More Intro to AI

Dr. Melanie Martin

CS 4480

August 25, 2010

What is AI?

- “The practice of designing systems that possess and acquire knowledge and reason with knowledge.” (Tanimoto 1987)
- “The design and study of computer programs that behave intelligently.” (Dean, Allen, Aloimonos 1995)
- “The branch of computer science concerned with making computers behave like humans.” (Webopedia)

What is AI?

- But then, what is intelligence???
  - "the capacity for learning, reasoning, understanding, and similar forms of mental activity; aptitude in grasping truths, relationships, facts, meanings, etc."
    (Webster's Encyclopedic Unabridged Dictionary of the English Language 1996)

What is AI?

Categories under AI on Cora (~1999-2001)

Domain Specific Search Engine for CS papers

- Agents
- Data Mining
- Expert Systems
- Games and Search
- Knowledge Representation
- Machine Learning
- Theory, Case-Based, Rule Learning, ...
- Natural Language Processing
- Planning
- Robotics
- Speech
- Theorem Proving
- Vision & Pattern Recognition

What is AI?

- Goals in AI
  - Engineering: Solve real-world problems. Build systems that exhibit intelligent behavior.
  - Scientific: Understand what kind of computational mechanisms and knowledge are needed for modeling intelligent behavior.

What is AI?

- Do we really want to model humans?
  - Seem like our best example, but….
  - Should we build airplanes with wings that flap like birds?
- How do we know we did it?
  - Turing test?
    - Focus on behavior instead of internal algorithm
    - Defines success in terms of human intelligence

The Turing Test

- Test proposed by Alan Turing in 1950
- The computer is asked questions by a human interrogator. It passes the test if the interrogator cannot tell whether the responses come from a person
- Required capabilities: natural language processing, knowledge representation, automated reasoning, learning,...
- No physical interaction

But, does this show intelligence???

The Chinese Room

- Searle (1980) p. 958

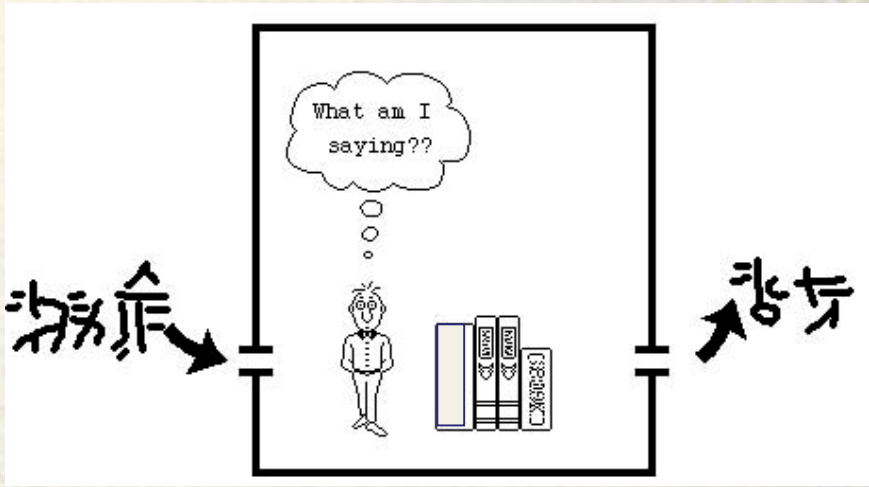

- Human: CPU
- Rule Book: Program
- Paper: Memory

- Human understands only English
- Input symbols, output symbols based on rules
- Appears to have conversation in Chinese

How strong do you like your AI?

- Weak AI
  - Machines could act as if they were intelligent
- Strong AI
  - Machines that act intelligent are actually thinking

What is AI?

- A few recurring issues:
  - How important is cognitive modeling in our systems?
  - How do we balance scientific and engineering goals?
  - How do we evaluate our system?

Acting rationally: rational agent

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: \mathcal{P}^{\star} \rightarrow \mathcal{A}]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable

  → design best program for given machine resources

Coming Next

- By Friday
  - Read Chapter 1 and 2 in R&N
- We'll cover Intelligent Agents briefly (Ch. 2)
- Then to Problem Solving (search)
- Some Lisp programming in the next few weeks
- Now on to Chapter 2