Introduction

CS 486/686: Introduction to Artificial Intelligence
Fall 2013

Outline

- Course administration

- What is AI? (Chapter 1)

  - Definitions

  - History

  - What we will cover

- Rational Agents (Chapter 2)

Course Administration

- CS 486/686: Introduction to Artificial Intelligence
  - Section 1: MC 4045 Tues/Thurs 4:0-5:20

- Course Personnel:
  - Kate Larson (klarson@uwaterloo.ca)
    - Office: DC 2518
    - Office Hours: Mondays 3:00-4:00 in DC 2518
  - TAs: Arthur Carvalho, Hadi Hosseini, Milad Khaki

Course Administration

- Website:

  - http://www.cs.uwaterloo.ca/~klarson/teaching/F13-486

- Newsgroup:

  - We will be using Piazza for the newsgroup. Details on how to sign up are in the syllabus.

- Texts:

  - Artificial Intelligence: A Modern Approach by S. Russell and P. Norvig (3rd Edition)

  - Artificial Intelligence: Foundations of Computational Agents, D. Poole and A. Mackworth (available online)

Evaluation

CS 486

- 4 Assignments: 40%

- Midterm (in class on October 10): 20%

- Final: 40%

- Project (Optional): up to 5 bonus marks

CS 686

- 4 Assignments: 28%

- Midterm (in class on October 10): 12%

- Final: 35%

- Project: 25%

Assignment Late Policy

- Assignments are due as announced

  - For each assignment, you can pass it in up to 48 hours late

    - No doctor's note required, etc

- **<u>BUT</u>**

  - No assignment will be accepted after the 48 hour grace period

  - No questions about the assignment will be answered during the 48 hour period

Outline

- Course administration

- What is AI? (Chapter 1)

  - Definitions

  - History

  - What we will cover

- Rational Agents (Chapter 2)

What is AI?

- According to media/
  popular perception

  - What socially-inept
    hackers do

  - Hal, Data, Sonny

  - "When Robots Attack"

  - ...

HOW TO SURVIVE A ROBOT UPRISING

TIPS ON DEFENDING YOURSELF
AGAINST THE COMING REBELLION

Daniel H. Wilson

What is AI?

- Definition of AI differ along two dimensions

  - Reasoning vs behaviour

  - Fidelity to human behaviour vs rationality

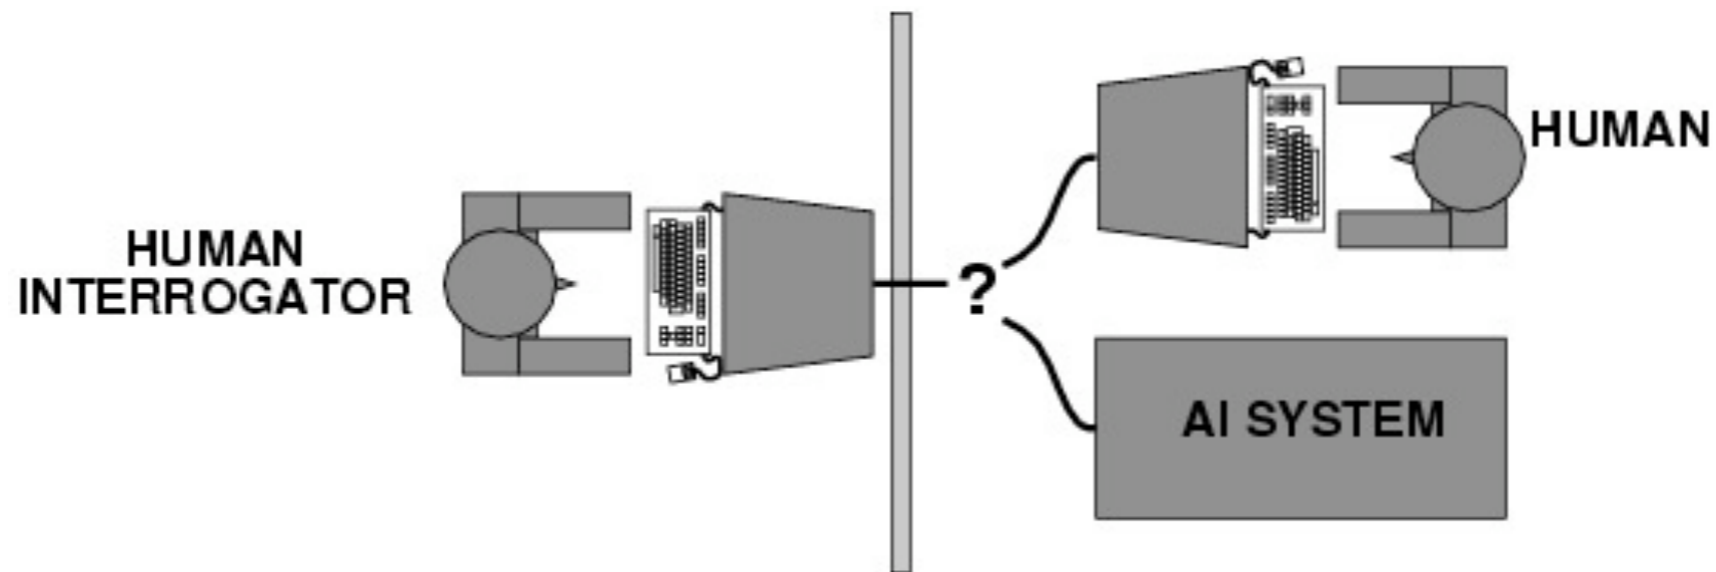| Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

