Artificial Intelligence

Berlin Chen 2005

Course Contents

- The theoretical and practical issues for all disciplines of Artificial Intelligence (AI) will be considered
  - AI is interdisciplinary !

- Foundational Topics to Covered
  - Intelligent Agents
  - Search, Advanced Search, Adversarial Search (Game Playing), Constraint Satisfaction Problems (CSP)
  - Propositional and Predicate Logic, Inference and Resolution
  - Rules and Expert Systems
  - Probabilistic Reasoning and Bayesian Belief Networks
  - Others (Hidden Markov Models, Graphical Models, Neural Networks, Genetic Algorithms, etc.)

Textbook and References

- Textbook:
  - S Russell and P. Norvig. ***Artificial Intelligence: A Modern Approach***. Prentice Hall, 2003

    http://aima.cs.berkeley.edu/

- References:
  - M. Negenevitsky. ***Artificial Intelligence: A Guide to Intelligence Systems***. Addison-Wesley, 2005
  - Nils J. Nilsson. ***Artificial Intelligence: A New Synthesis***. Morgan Kaufmann, 1998
  - B. Coppin. ***Artificial Intelligence Illuminated***. Jones and Bartlett, 2004
  - E. Alpaydin, ***Introduction to Machine Learning***, MIT Press, 2004
  - T.M. Mitchell. ***Machine Learning***. McGraw-Hill, 1997

Grading

- Midterm or Final: 30%
- Homework: 25%
- Project/Presentation: 30%
- Attendance/Other: 15%

Introduction

Berlin Chen 2005

Reference:

1. S. Russell and P Norvig. ***Artificial Intelligence: A Modern Approach***. Chapter 1

What is AI ?

- "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning…" (Bellman, 1978)
- "The exciting new effort to make computer think … machines with mind, in the full and literal sense." (Haugeland, 1985)
- "The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985)
- "The study of how to make computers do things at which, at the moment, people do better." (Rich and Knight, 1991)

What is AI ?

- The study of the computations that it possible to perceive, reason, and act." (Winston, 1992)
- "AI…is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

AI systemizes and automates intellectual tasks

as well as any sphere of human intellectual activities.

- Duplicate human facilities like creativity, self-improvement, and language use
- Function autonomously in complex and changing environments

AI still has openings for several full-time Einsteins !

Categorization of AI

	fidelity	rationality
Thought/ reasoning	Systems that **think** like humans	Systems that **think** rationally
behavior	Systems that **act** like humans	Systems that **act** rationally

- Physical simulation of a person is unnecessary for intelligence ?
  - Humans are not necessarily “rational”

Acting Humanly: The Turing Test

- Turing test: proposed by Alan Turing, 1950

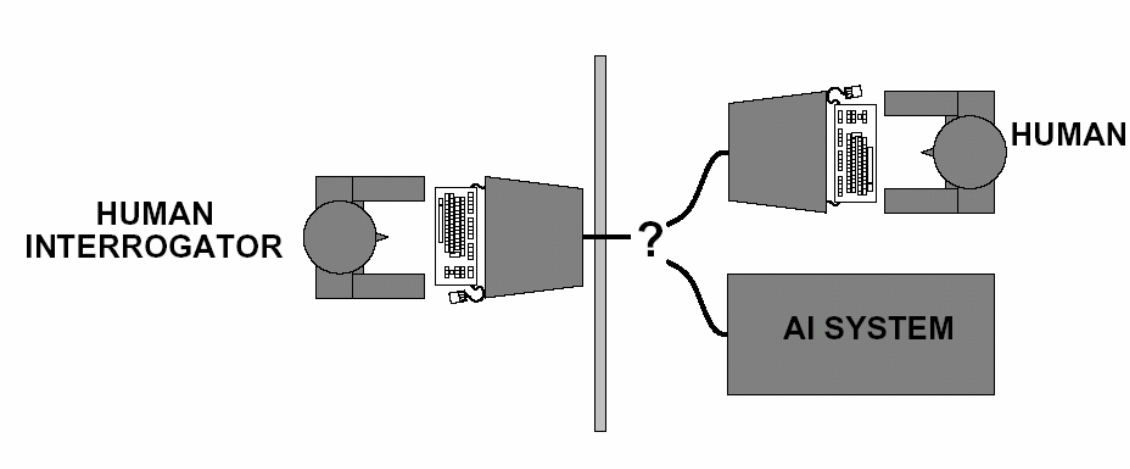

  - The test is for a program to have a conversation (via online typed messages) with an interrogator for 5 minutes
  - The interrogator has to guess if the conversation is with a machine or a person
  - The program passes the test if it fools the interrogator 30% of the time

Acting Humanly: The Turing Test

- Turing's Conjecture
  - At the end of 20 century a machine with 10 gigabytes of memory would have 30% chance of fooling a human interrogator after 5 minutes of questions
- Problems with Turing test
  - The interrogator may be incompetent
  - The interrogator is too lazy to ask the questions
  - The human at the other hand may try to trick the interrogator
  - The program doesn't have to think like a human
  - ….

Acting Humanly: The Turing Test

- The computer would possess the following capabilities to pass the Turing test
  - **Natural language** (Speech processing?)
  - **Knowledge representation**
  - **Automated reasoning**
  - **Machine learning/adaptation**
  - Computer vision
  - Robotics

physical simulation

Six disciplines compose most of AI

So-called "total Turing Test"

Imitate humans or learn something from humans ?

