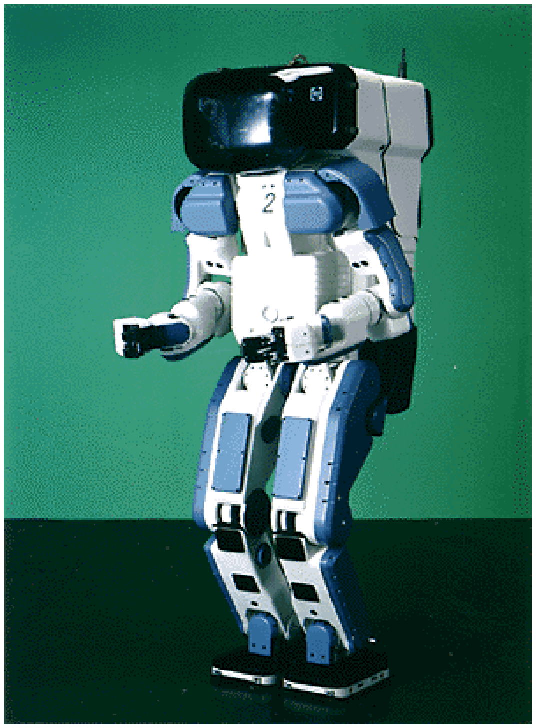

Labor

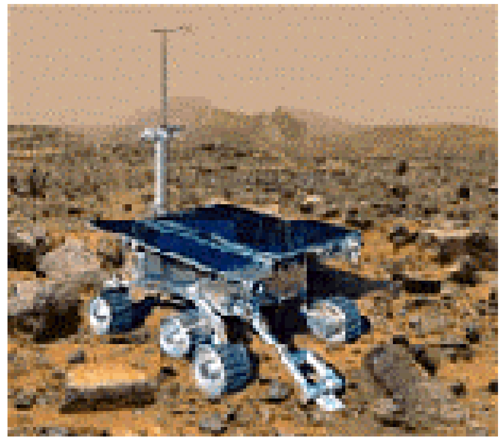

Science

Google

YAHOO!

Search engines

Medicine/
Diagnosis

Appliances

What else?

Honda Humanoid Robot

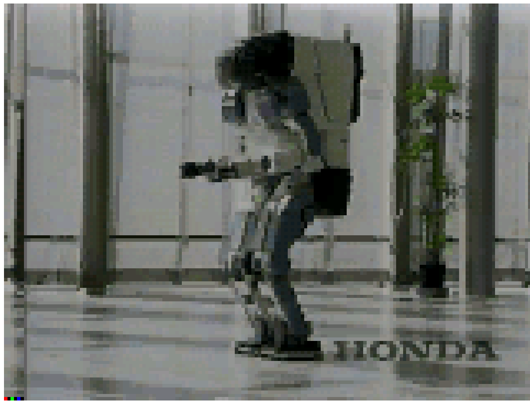

Walk

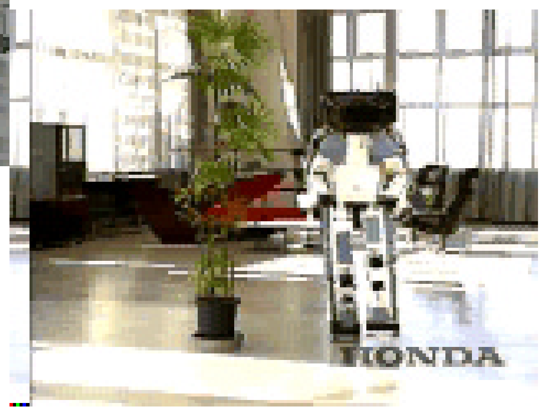

Turn

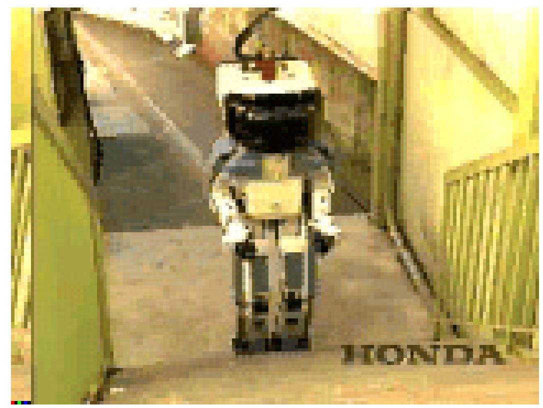

Stairs

http://world.honda.com/robot/

Sony AIBO

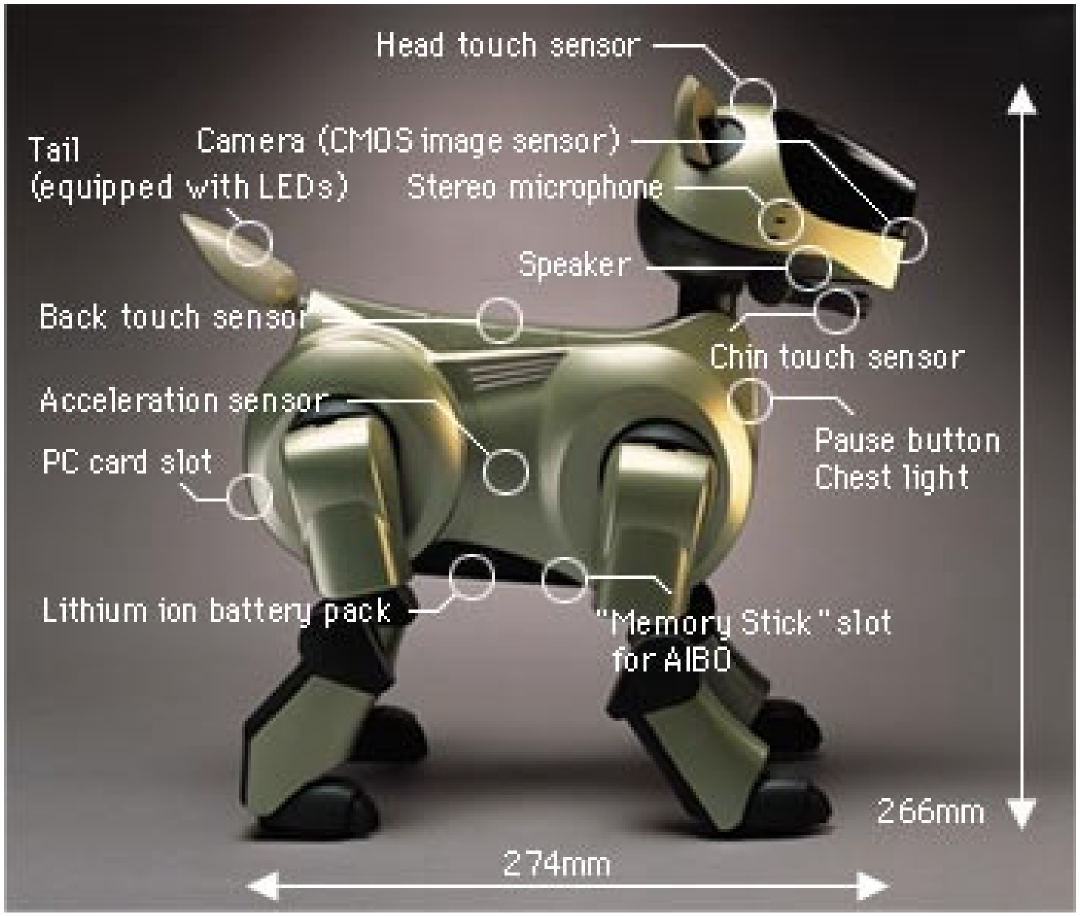

http://www.aibo.com

Natural Language Question Answering

http://aimovie.warnerbros.com http://www.ai.mit.edu/projects/infolab/

Robot Teams

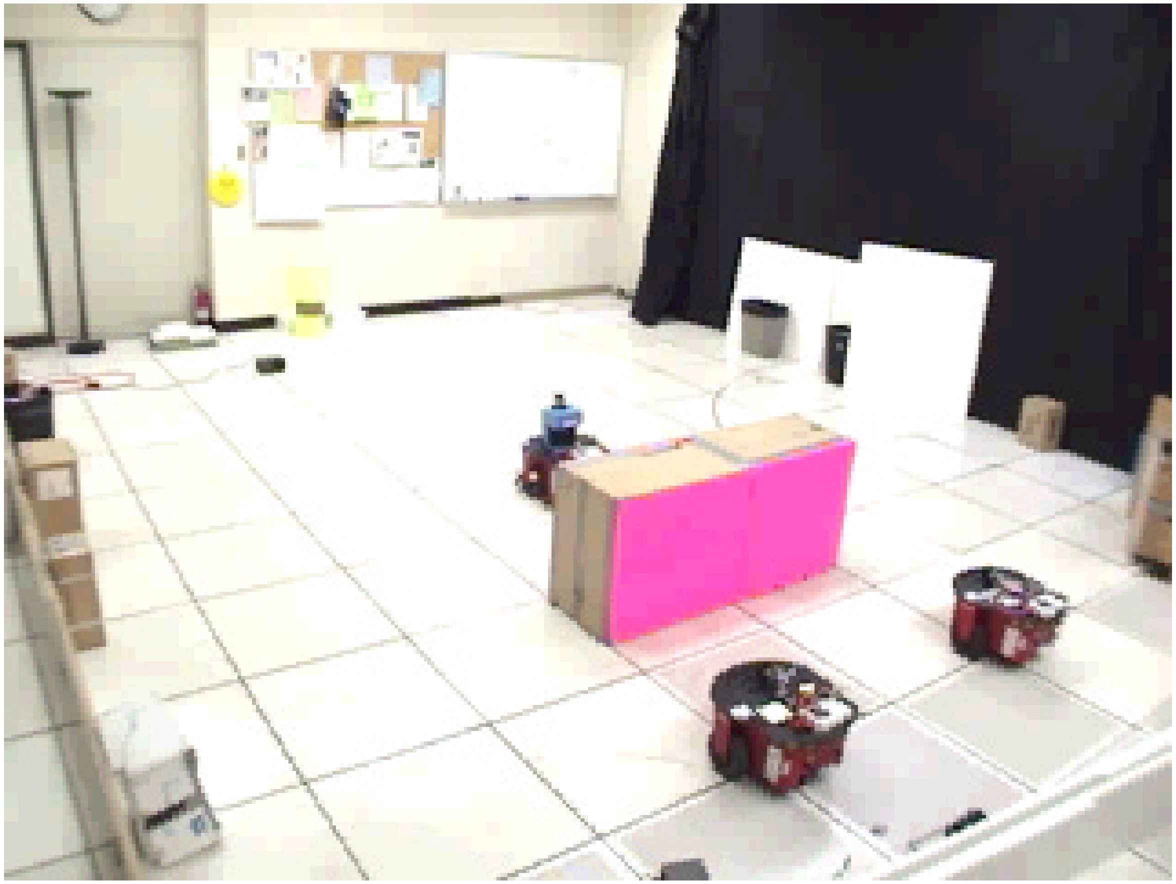

USC robotics Lab

What is AI?

The exciting new effort to make computers thinks … *machine with minds,* in the full and literal sense" (Haugeland 1985)	"The study of mental faculties through the use of computational models" (Charniak et al. 1985)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990)	A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkol, 1990)

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Acting Humanly: The Turing Test

- Alan Turing's 1950 article *Computing Machinery and Intelligence* discussed conditions for considering a machine to be intelligent
  - “Can machines think?” ←→ “Can machines behave intelligently?”
  - The Turing test (The Imitation Game): Operational definition of intelligence.

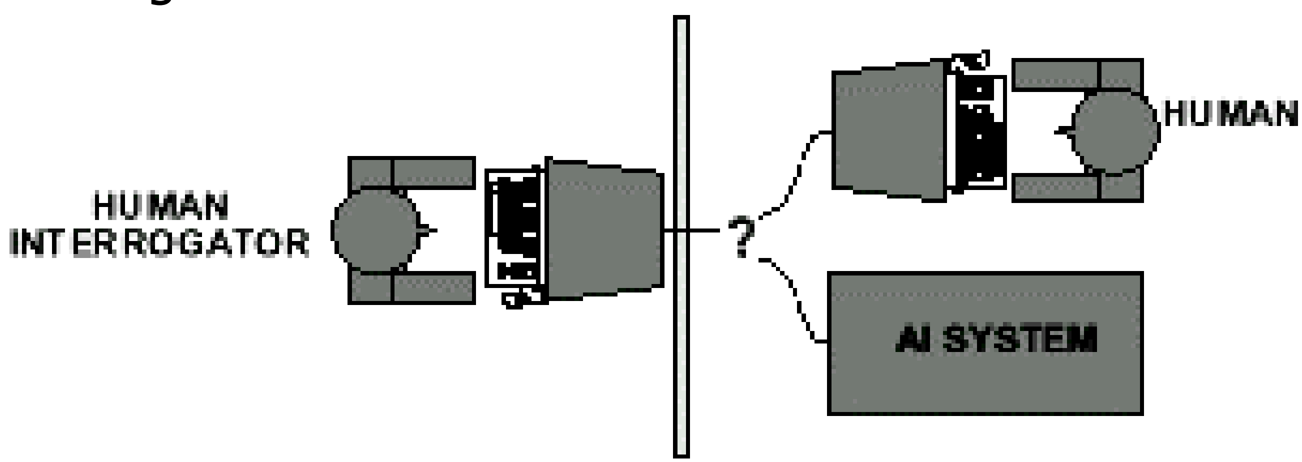

Acting Humanly: The Turing Test

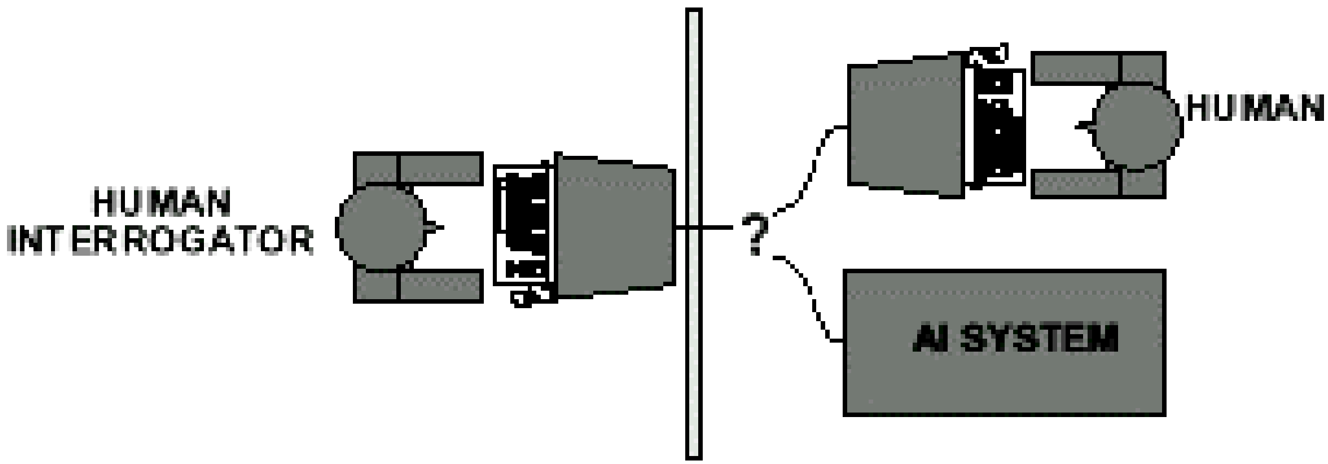

- Computer needs to possess: Natural language processing, Knowledge representation, Automated reasoning, and Machine learning

- Are there any problems/limitations to the Turing Test?