What are the Goals of AI?

- AI is about duplicating what the (human) brain **DOES**

  - Turing Test

- AI is about duplicating what the human brain **SHOULD DO**

  - Rationality

Behaving Like a Human

- Alan Turing (1950) "Computing machinery and intelligence"

Computing Machinery and Intelligence

- Predicted that by 2000 a computer would have a 30% chance of fooling a lay person for 5 minutes

- Anticipated all major arguments against AI

- Suggested major components of AI:

  - Knowledge, Reasoning, Language Understanding, Learning

The Turing Test

- The test is still relevant today

    - The Loebner Prize

- However, AI researchers spend limited effort on it

    - It is not reproducible or amenable to mathematical analysis

    - More important to understand underlying principles of intelligence that copy them?

What are the Goals of AI?

- AI is about duplicating what the (human) brain **DOES**

  - Turing Test

- AI is about duplicating what the human brain **SHOULD DO**

  - **Rationality**

Rational Behaviour

- Rational behaviour

  Doing the Right Thing

- Doing what is expected to maximize goal achievement, given available information

  - Does not necessarily require thinking

    - But often thinking serves rational behaviour

Abridged History of AI

- 1943: McCulloch & Pitts: Boolean circuit of the brain

- 1950: Turing's "Computing machinery and intelligence"

- 1950s: Early AI Programs including Samuel's checkers, Newell and Simon's Logic Theorist

- **1956: Dartmouth meeting: "Artificial Intelligence"**

- 1966-1973: Problems with scaleability, Perceptron paper

- 1970s: Knowledge-based systems

- 1980's: Expert-systems industry

- 1988-now: Probabilistic and decision theoretic methods

"The AI Winter"

Classical AI

- Reasoning was seen as THE AI problem

  – Chess was considered pivotal to understanding intelligence

- **Goal**: General Problem Solver

Recent AI

- Focus on solving specific problems

- Heavy use of probability theory, decision theory, statistics,...

- Collection of subfields

  - Perception (including vision) is usually separate

  - Robotics is mostly separate

  - Deliberative reasoning is "AI"

    - But lots of different approaches

Course Contents

- Search

- Knowledge Representation and Reasoning

- Planning

- Reasoning Under Uncertainty

- Learning

Rational Agents

- An entity that perceives and acts

  - Function from percept to actions $f:P \rightarrow A$

- Performance measures

  - Goal achievement, resource consumption,...

- **Caveat**: Computational limitations and environmental constraints mean we do not have perfect rationality
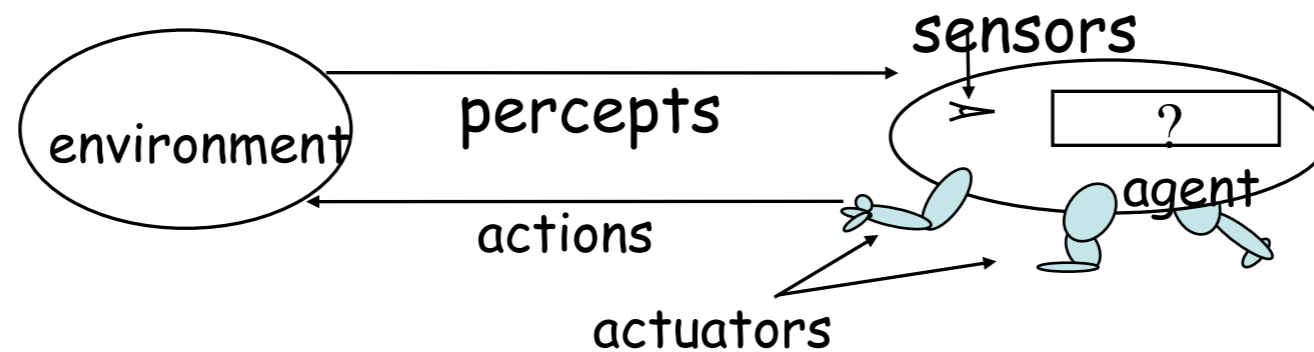
Task Environment