Acting Humanly: The Turing Test

- However, scientists devoted much effort to studying the underlying principles of intelligence than passing Turing test !

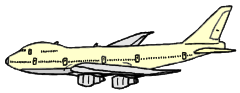

  - E.g. Aircrafts vs. birds
  - E.g. Boats/submarines vs. fishes/dolphins/whales
  - E.g. Perception in speech/vision

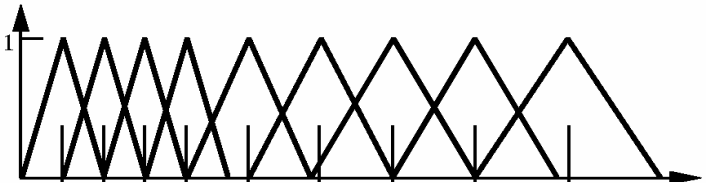

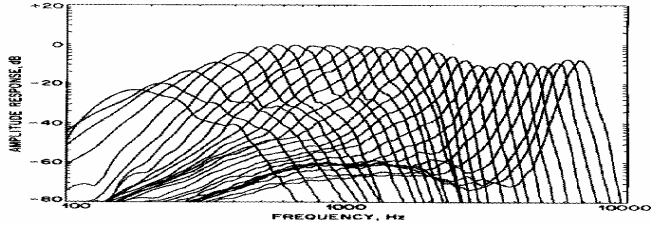

Thinking Humanly: Cognitive Modeling

- Get inside the actual workings of human minds through
  - Introspection
  - Psychological experiments

  } find the theory of the mind or trace the steps of humans' reasoning

- Once having a sufficiently precise theory of the mind, we can express the theory as a computer program !

- Cognitive science - interdisciplinary
  - Computer models from **AI**
  - Experimental techniques from **psychology**

An algorithm performs well $\overset{?}{\longleftrightarrow}$ A good model of human performance

Thinking Rationally: Laws of Thought

- Correct inference
  - "Socrates is a man; all men are mortal; therefore, Socrates is mortal"
  - – Correct premises yield correct conclusions

- Formal logic
  - – Define **a precise notion** for statements all things and the relations among them
    - Knowledge encoded in logic forms
  - – Main considerations
    - Not all things can be formally repressed in logic forms
    - Computation complexity is high

Acting Rationally: Rational Agents

- An agent is just something that perceives and acts
  - E.g., computer agents vs. computer programs
  - Autonomously, adaptively, goal-directly
- Acting rationally: doing the right thing
  - The right thing: that which is expected to maximize the goal achievement, given the available information
  - Don't necessarily involve thinking/inference
- Rationality $\longleftrightarrow$ Inference
- The study of AI as rational-agent design