What tasks require AI?

- "AI is the science and engineering of making intelligent machines which can <u>perform tasks that require intelligence when performed by humans</u> …"

- What tasks require AI?

What tasks require AI?

- Tasks that require AI:
  - Solving a differential equation
  - Brain surgery
  - Inventing stuff
  - Playing Jeopardy
  - Playing Wheel of Fortune
  - What about walking?
  - What about grabbing stuff?
  - What about pulling your hand away from fire?
  - What about watching TV?
  - What about day dreaming?

Acting Humanly: The Full Turing Test

- Alan Turing's 1950 article *Computing Machinery and Intelligence* discussed conditions for considering a machine to be intelligent
  - "Can machines think?" ←→ "Can machines behave intelligently?"
  - The Turing test (The Imitation Game): Operational definition of intelligence.

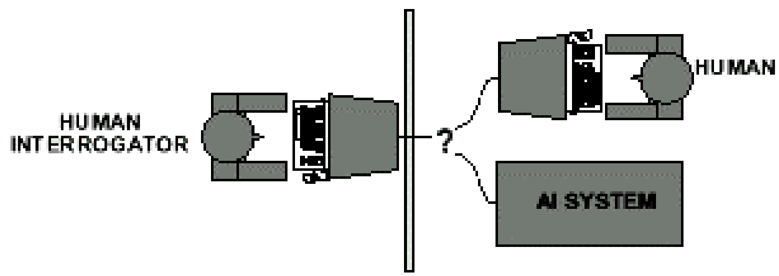

- Computer needs to posses: Natural language processing, Knowledge representation, Automated reasoning, and Machine learning
- <u>Problem:</u> 1) Turing test is not reproducible, constructive, and amenable to mathematic analysis. 2) What about physical interaction with interrogator and environment?
- Total Turing Test: Requires physical interaction and needs perception and actuation.

Acting Humanly: The Full Turing Test

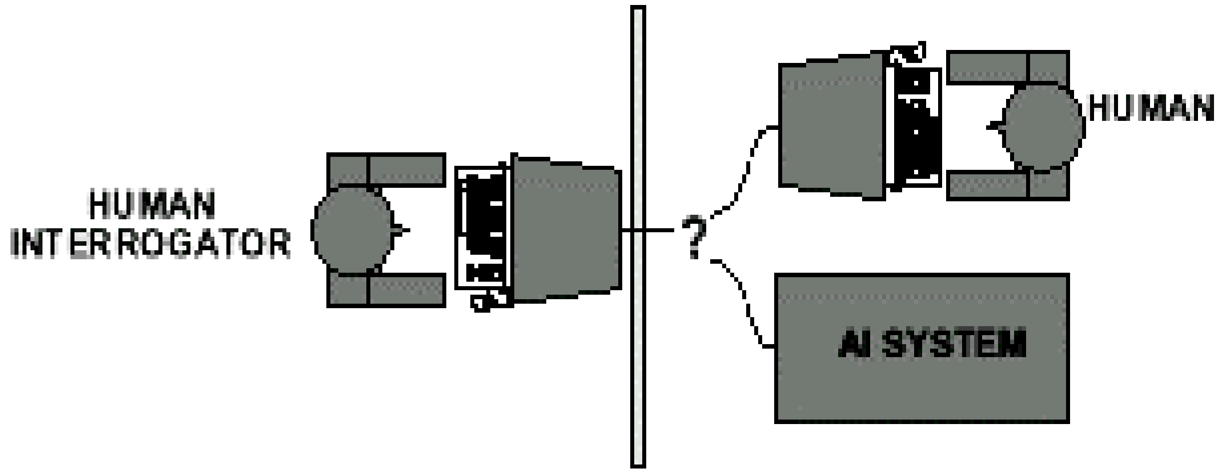

- <u>Problem:</u>

  1) Turing test is not reproducible, constructive, and amenable to mathematic analysis.

  2) What about physical interaction with interrogator and environment?

Acting Humanly: The Full Turing Test

Problem:

1) Turing test is not reproducible, constructive, and amenable to mathematic analysis.

2) What about physical interaction with interrogator and environment?

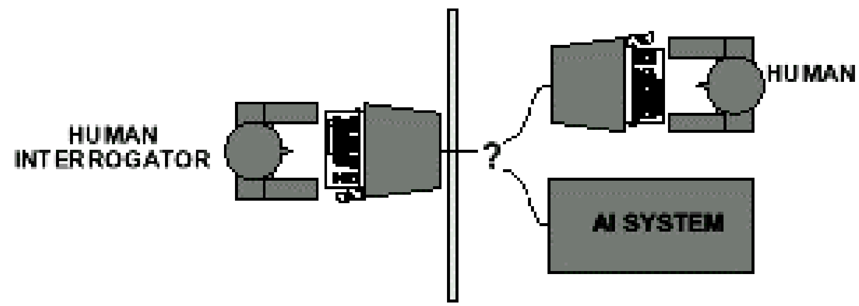

Trap door

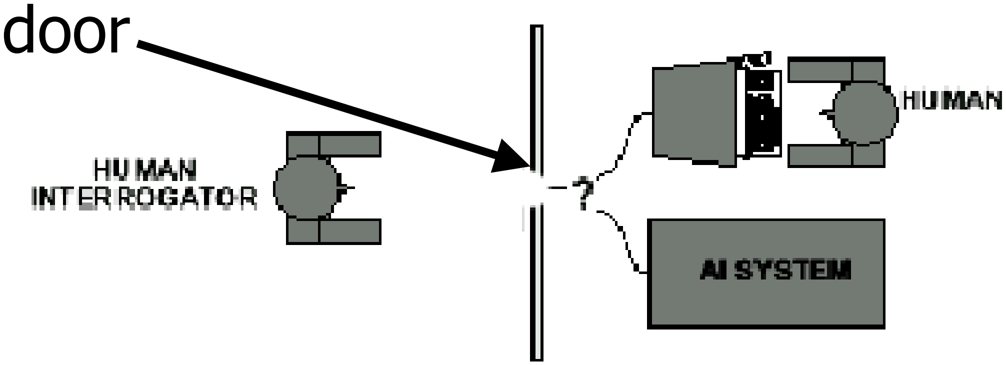

What would a computer need to pass the Turing test?

- Natural language processing: to communicate with examiner.
- Knowledge representation: to store and retrieve information provided before or during interrogation.
- Automated reasoning: to use the stored information to answer questions and to draw new conclusions.
- Machine learning: to adapt to new circumstances and to detect and extrapolate patterns.

What would a computer need to pass the Turing test?

- Vision (for Total Turing test): to recognize the examiner's actions and various objects presented by the examiner.
- Motor control (total test): to act upon objects as requested.
- Other senses (total test): such as audition, smell, touch, etc.

Thinking Humanly: Cognitive Science

- 1960 “Cognitive Revolution”: information-processing psychology replaced behaviorism

- Cognitive science brings together theories and experimental evidence to model internal activities of the brain
  - What level of abstraction? “Knowledge” or “Circuits”?
  - How to validate models?
    - Predicting and testing behavior of human subjects (top-down)
    - Direct identification from neurological data (bottom-up)
    - Building computer/machine simulated models and reproduce results (simulation)

Thinking Rationally: Laws of Thought

- Aristotle (~ 450 B.C.) attempted to codify "right thinking"
  What are correct arguments/thought processes?

- E.g., "Socrates is a man, all men are mortal; therefore Socrates is mortal"

- Several Greek schools developed various forms of logic: <u>notation plus rules of derivation for thoughts.</u>

Thinking Rationally: Laws of Thought

- <u>Problems:</u>
  1) Uncertainty: Not all facts are certain (e.g., *the flight might be delayed).*

  2) Resource limitations:
     - Not enough time to compute/process
     - Insufficient memory/disk/etc
     - Etc.

Acting Rationally: The Rational Agent

- Rational behavior: Doing the right thing!
- The right thing: That which is expected to maximize the expected return
- Provides the most general view of AI because it includes:
  - Correct inference ("Laws of thought")
  - Uncertainty handling
  - Resource limitation considerations (e.g., reflex vs. deliberation)
  - Cognitive skills (NLP, AR, knowledge representation, ML, etc.)
- <u>Advantages:</u>
  1) More general
  2) Its goal of rationality is well defined

How to achieve AI?

- How is AI research done?
- AI research has both theoretical and experimental sides. The experimental side has both basic and applied aspects.
- There are two main lines of research:
  - One is biological, based on the idea that since humans are intelligent, AI should study humans and imitate their psychology or physiology.
  - The other is phenomenal, based on studying and formalizing common sense facts about the world and the problems that the world presents to the achievement of goals.
- The two approaches interact to some extent, and both should eventually succeed. It is a race, but both racers seem to be walking. [**John McCarthy]**

Branches of AI

- **Logical AI**
- **Search**
- **Natural language processing**
- **pattern recognition**
- **Knowledge representation**
- **Inference** From some facts, others can be inferred.
- **Automated reasoning**
- **Learning from experience**
- **Planning** To generate a strategy for achieving some goal
- **Epistemology** Study of the kinds of knowledge that are required for solving problems in the world.
- **Ontology** Study of the kinds of things that exist. In AI, the programs and sentences deal with various kinds of objects, and we study what these kinds are and what their basic properties are.
- **Genetic programming**
- **Emotions???**
- …

AI Prehistory

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Linguistics	knowledge representation grammar
Neuroscience	physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

AI History

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952–69	Look, Ma, no hands!
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1969–79	Early development of knowledge-based systems
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probabilistic and decision-theoretic methods Rapid increase in technical depth of mainstream AI "Nouvelle AI": ALife, GAs, soft computing

AI State of the art

- Have the following been achieved by AI?
  - World-class chess playing
  - Playing table tennis
  - Cross-country driving
  - Solving mathematical problems
  - Discover and prove mathematical theories
  - Engage in a meaningful conversation
  - Understand spoken language
  - Observe and understand human emotions
  - Express emotions
  - …

Course Overview

General Introduction

- ***01-Introduction.*** [AIMA Ch 1] Course Schedule. Homeworks, exams and grading. Course material, TAs and office hours. Why study AI? What is AI? The Turing test. Rationality. Branches of AI. Research disciplines connected to and at the foundation of AI. Brief history of AI. Challenges for the future. Overview of class syllabus.

- ***02-Intelligent Agents.*** [AIMA Ch 2] What is an intelligent agent? Examples. Doing the right thing (rational action). Performance measure. Autonomy. Environment and agent design. Structure of agents. Agent types. Reflex agents. Reactive agents. Reflex agents with state. Goal-based agents. Utility-based agents. Mobile agents. Information agents.

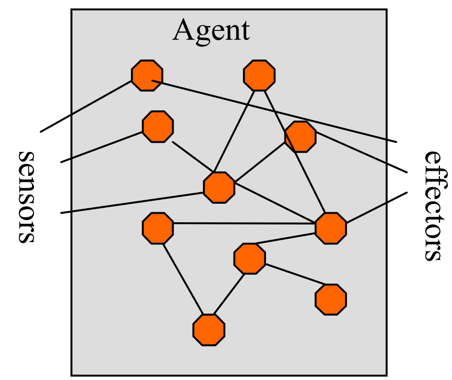

Course Overview (cont.)

How can we solve complex problems?

- ***03/04-Problem solving and search.*** [AIMA Ch 3] Example: measuring problem. Types of problems. More example problems. Basic idea behind search algorithms. Complexity. Combinatorial explosion and NP completeness. Polynomial hierarchy.

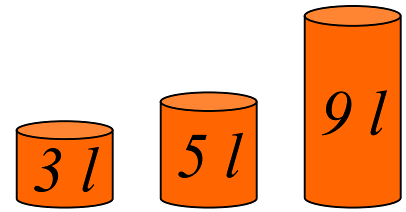

Using these 3 buckets, measure 7 liters of water.

- ***05-Uninformed search.*** [AIMA Ch 3] Depth-first. Breadth-first. Uniform-cost. Depth-limited. Iterative deepening. Examples. Properties.