- To design a rational agent, the **task environment** must be specified

  - Performance measure

  - Environment

  - Actuators

  - Sensors

Performance Measures
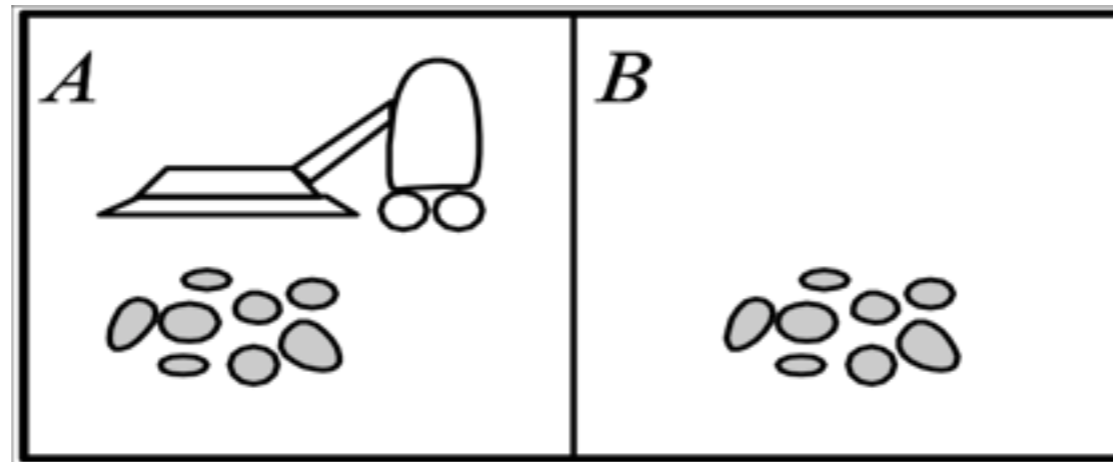
- **Percepts:** [Location, Dirty or Clean]

- **Actions:** Right, Left, Vacuum, NoOp, Dump

- **Function**: ([A,Clean],Right), ([A, Dirty], Vacuum), ([B, Dirty], Vacuum), ((([B, Clean],[B, Clean]), Left)...

Properties of Task Environment

- Fully Observable vs Partially Observable

- Deterministic vs Stochastic

- Episodic vs Dynamic

- Discrete vs Continuous

- Single agent vs Multi agent

State of the Art

- Chess was **THE** AI challenge for decades

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

-Gary Kasparov

MONTY NEWBORN

KASPAROV VERSUS DEEP BLUE

COMPUTER CHESS COMES OF AGE
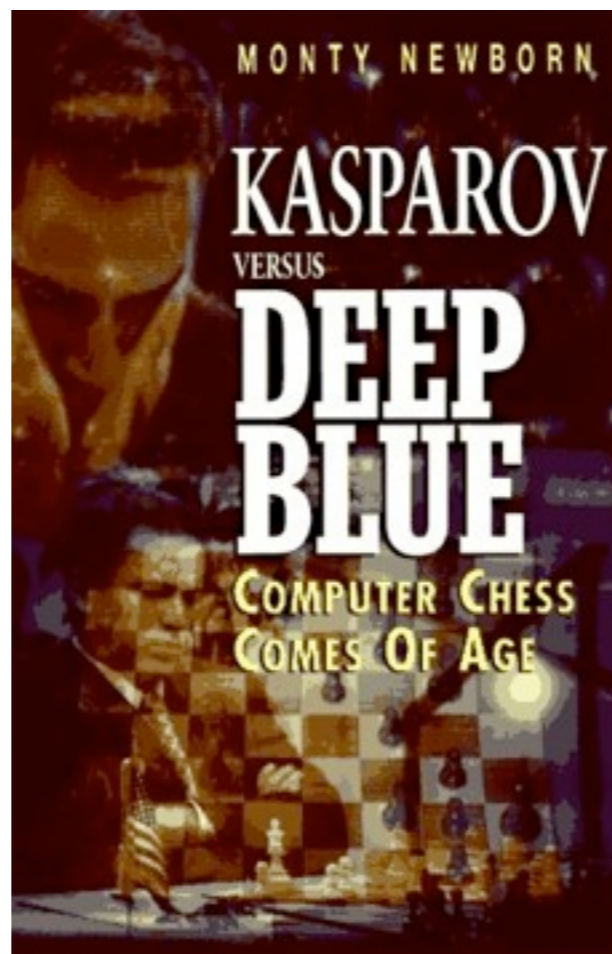
"Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings."

– Drew McDermott

State of the Art

- Checkers?

- Poker?

- Jeopardy?

State of the Art

- Airport Security?

State of the Art

- Autonomous Vehicles?

State of the Art

- Real Time Translation?

State of the Art

- Journalism?

Questions?

- Next lecture: Problem Solving Agents (Chapter 3)