Acting Rationally: Rational Agents

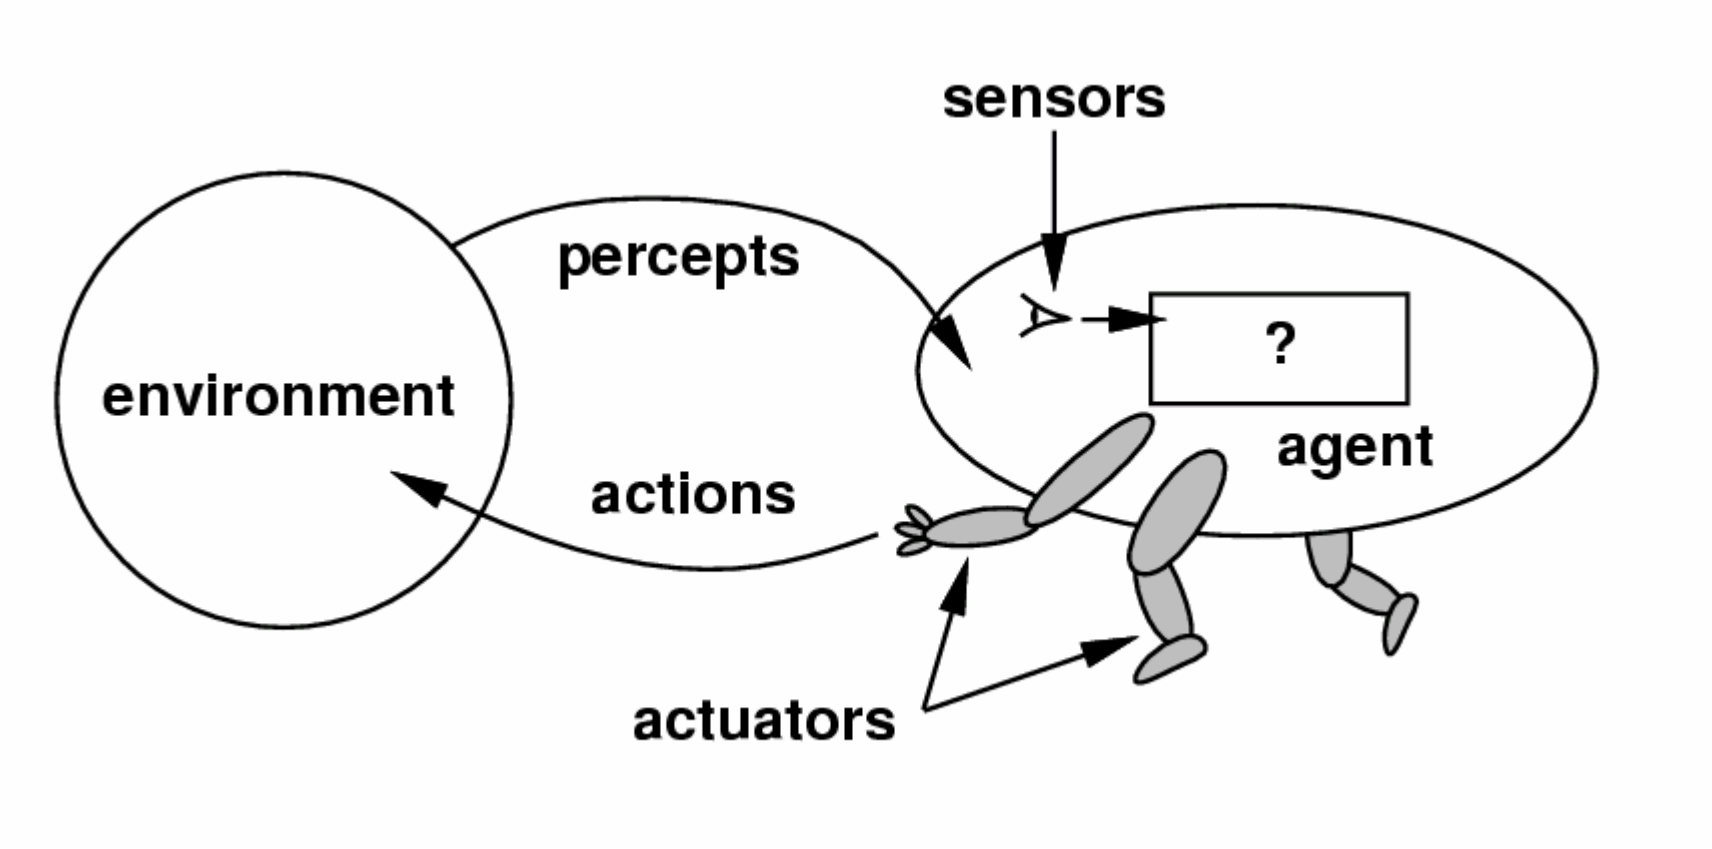

Foundations of AI

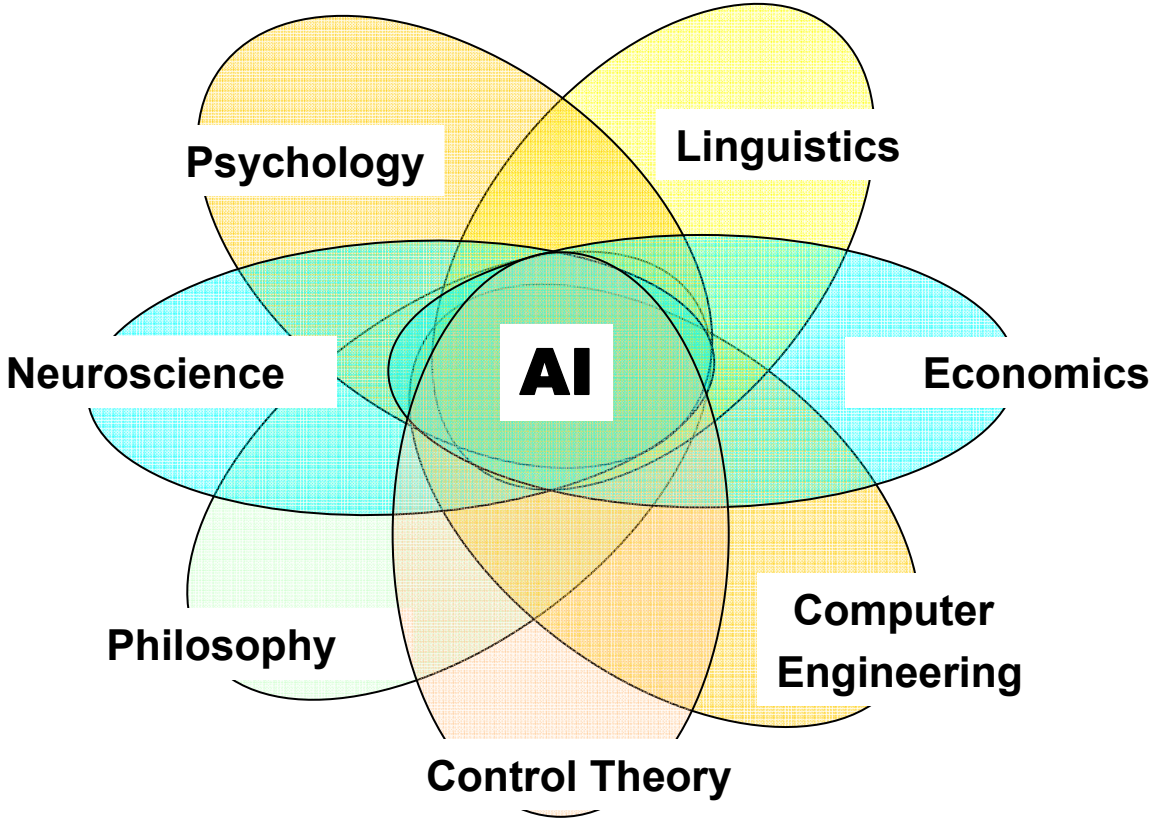

Foundations of AI

- **Philosophy** : (428 B.C. - present)