- ***06/07-Informed search.*** [AIMA Ch 4] Best-first. A* search. Heuristics. Hill climbing. Problem of local extrema. Simulated annealing.

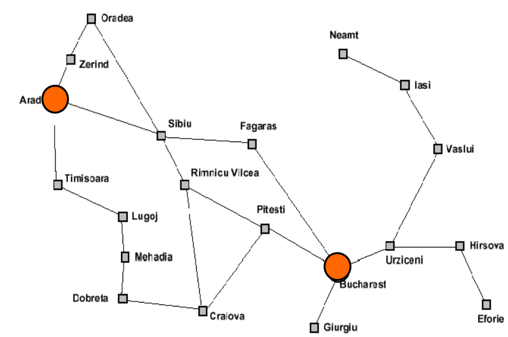

Traveling salesperson problem

Course Overview (cont.)

Practical applications of search.

- ***08/09-Game playing.*** [AIMA Ch 5] The minimax algorithm. Resource limitations. Aplha-beta pruning. Elements of chance and non-deterministic games.

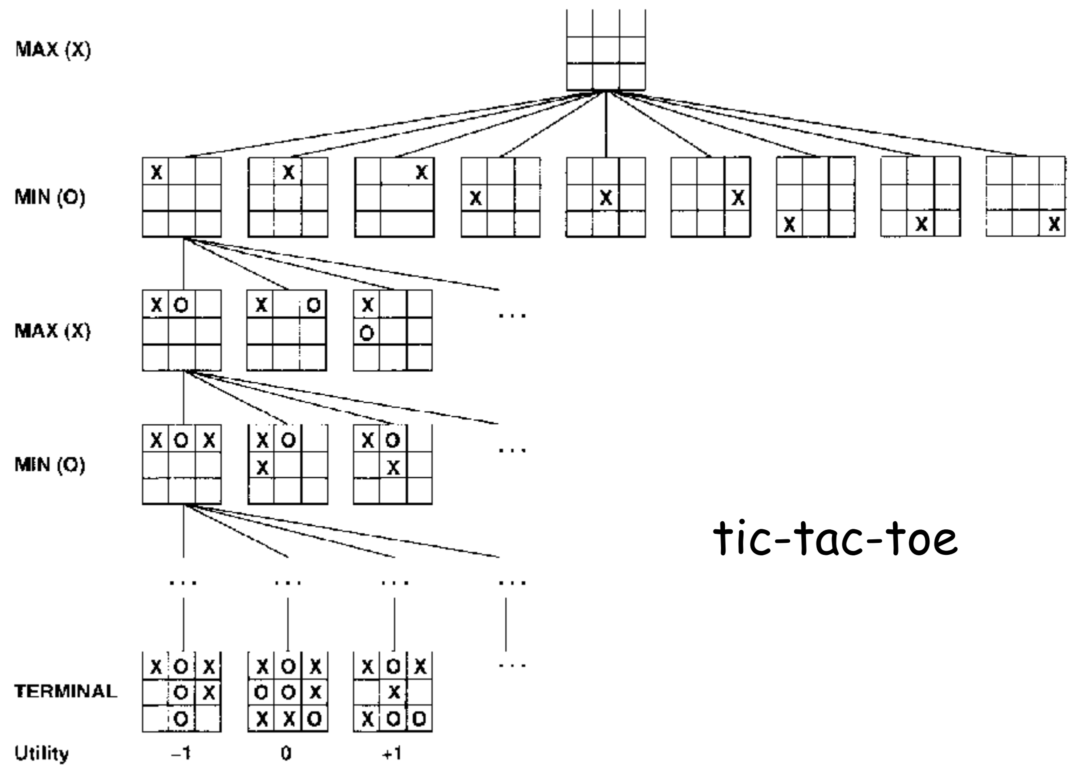

tic-tac-toe

Course Overview (cont.)

Towards intelligent agents

- ***10-Agents that reason logically 1.*** [AIMA Ch 6] Knowledge-based agents. Logic and representation. Propositional (boolean) logic.
- ***11-Agents that reason logically 2.*** [AIMA Ch 6] Inference in propositional logic. Syntax. Semantics. Examples.

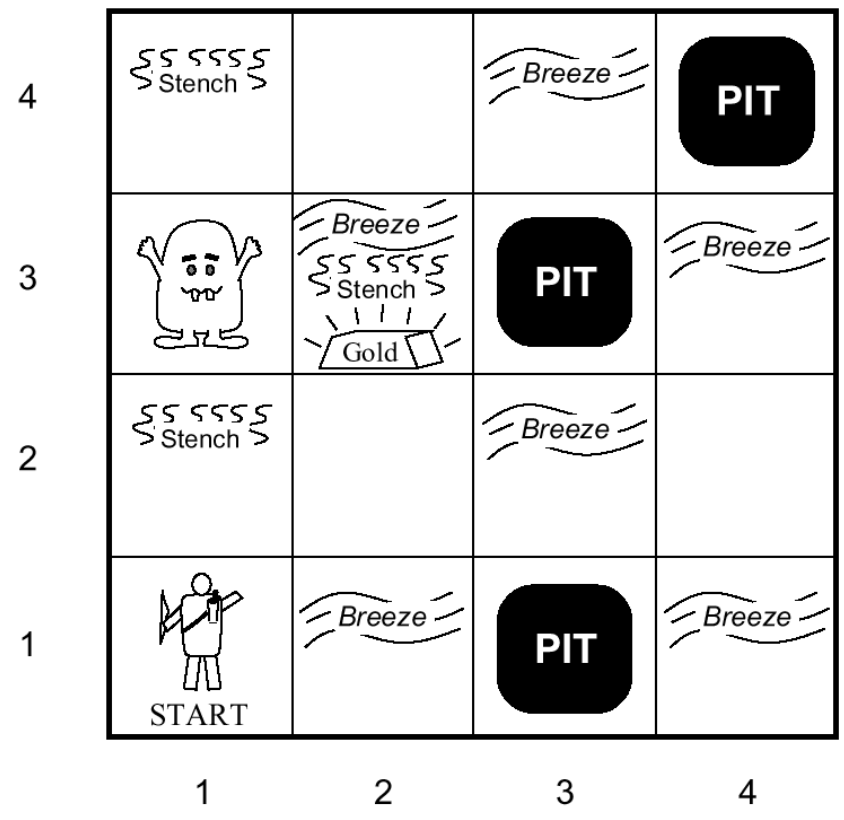

wumpus world

Course Overview (cont.)

Building knowledge-based agents: 1st Order Logic

- ***12-First-order logic 1.*** [AIMA Ch 7] Syntax. Semantics. Atomic sentences. Complex sentences. Quantifiers. Examples. FOL knowledge base. Situation calculus.

- ***13-First-order logic 2.*** [AIMA Ch 7] Describing actions. Planning. Action sequences.

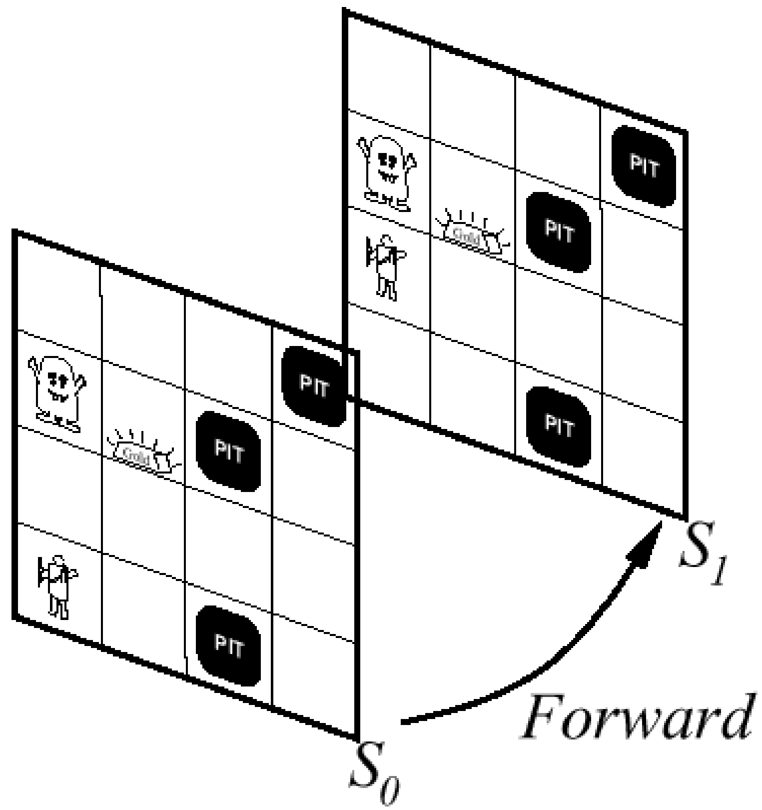

Course Overview (cont.)

Representing and Organizing Knowledge

- ***14/15-Building a knowledge base.*** [AIMA Ch 8] Knowledge bases. Vocabulary and rules. Ontologies. Organizing knowledge.

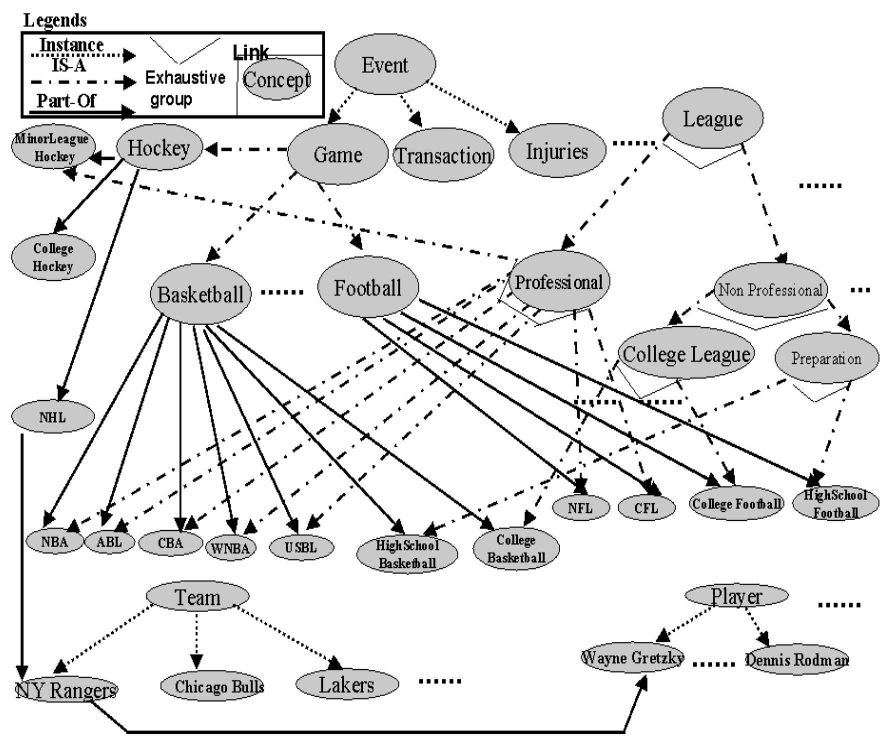

Kahn & Mcleod, 2000

An ontology for the sports domain

Course Overview (cont.)

Reasoning Logically

- ***16/17/18-Inference in first-order logic.*** [AIMA Ch 9] Proofs. Unification. Generalized modus ponens. Forward and backward chaining.

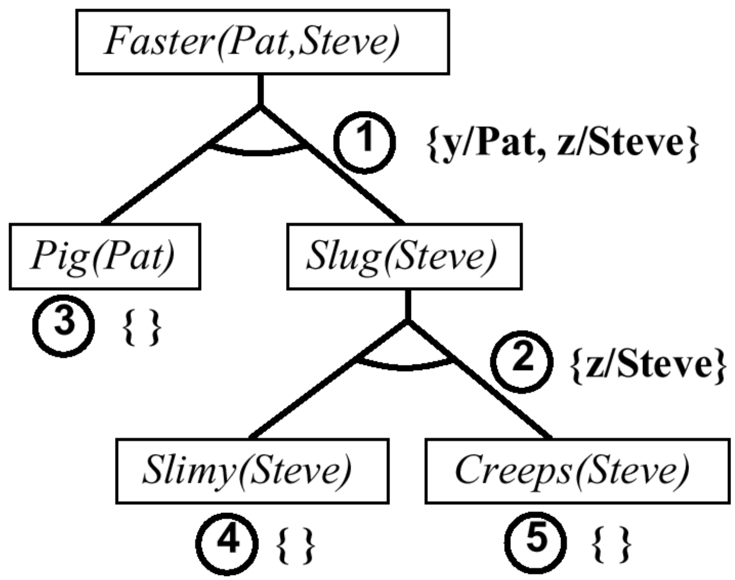

Example of backward chaining

Course Overview (cont.)

Examples of Logical Reasoning Systems

- ***19-Logical reasoning systems.*** [AIMA Ch 10] Indexing, retrieval and unification. The Prolog language. Theorem provers. Frame systems and semantic networks.