  Logic, methods of reasoning

  – A set of rules that can describe the formal/rational parts of mind
  – Mind as a physical system / computation process
  – Knowledge acquired from experiences and encoded in mind, and used to choose right actions
  – Learning, language, rationality

Foundations of AI

- **Mathematics** (C. 800 - present)

  Formal representation and proof

  – Tools to manipulate logical/probabilistic statements
  – Groundwork for computation and algorithms

Three main contributions:

- (decidability of) logic, (tractability of) computation, and probability (for uncertain reasoning)

Foundations of AI

- **Economics** (1776 - present)

  Formal theory for the problem of making decisions

  – Utility: the preferred outcomes
  – Decision theory
  – Game theory（賽局）
  – Operations research
    - Payoffs from actions may be far in the future

Maximize the utility
Right actions under competition

Foundations of AI

- **Neuroscience** (1861- present)

  Brains cause minds

  – The mapping between areas of the brain and the parts of body they control or from which they receive sensory input

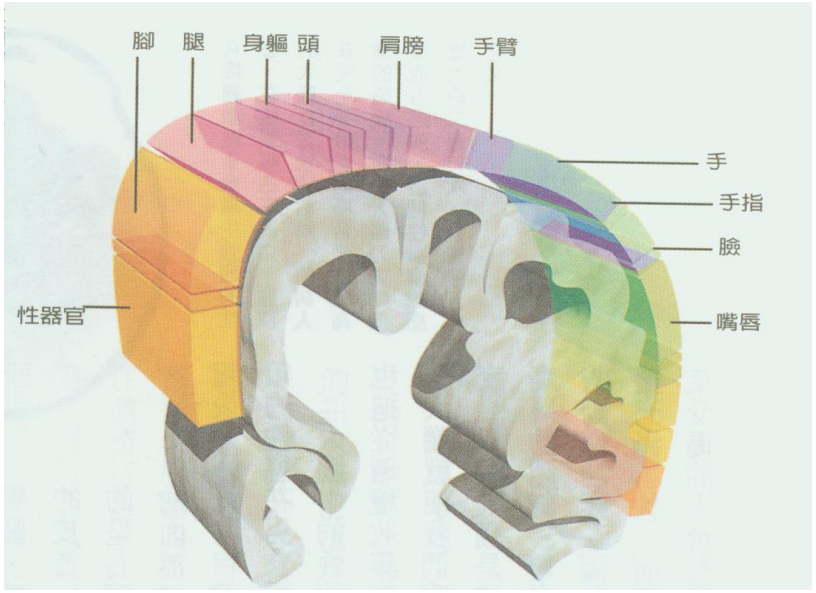

Ramón y Cajál (拉蒙卡哈)

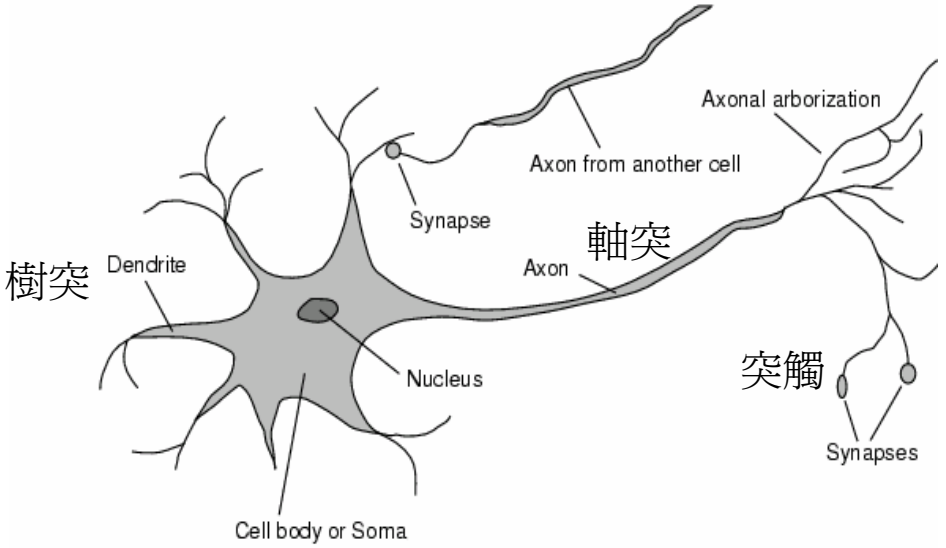

Foundations of AI

- **Psychology** (1879- present)

  Brains as information-processing devices

  – Knowledge-based agent
    - Stimulus translated into an internal representation
    - Cognitive process derive new international representations from it
    - These representations are in turn retranslated back into action
- **Computer engineer** (1940- present)

  Artifacts for implementing AI ideas/computation
  - (Software) programming languages
  - The increase in speed and memory

Foundations of AI

- **Control theory** (1948- present)

  Maximizing an objective function over time

  – Minimize the different between current and goal states

- **Linguistics** (1957- present)

  How does language relate to thought?

  – Languages fit information processing model

  – Understanding languages requires an understanding of subject matter and context

History of AI

- 1943-55 Gestation of Artificial Intelligence
  McCulloch & Pitt: Boolean circuit model of neurons
  Turing's "Computing Machinery and Intelligence"
- 1956 The birth of Artificial Intelligence
  Dartmouth meeting: "Artificial Intelligence" adopted
  (McCarthy, Minsky, Shannon, …)
- 1966-85 Neural network research almost disappears
  No efficient Training Algorithms for Layered networks
- 1969-79 Knowledge-based systems
- 1980-88 Expert system industry booms
- A million dollars to billions of dollars
- 1986- Neural networks return to popularity
- 1988-93 Expert system industry busts: "AI winter"
- 1995- Agents everywhere …

Advances in AI

- **Search Algorithms** like genetic algorithms

- Machine Learning
  - **Hidden Markov Models**
  - **Neural Networks**
  - **Bayesian Networks**
  - **Support Vector Machines**

- Data Mining

Research Prototypes

Marvin Minsky

- MIT Microworlds: “analogy problems” (1968)

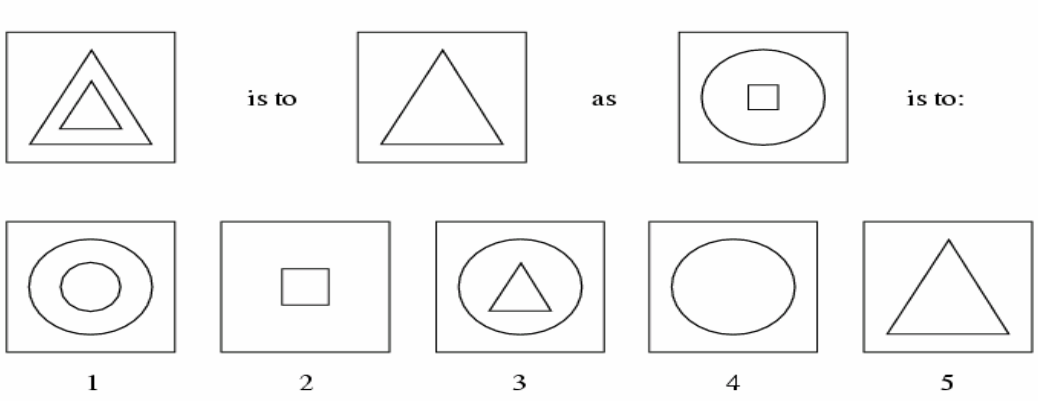

- MIT Microworlds: “blocks world” (1972)

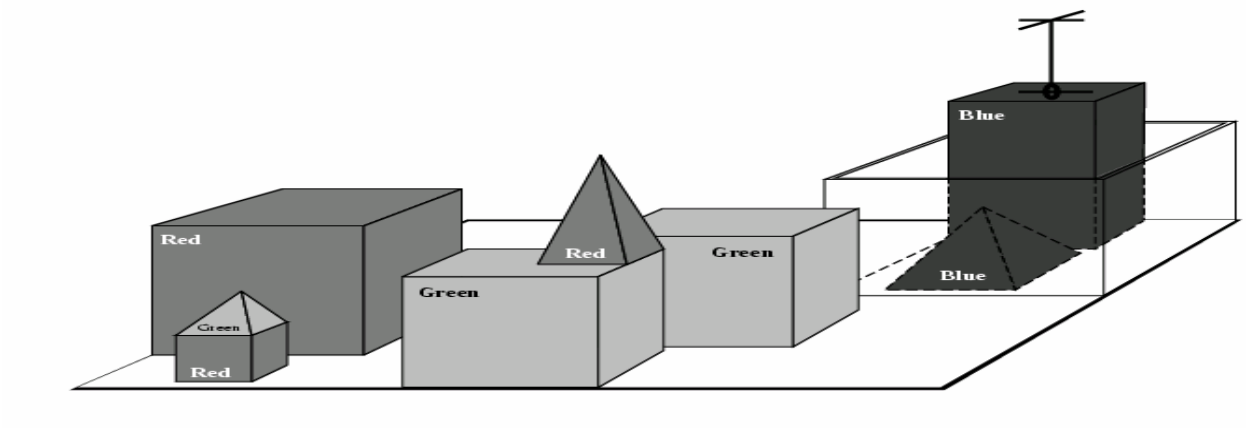

Research Prototypes

- ELIZA: A simple natural-language understanding program (1966)

User1:	Men are all alike.
ELIZA1:	IN WHAT WAY
User2:	They're always bugging us about something or other.
ELIZA2:	CAN YOU THINK OF A SPECIAL EXAMPLE
User3:	Well, my boyfriend made me come here.
ELIZA3:	YOUR BOYFRIEND MADE YOU COME HERE
User4:	He says I'm depressed much of the time.
ELIZA4:	I AM SORRY TO HERE YOU ARE DEPRESSED

s/.* all .*/IN WHAT WAY/

s/.* always .*/CAN YOU THINK OF A SPECIFIC EXAMPLE/

s/.* I'm (depressed|sad) .*/I AM SORRY TO HEAR YOU ARE \1/

Research Prototypes

- CMU ALVIN project, 1989 (Autonomous Land Vehicle In a Neural Network)
  - 1200 computer-generated images as training examples
    - Half-hour training
    - The salient features have been directly acquired by the network itself

An additional information from previous image indicating the darkness or lightness of the road