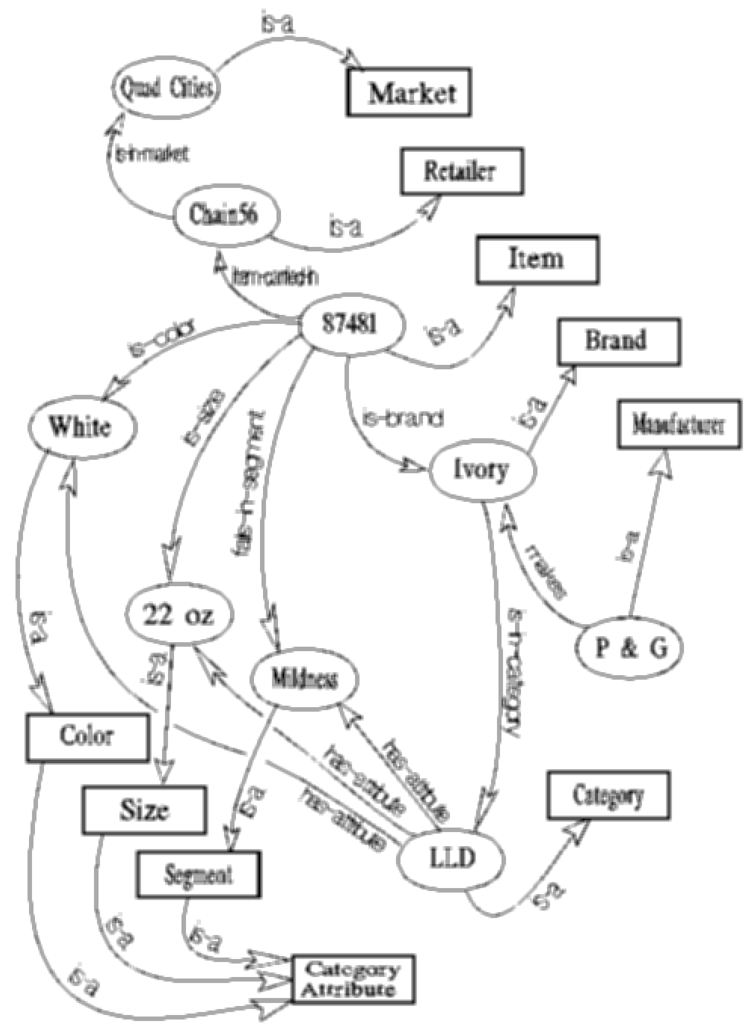

Semantic network used in an insight generator (Duke university)

Course Overview (cont.)

Systems that can Plan Future Behavior

- ***20-Planning.*** [AIMA Ch 11] Definition and goals. Basic representations for planning. Situation space and plan space. Examples.

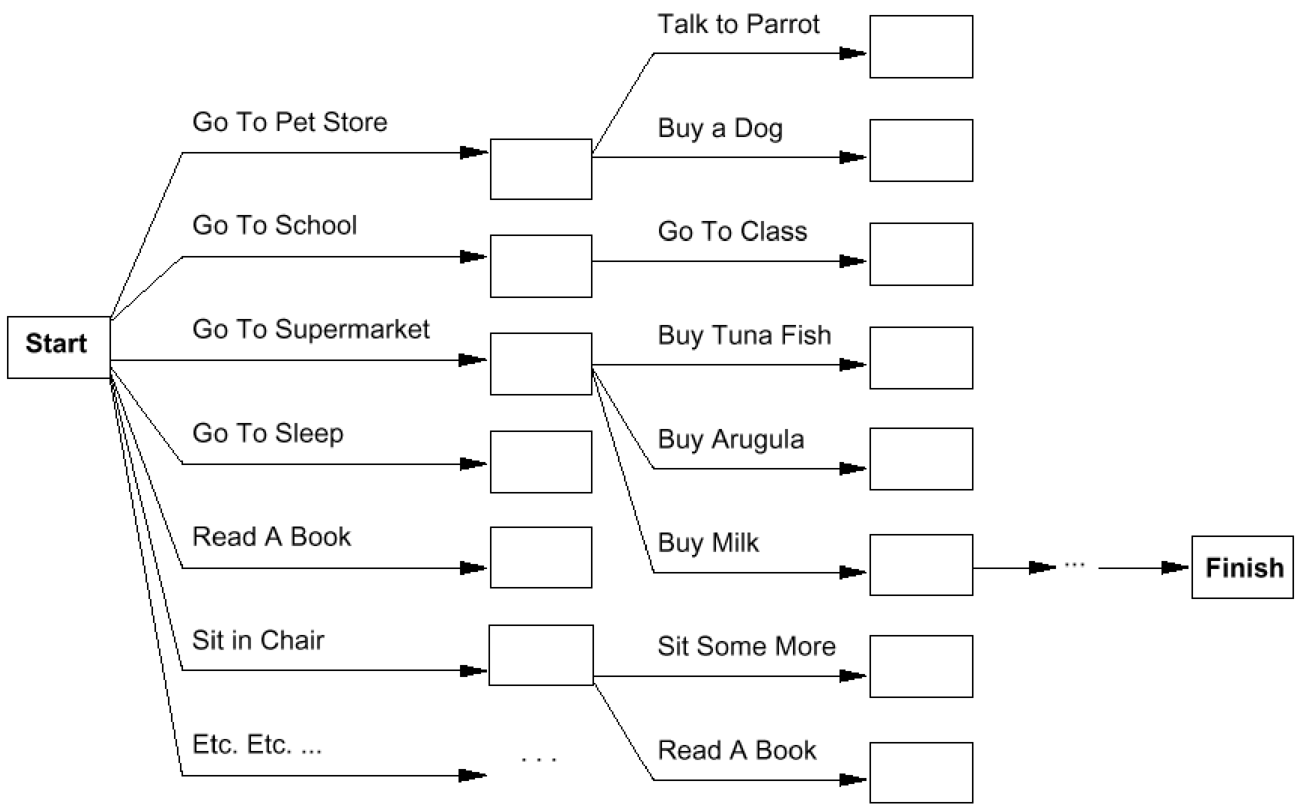

Course Overview (cont.)

Expert Systems

- ***21-Introduction to CLIPS.*** [handout]
  Overview of modern rule-based expert systems. Introduction to CLIPS (C Language Integrated Production System). Rules. Wildcards. Pattern matching. Pattern network. Join network.

```
CLIPS> (clear)
CLIPS> (assert (animal-is duck))
<Fact-0>
CLIPS> (assert (animal-sound quack))
<Fact-1>
CLIPS> (assert (The duck says "Quack."))
<Fact-2>
CLIPS> (facts)
f-0     (animal-is duck)
f-1     (animal-sound quack)
f-2     (The duck says "Quack.")
For a total of 3 facts.
CLIPS>
```

CLIPS expert system shell

Course Overview (cont.)

Logical Reasoning in the Presence of Uncertainty

- ***22/23-Fuzzy logic.***
  [Handout] Introduction to fuzzy logic. Linguistic Hedges. Fuzzy inference. Examples.

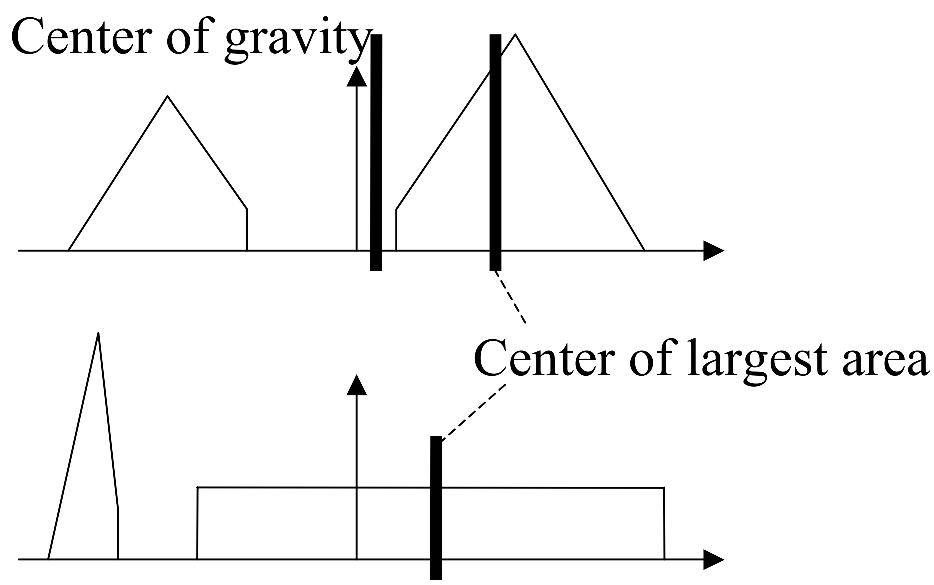

Course Overview (cont.)

AI with Neural networks

- ***24/25-Neural Networks.***
  [Handout] Introduction to perceptrons, Hopfield networks, self-organizing feature maps. How to size a network? What can neural networks achieve?

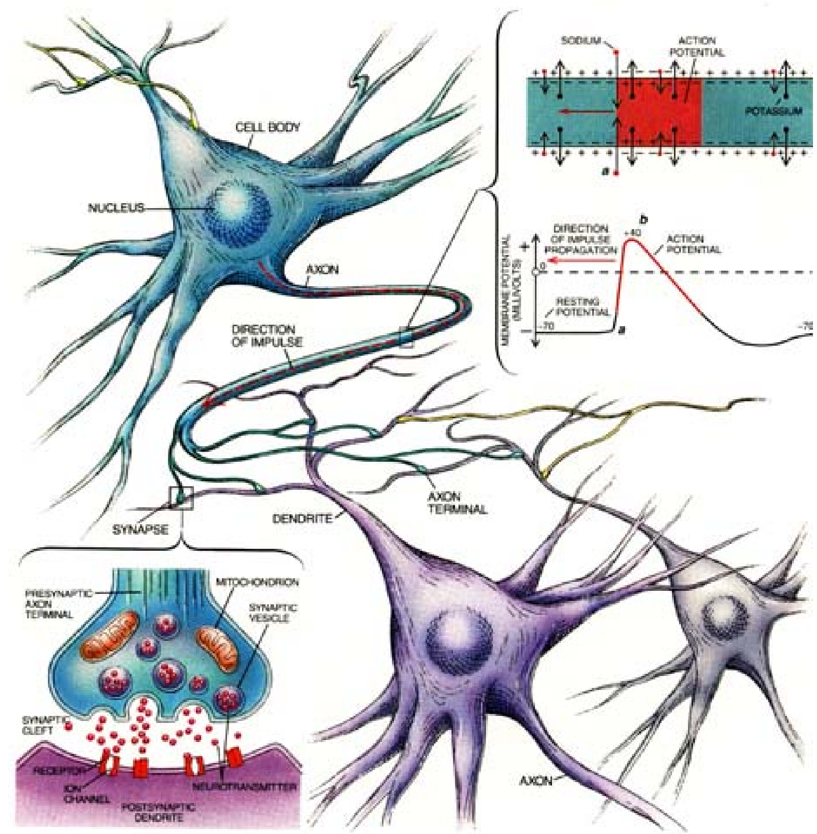

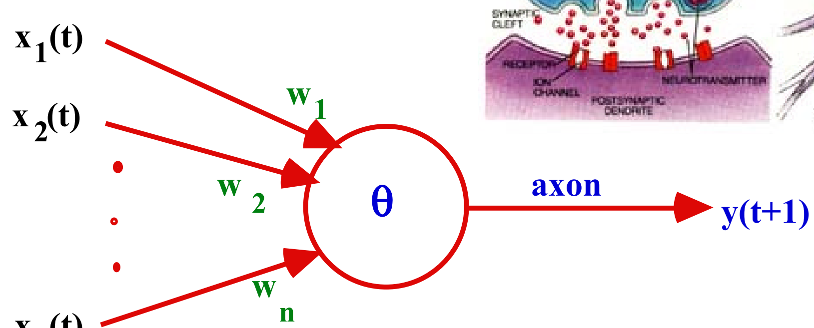

Course Overview (cont.)

Evolving Intelligent Systems

- ***26-Genetic Algorithms.*** [Handout] Introduction to genetic algorithms and their use in optimization problems.

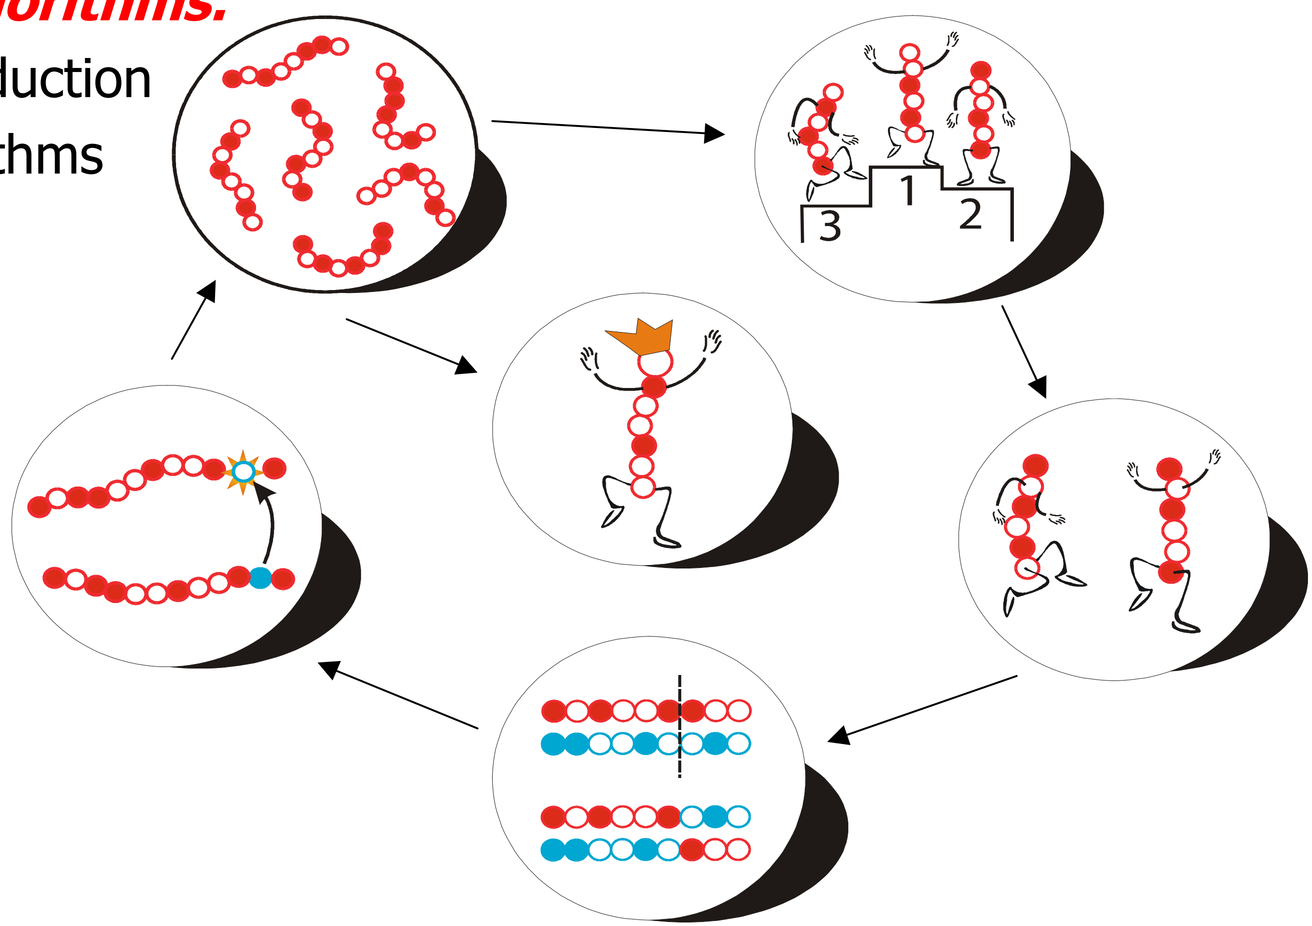

Course Overview (cont.)

What challenges remain?

- ***27-Towards intelligent machines.*** [AIMA Ch 25] The challenge of robots: with what we have learned, what hard problems remain to be solved? Different types of robots. Tasks that robots are for. Parts of robots. Architectures. Configuration spaces. Navigation and motion planning. Towards highly-capable robots.
- ***28-Overview and summary.*** [all of the above] What have we learned. Where do we go from here?

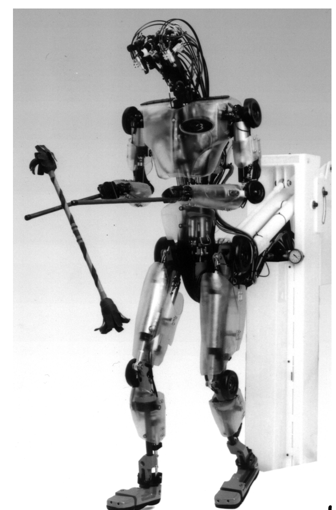

A driving example: Beobots

- **Goal:** build robots that can operate in unconstrained environments and that can solve a wide variety of tasks.

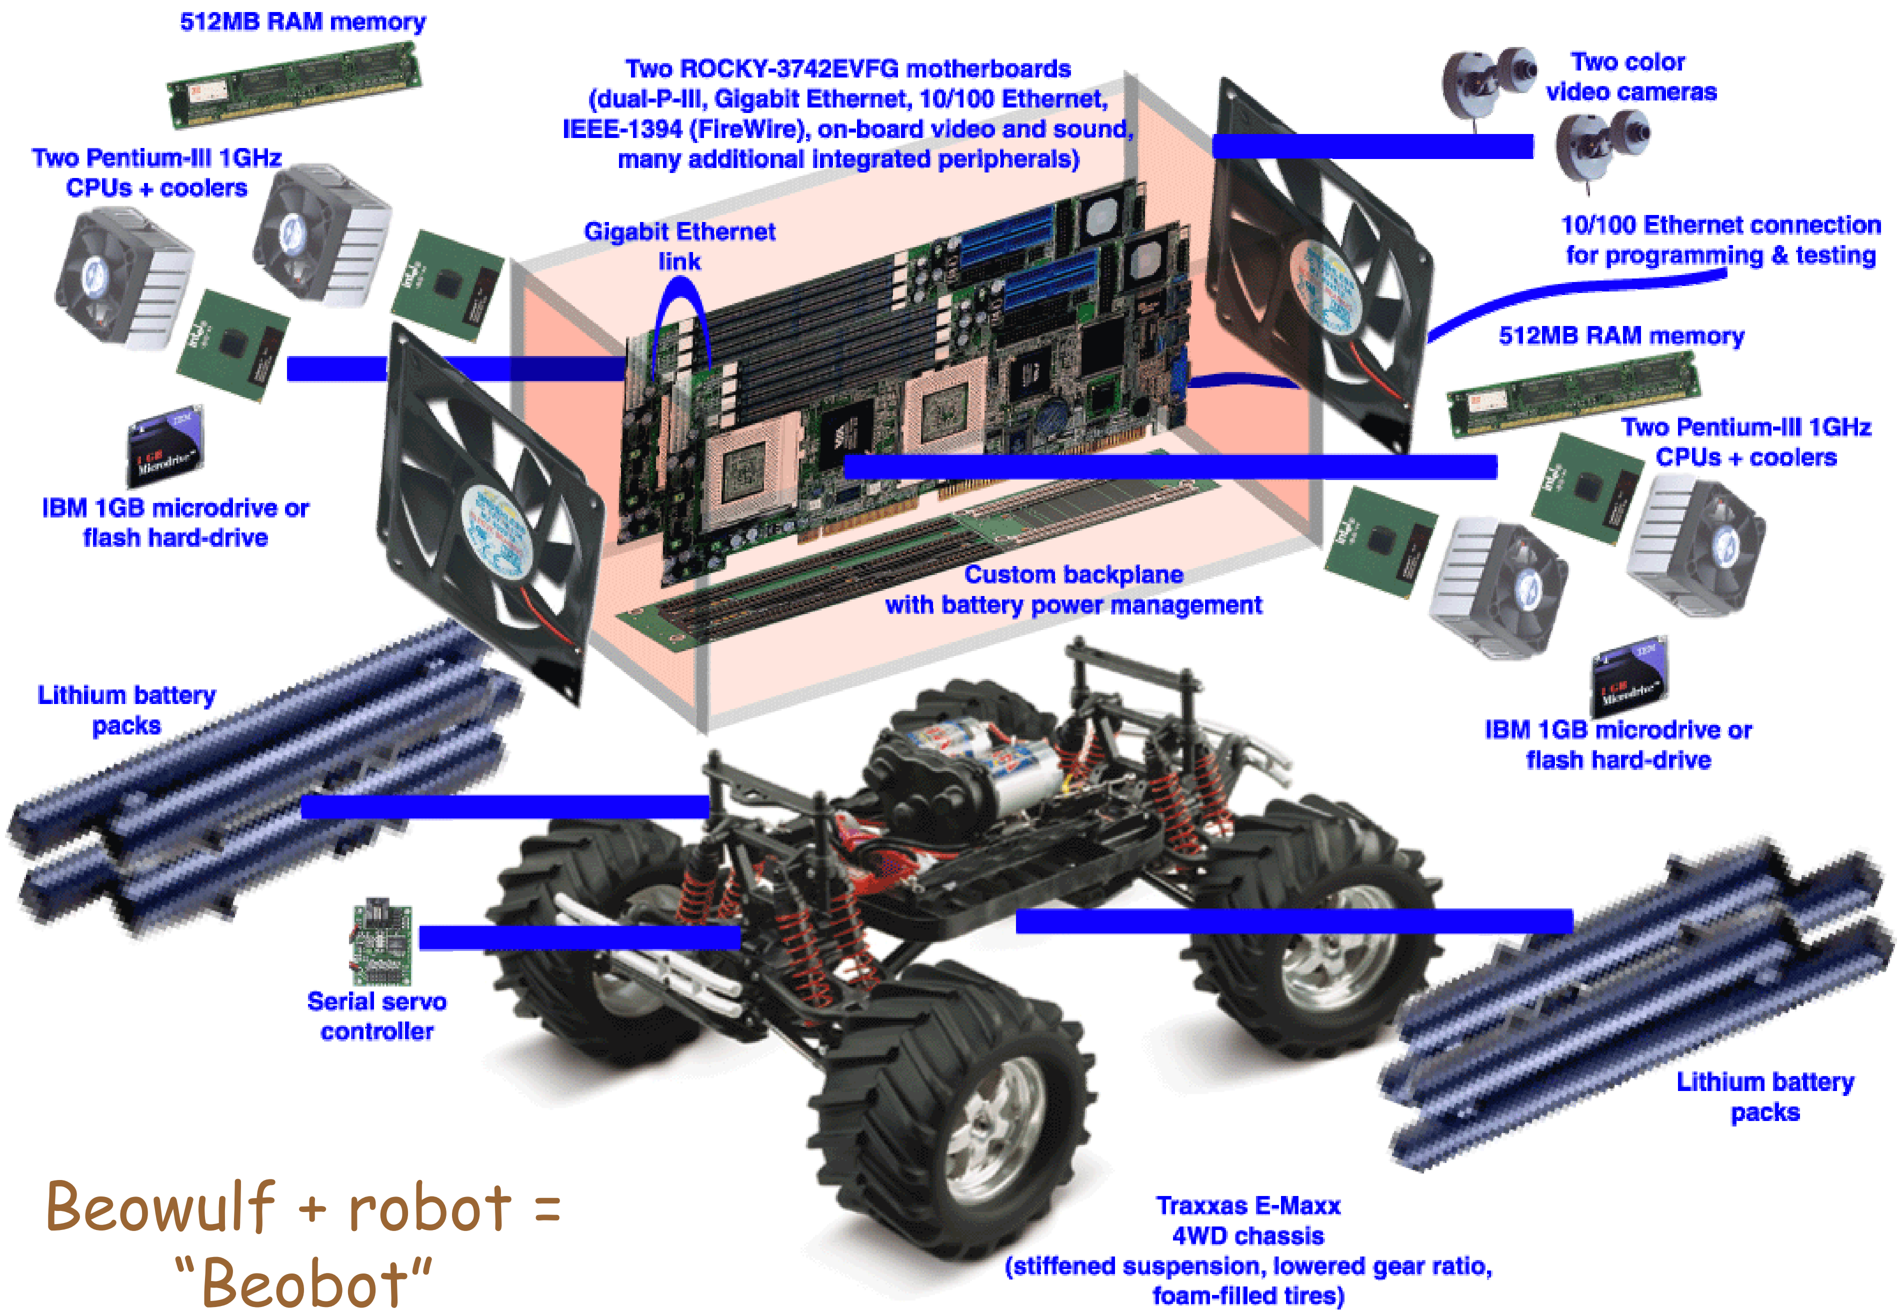

Beowulf + robot =
"Beobot"

A driving example: Beobots

- **Goal:** build robots that can operate in unconstrained environments and that can solve a wide variety of tasks.

- **We have:**
  - Lots of CPU power
  - Prototype robotics platform
  - Visual system to find interesting objects in the world
  - Visual system to recognize/identify some of these objects
  - Visual system to know the type of scenery the robot is in

- **We need to:**
  - Build an internal representation of the world
  - Understand what the user wants
  - Act upon user requests / solve user problems