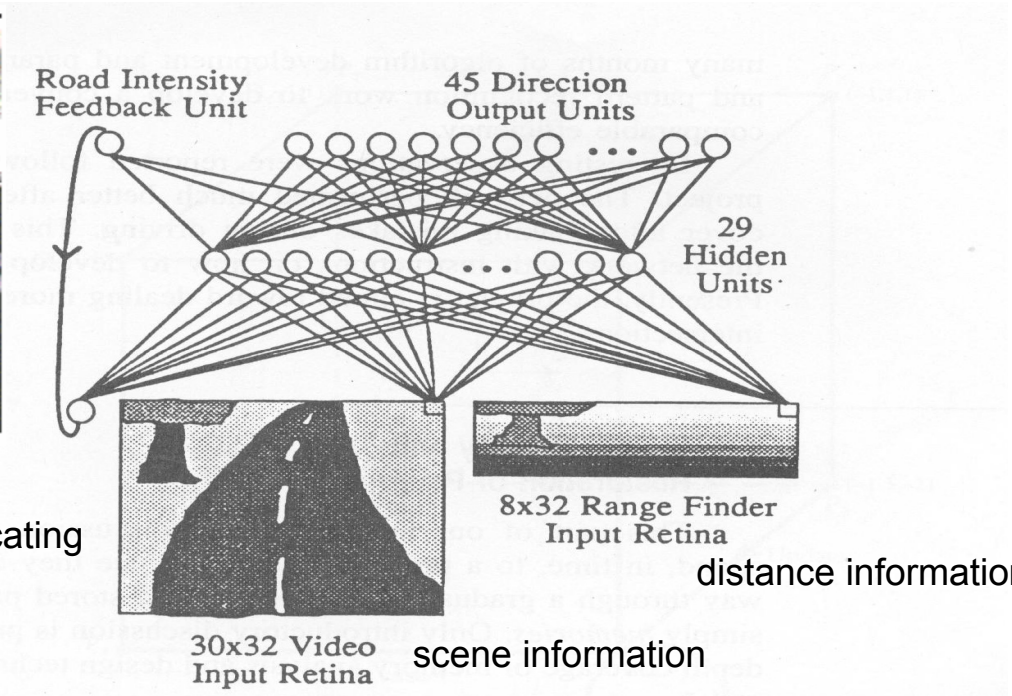

Research Prototypes

- IBM Deep Blue (1997)

  - Let IBM's stock increase by $18 billion at that year

Research Prototypes

Research Prototypes

- Sony AIBO robot
  - Available on June 1, 1999
  - Weight: 1.6 KG
  - Adaptive learning and growth capabilities
  - Simulate emotion such as happiness and anger

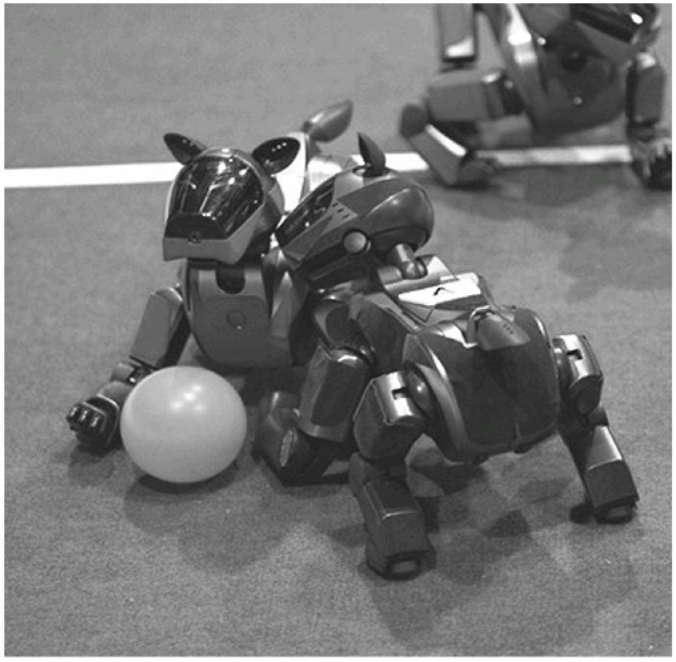

Research Prototypes

- Honda ASIMO (**A**dvanced **S**tep in **I**nnovate **MO**bility)
  - Born on 31 October, 2001
  - Height: 120 CM, Weight: 52 KG

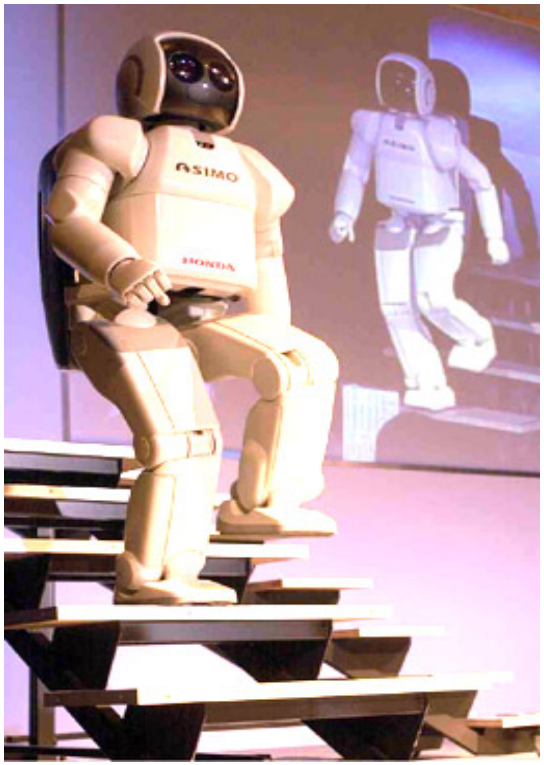

Toy examples ?
Real-world applications ?