The basic components of vision

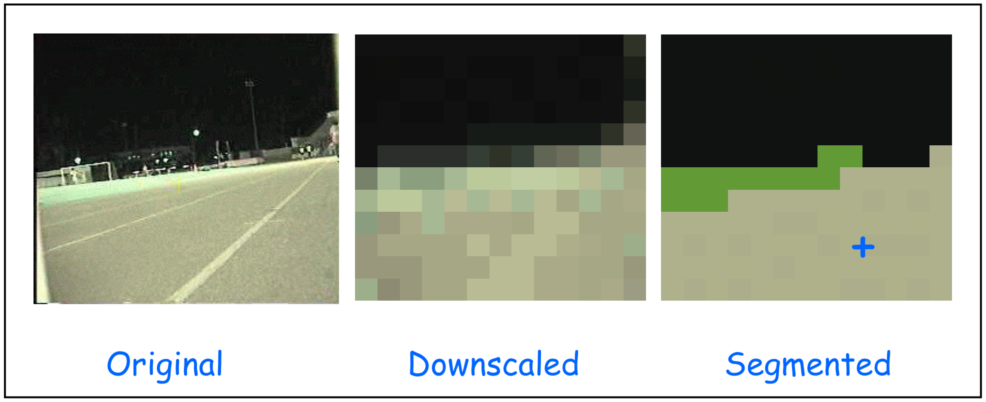

Scene Layout
& Gist

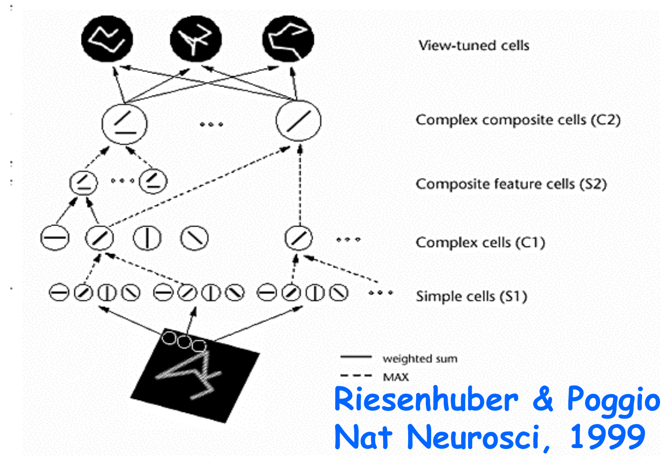

Riesenhuber & Poggio, Nat Neurosci, 1999

Localized Object Recognition

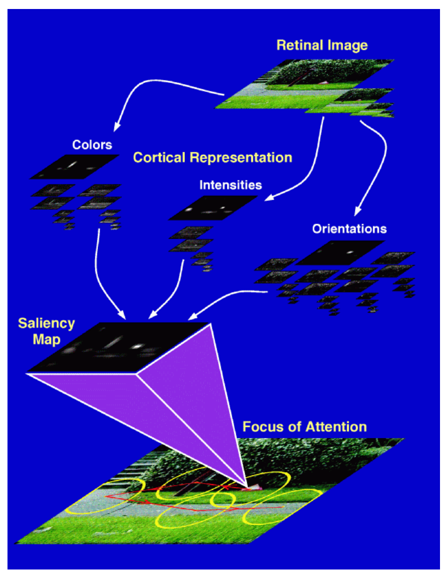

Attention

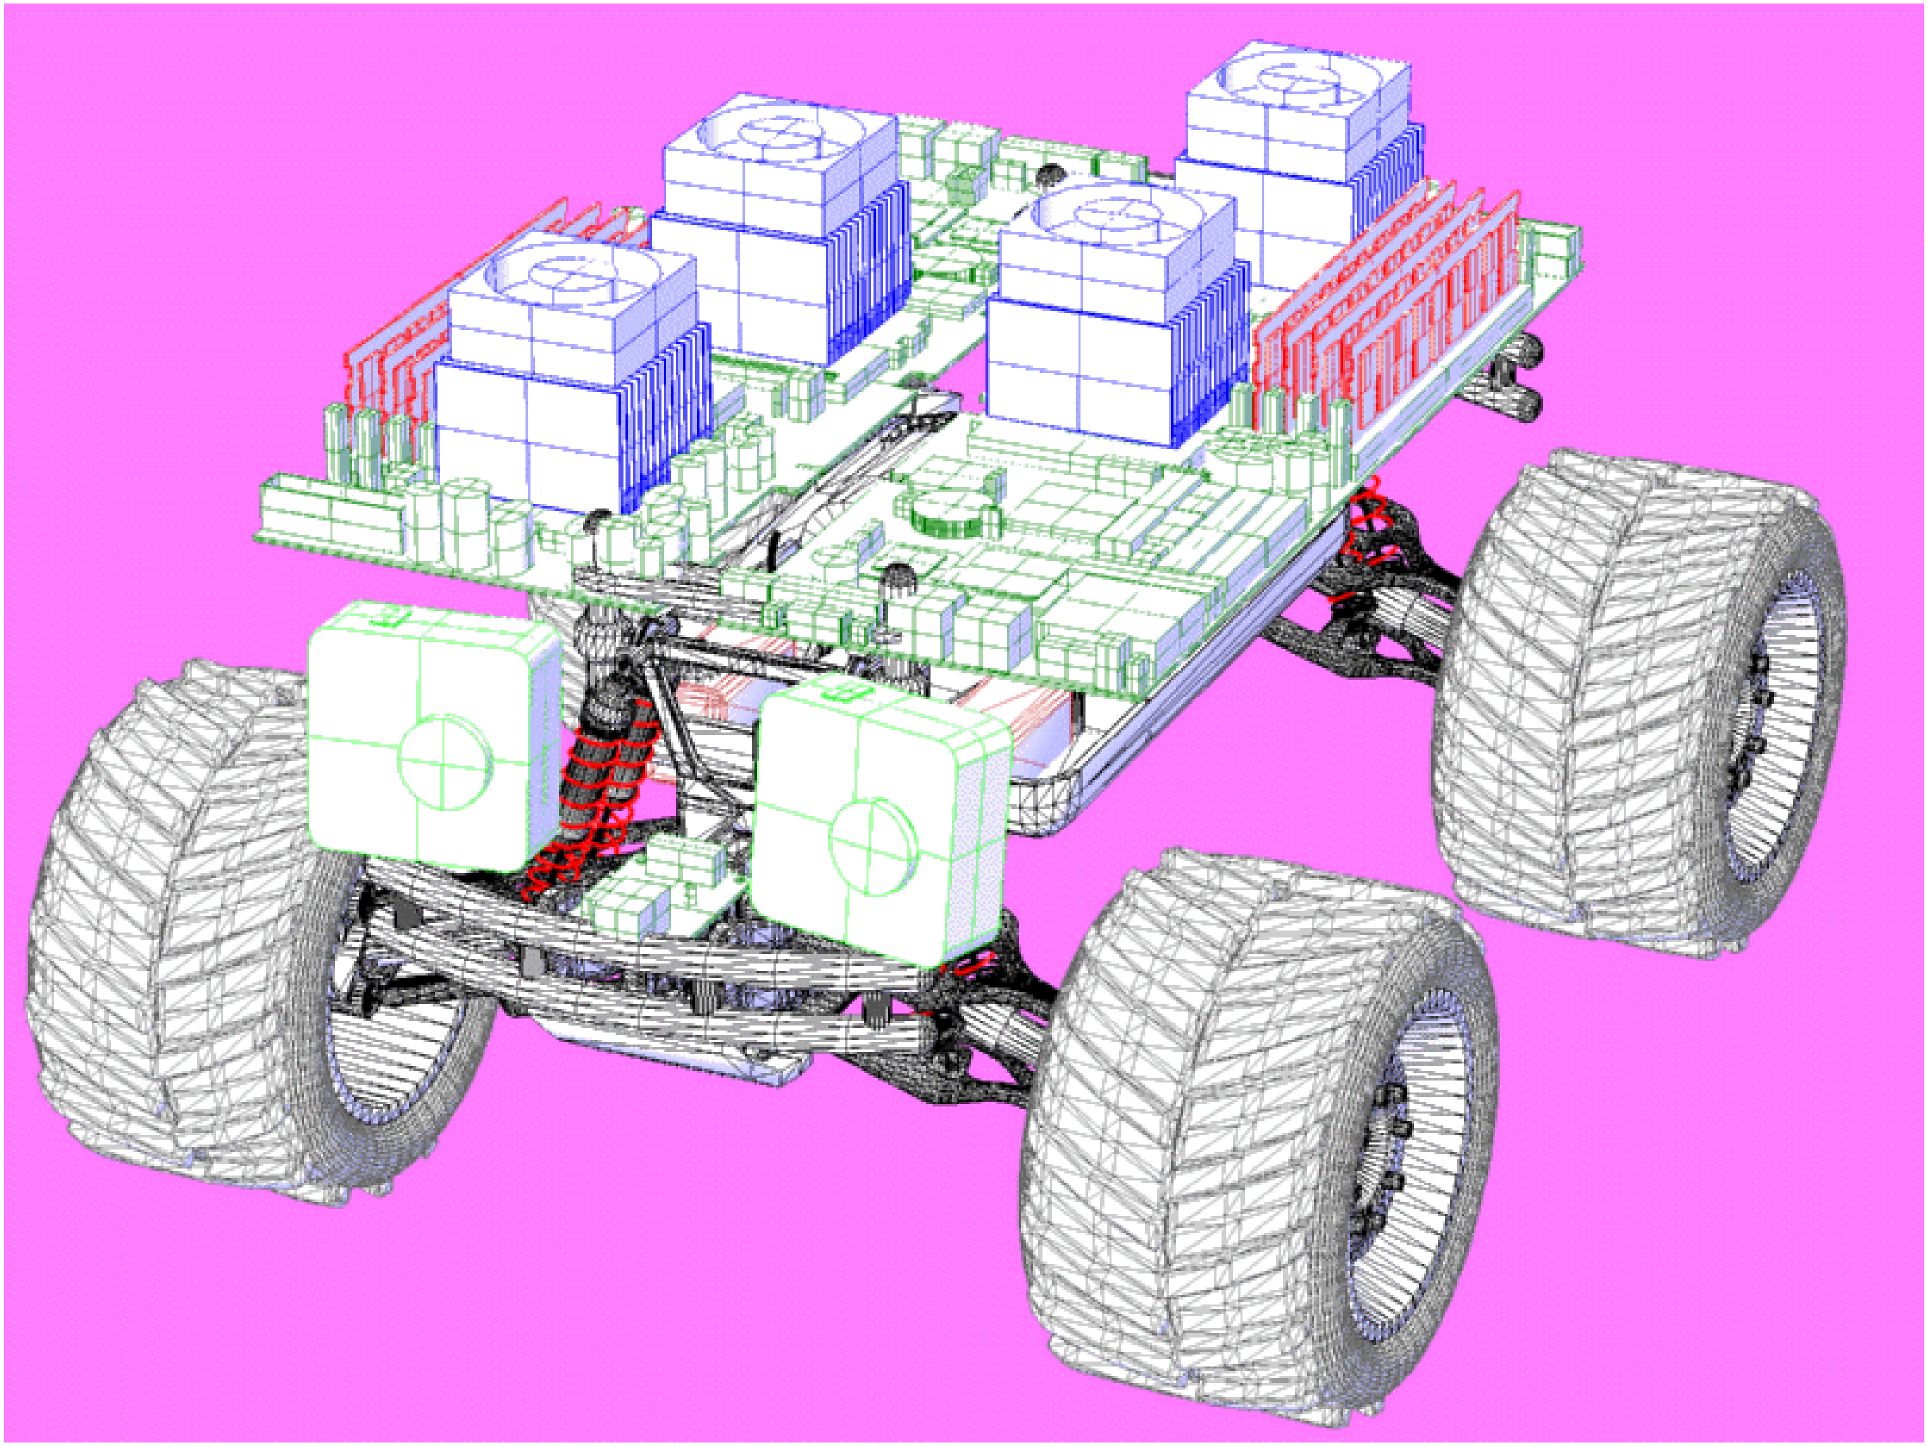

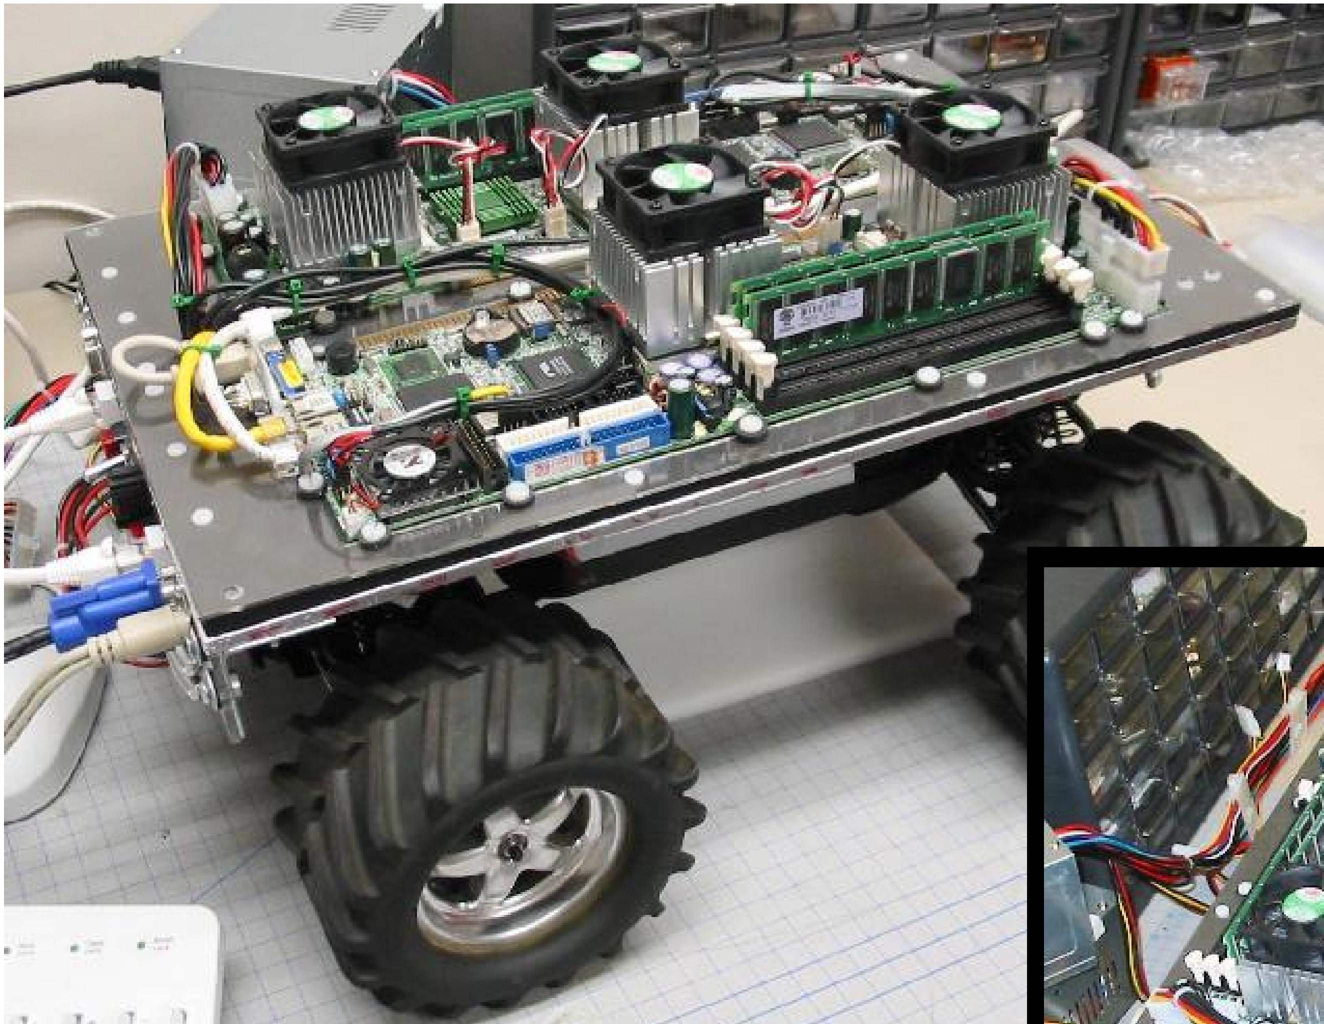

Beowulf + Robot =
"Beobot"

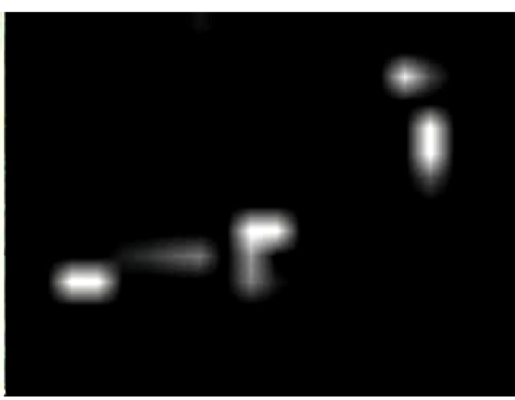

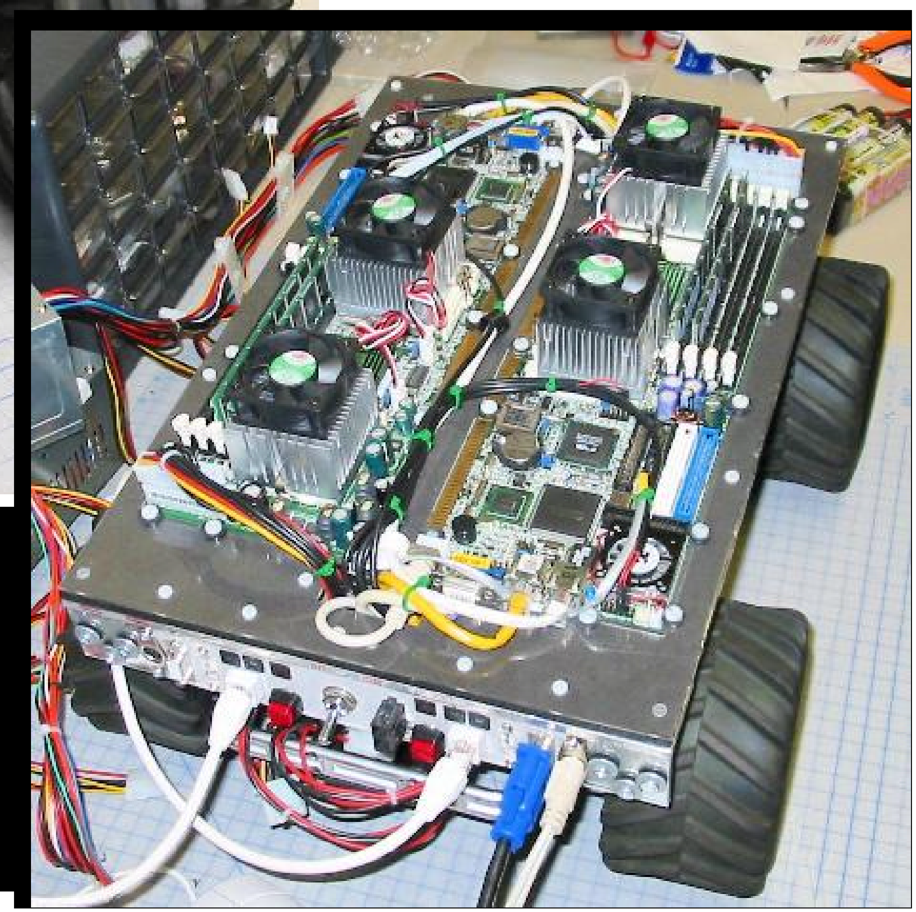

Main challenge: extract the "minimal subscene" (i.e., small number of objects and actions) that is relevant to present behavior from the noisy attentional scanpaths.

Achieve representation for it that is robust and stable against noise, world motion, and egomotion.

Prototype

Stripped-down version of proposed general system, for simplified goal: drive around USC olympic track, avoiding obstacles

Operates at 30fps on quad-CPU Beobot;

Layout & saliency very robust;

Object recognition often confused by background clutter.

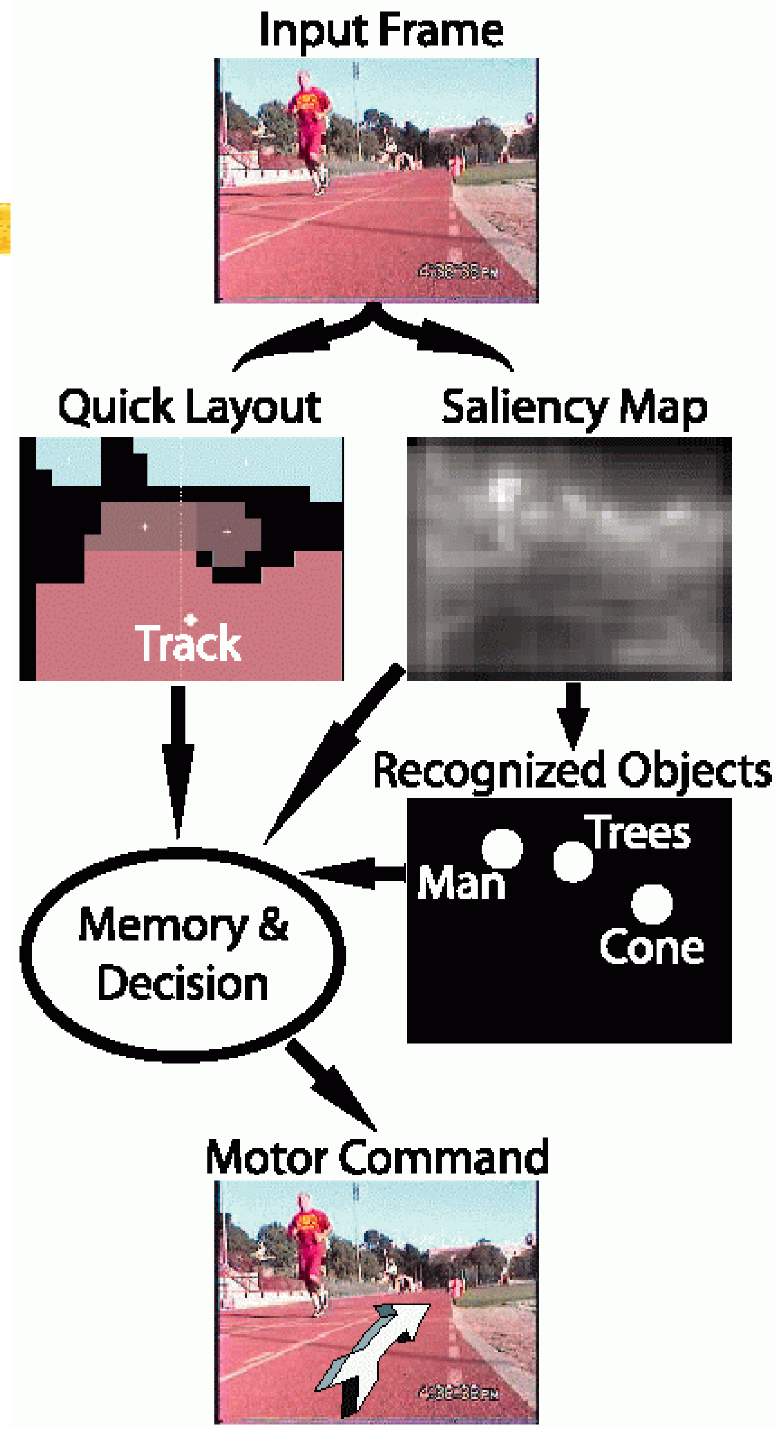

Major issues

- How to represent knowledge about the world?

- How to react to new perceived events?
- How to integrate new percepts to past experience?

- How to understand the user?
- How to optimize balance between user goals & environment constraints?
- How to use reasoning to decide on the best course of action?
- How to communicate back with the user?

- How to plan ahead?
- How to learn from experience?