Research Prototypes

- MIT CSAIL (電腦科學與人工智慧)

Research Prototypes

- MIT Oxygen Project: Spoken Interface (CMU, Delta)

  ubiquitous

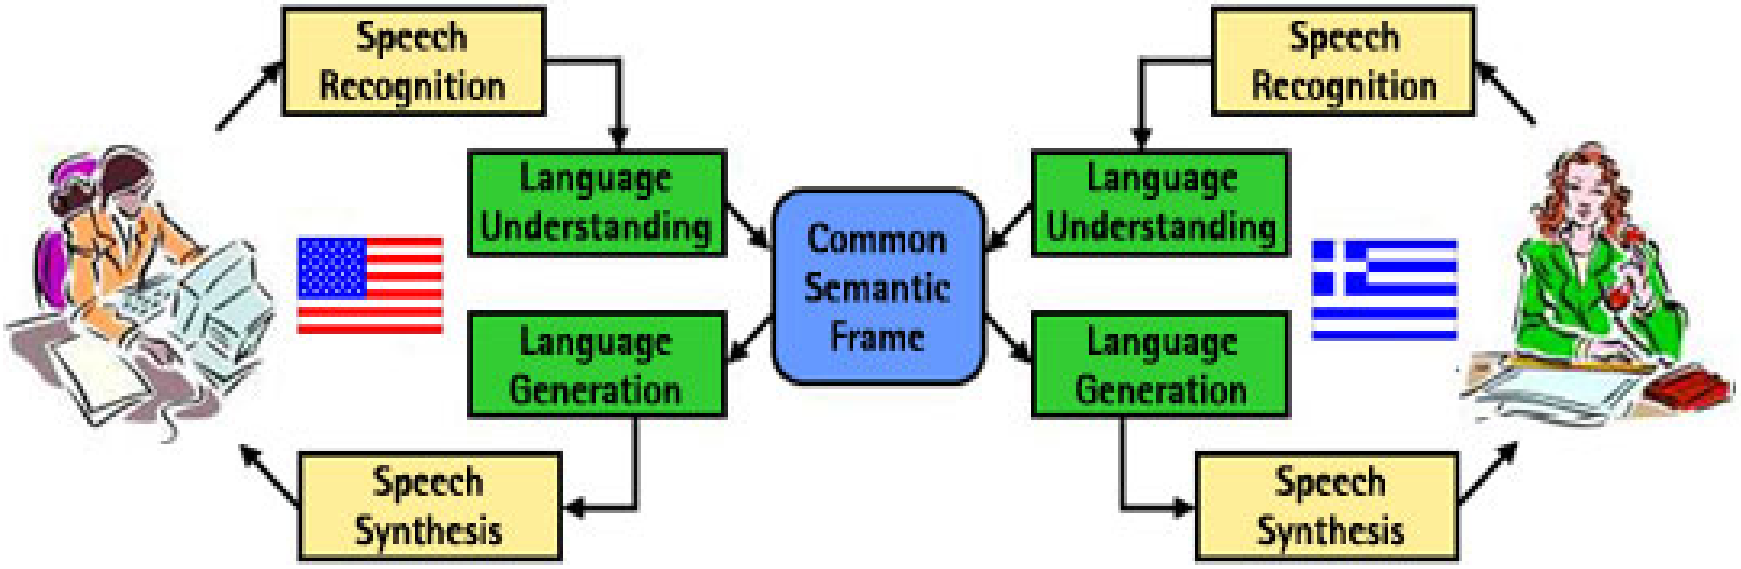

  - Speech recognition/synthesis
  - Natural language understanding/generation
  - Machine translation

SR & AI ?