General architecture

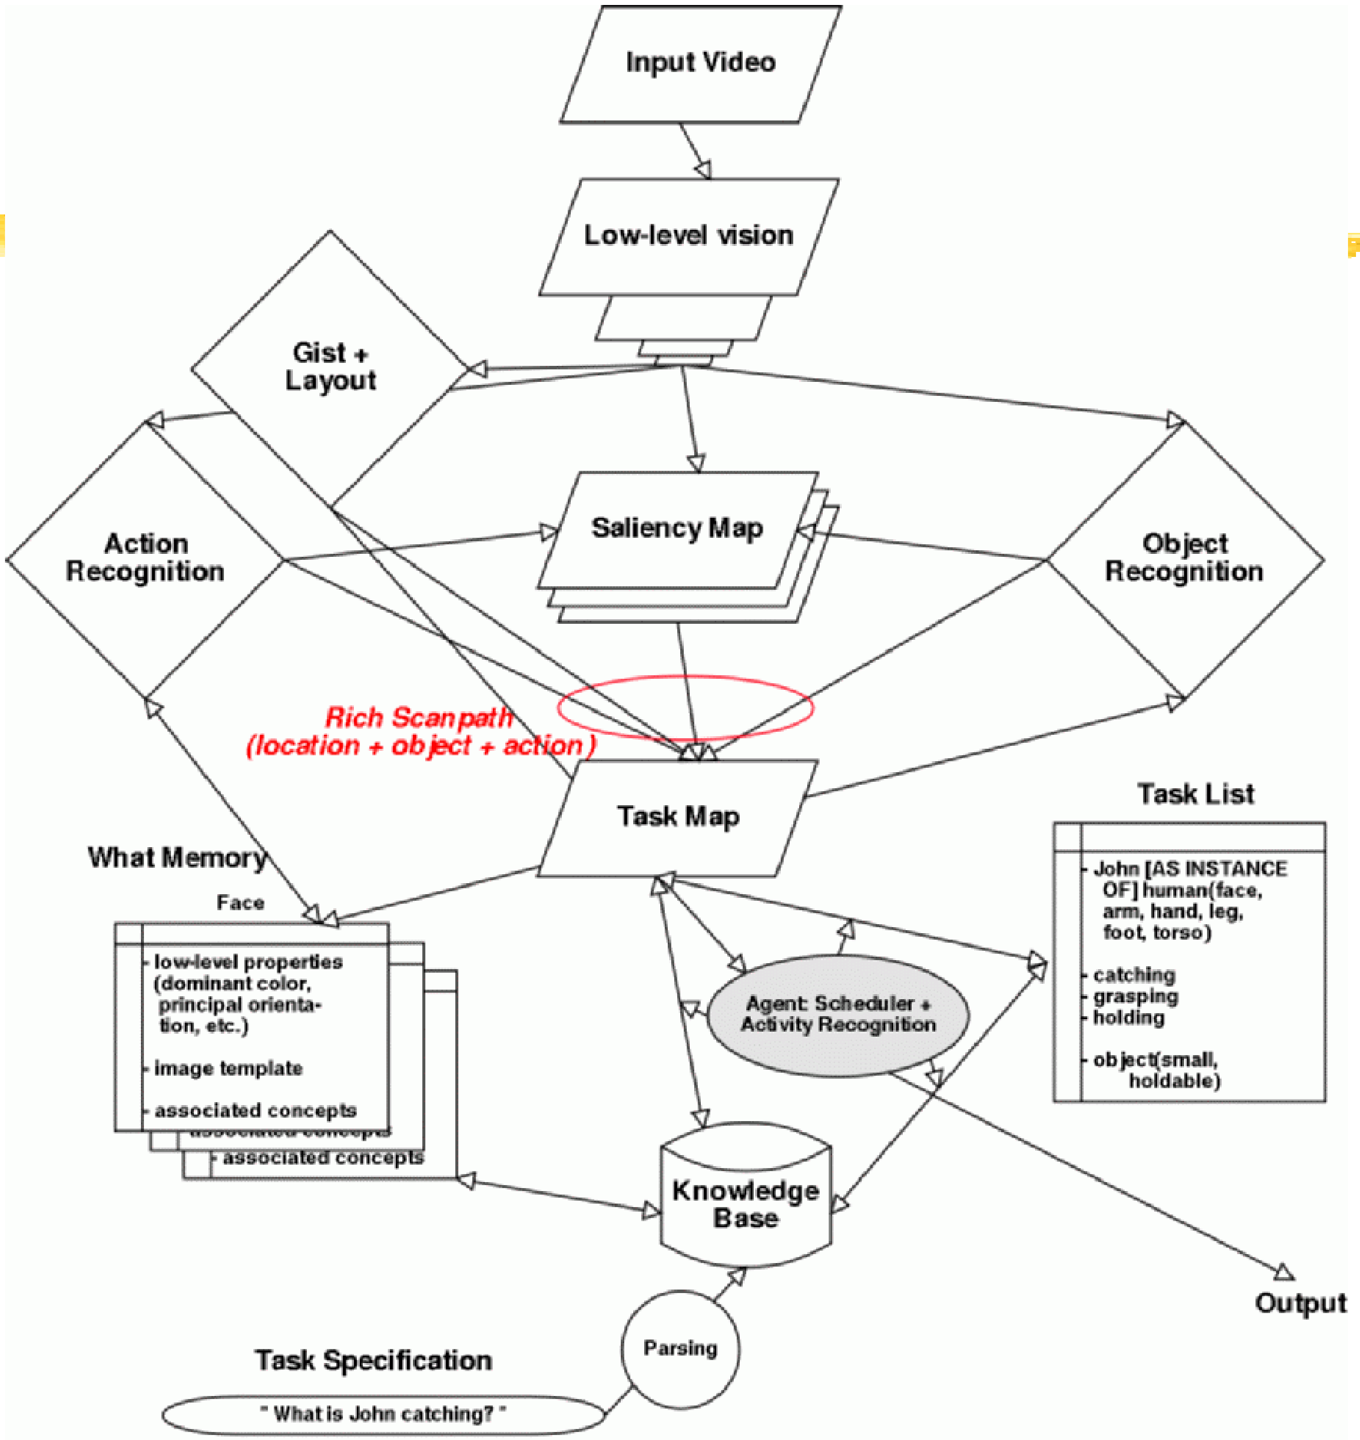

Ontology

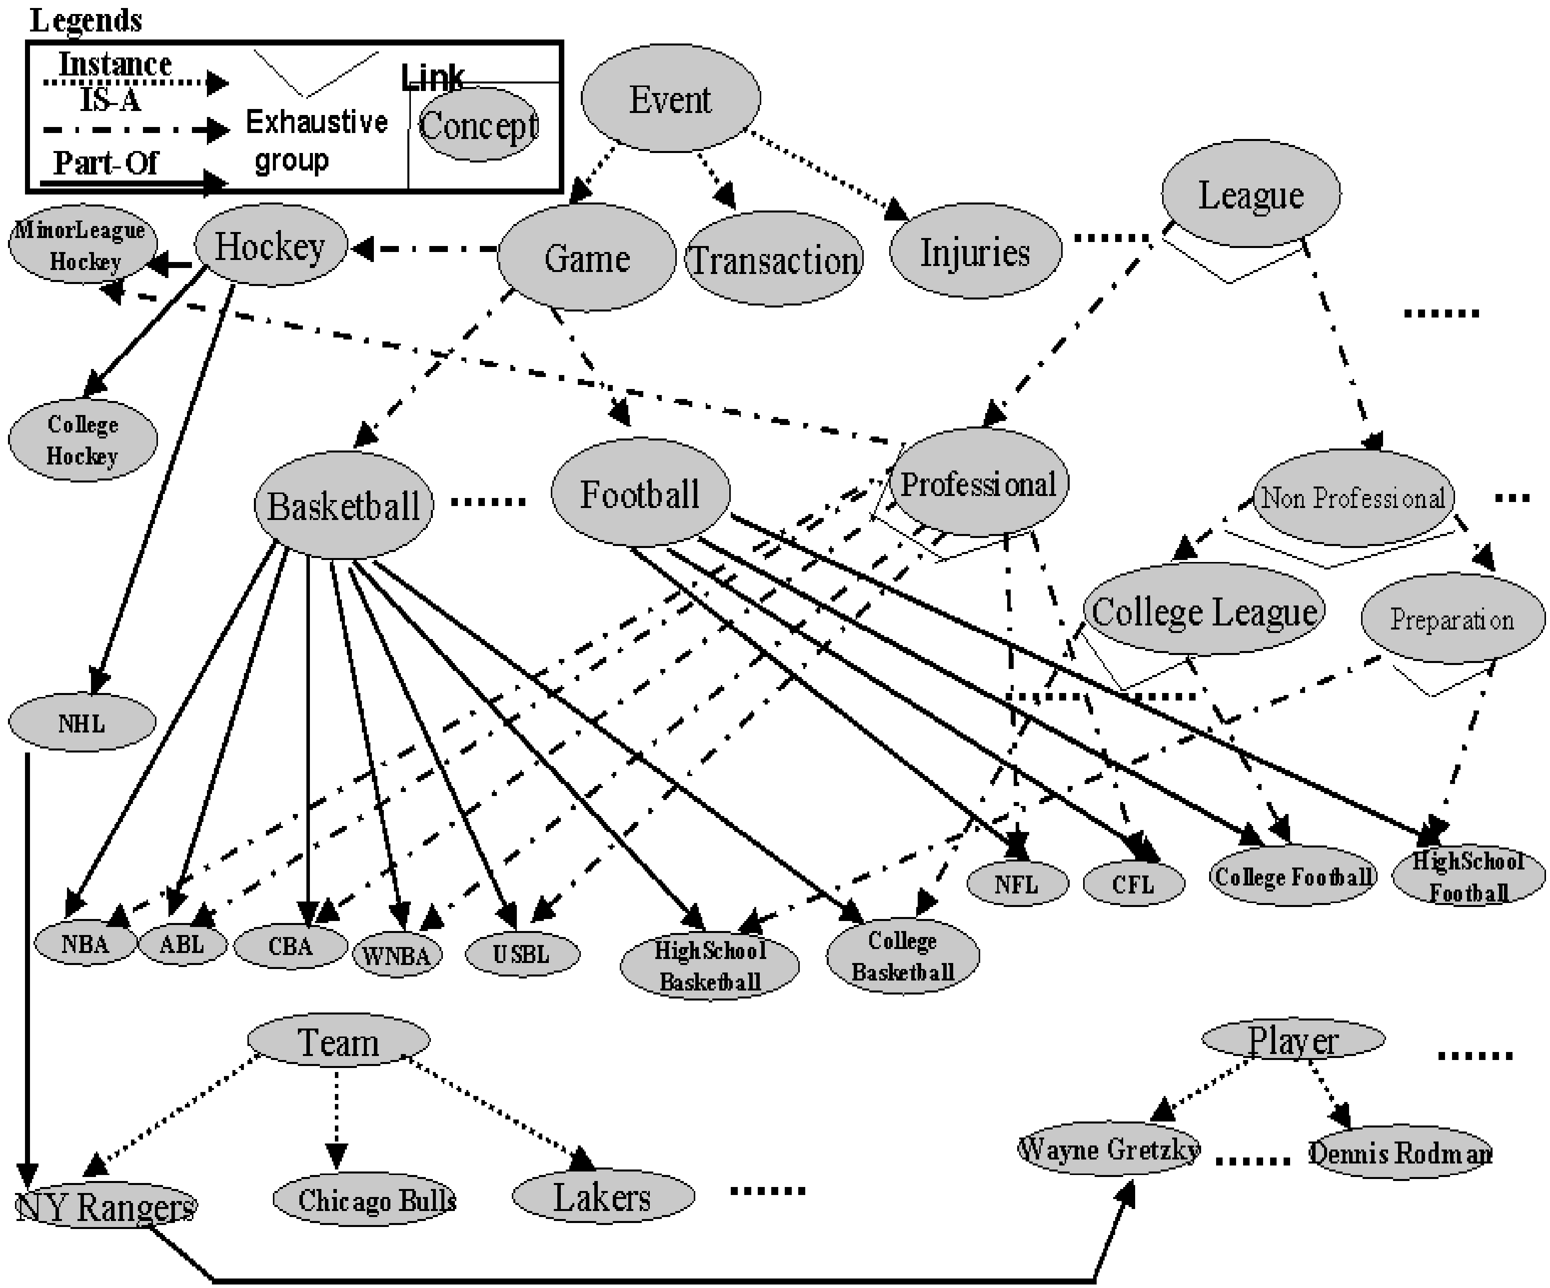

Khan & McLeod, 2000

The task-relevance map

Scalar topographic map, with higher values at more relevant locations

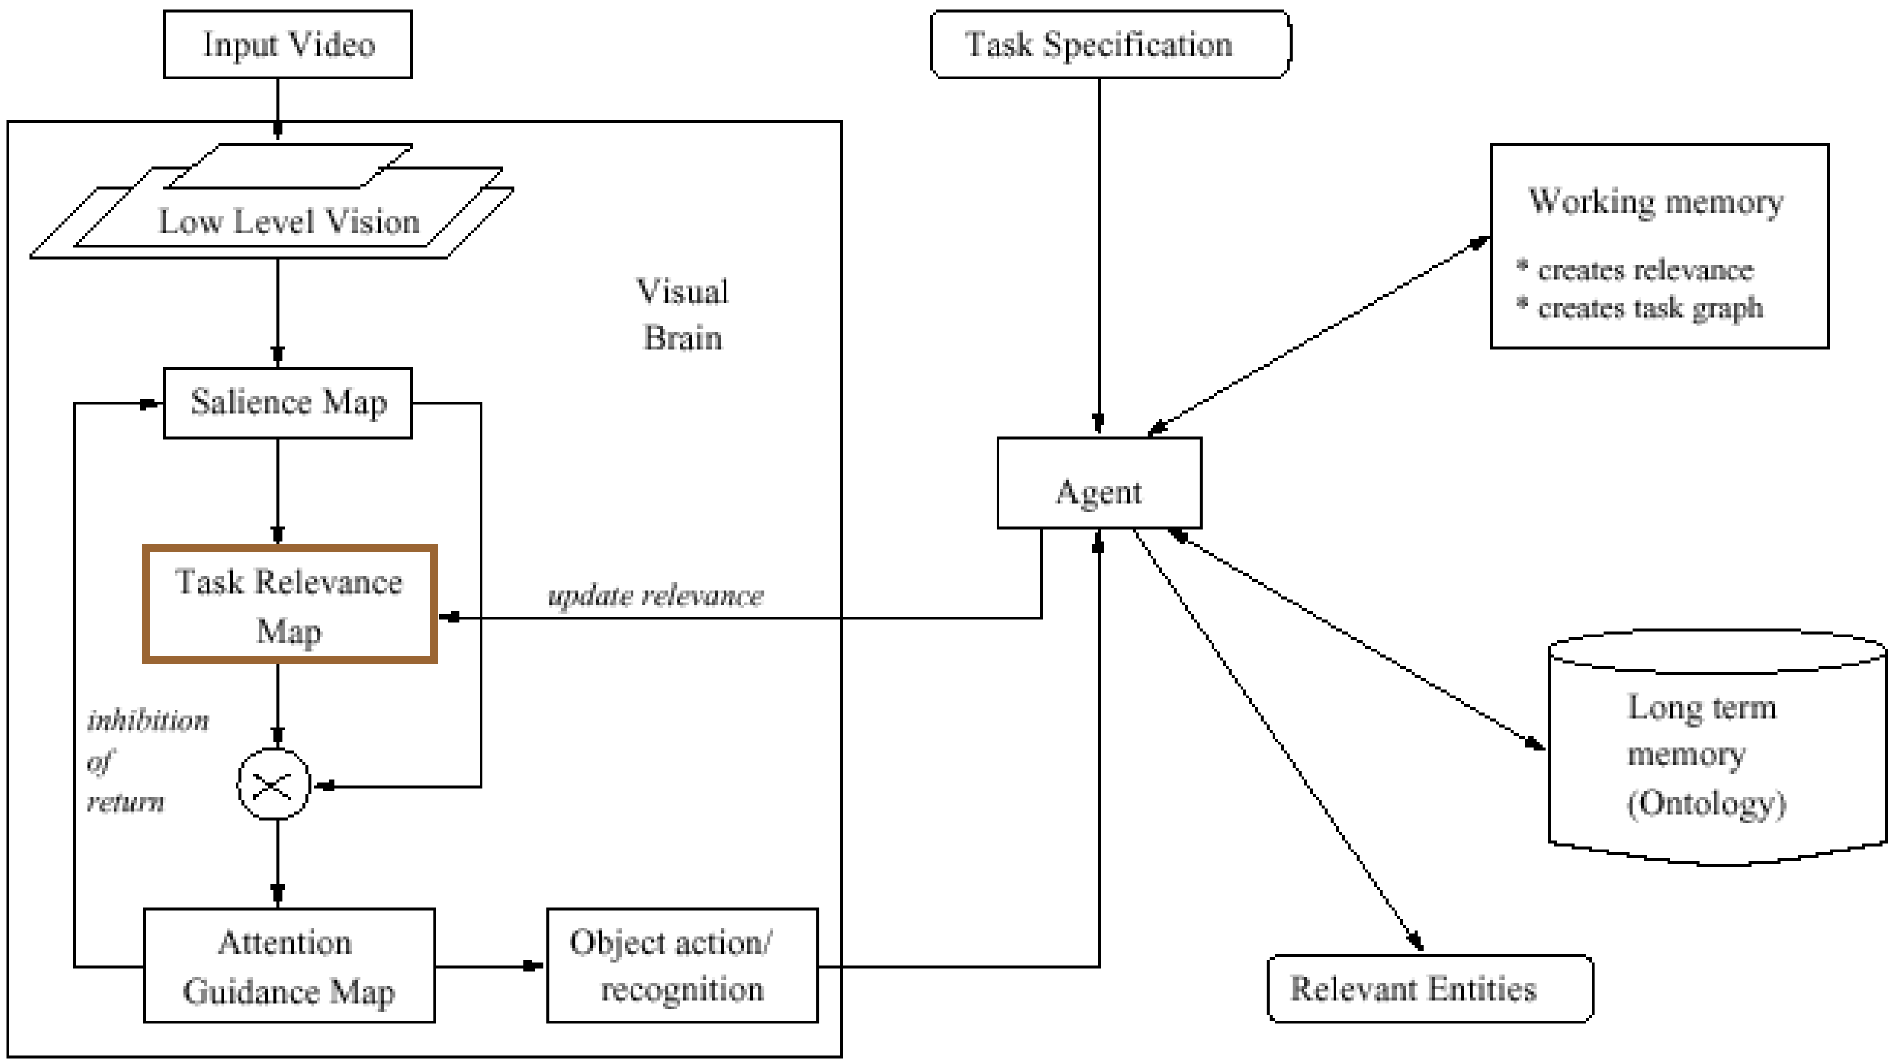

More formally: how do we do it?

- Use ontology to describe categories, objects and relationships:

  Either with unary predicates, e.g., Human(John),

  Or with reified categories, e.g., John $\in$ Humans,

  And with rules that express relationships or properties,

  e.g., $\forall x$ Human(x) $\Rightarrow$ SinglePiece(x) $\wedge$ Mobile(x) $\wedge$ Deformable(x)

- Use ontology to expand concepts to related concepts:

  E.g., parsing question yields "LookFor(catching)"

  Assume a category HandActions and a taxonomy defined by

  catching $\in$ HandActions, grasping $\in$ HandActions, etc.

  We can expand "LookFor(catching)" to looking for other actions in the category where catching belongs through a simple expansion rule:

  $\forall a,b,c \quad a \in c \wedge b \in c \wedge$ LookFor(a) $\Rightarrow$ LookFor(b)

Outlook

- AI is a very exciting area right now.
- This course will teach you the foundations.
- In addition, we will use the Beobot example to reflect on how this foundation could be put to work in a large-scale, real system.