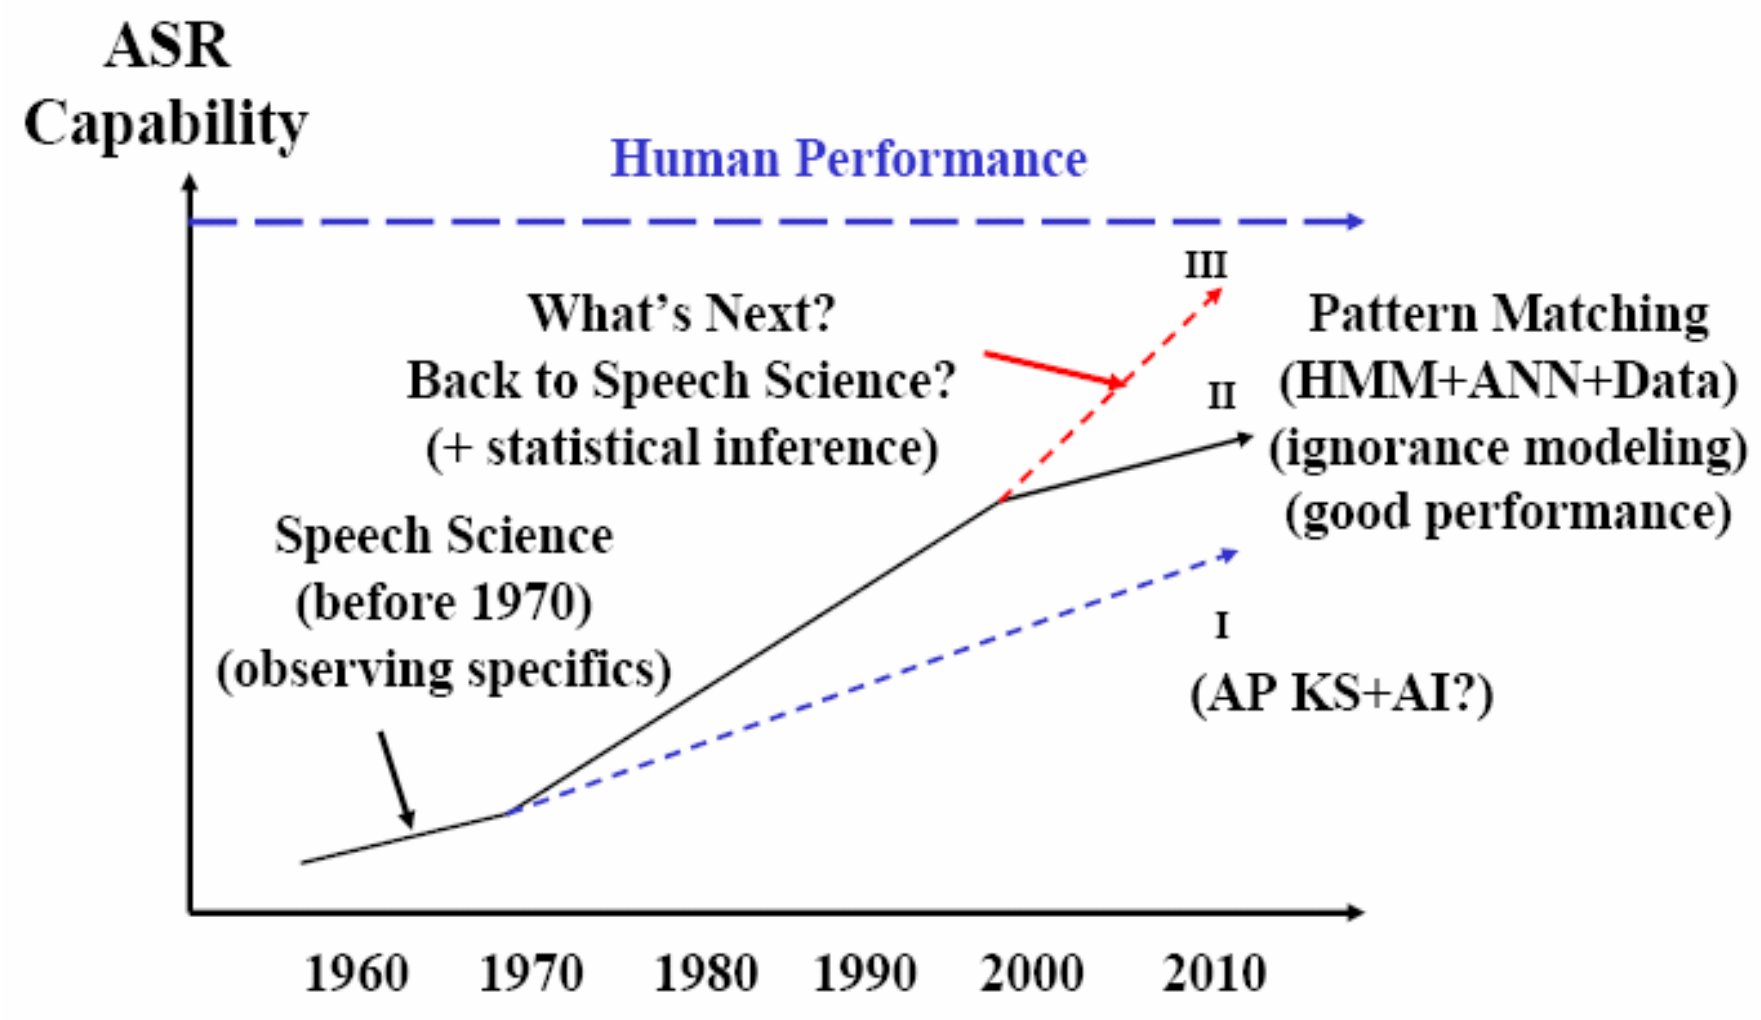

From Prof. Chin-Hui Lee

Tentative Schedule

9/22	Course Overview & Introduction
9/28	Agents
10/6	Searching: Uninformed Search: DFS, BFS, IDS, etc.
10/13	Searching: Informed Search: Greedy Best-First, A* Search, etc.
10/20	Searching: Informed Search: Local Search, Genetic algorithms, etc.
10/27	Searching: Constraint Satisfaction
11/3	Searching: Adversarial Search (Game Playing) (deterministic)
11/10	**Midterm**
11/17	Logical Agent & Propositional Logic
11/24	First-Order Logic
12/1	Inference and Planning
12/8	Rule-based/Fuzzy/Frame-based Expert Systems
12/15	Artificial Neural Networks
12/22	Paper Survey (1/2)
12/29	Paper Survey (2/2)
1/5	**Final**