Welcome to

CSCI 108

Artificial Intelligence:
Image and Reality

Prof. Andrea Danyluk

Why take this course?

- Practical (myopic?) view: to satisfy a Div III req'mt or a (Q) req'mt
- Technological view: to learn about AI and, in particular, robotics
  - To understand what's **really behind** the behavior of AI systems
  - To **implement** AI programs, with a focus on vehicular robots
- Philosophical view:
  - To consider the potential **impact** of AI on **society**
  - To consider the fundamental question of **intelligence** and its relationship to us as human beings
- General skills/knowledge view:
  - To develop analytical skills
  - To become a more informed consumer of scientific/technological information

Sample topics

- History of AI

Sample topics

- History of AI
- Robotics
  - Hardware: sensors and effectors
  - Control: planned vs reflexive

Sample topics

- History of AI
- Robotics
  - Hardware: sensors and effectors
  - Control: planned vs reflexive
- Vision
- Natural language

Sample topics

- History of AI
- Robotics
  - Hardware: sensors and effectors
  - Control: planned vs reflexive
- Vision
- Natural language
- Problem solving and reasoning (including games)
- Learning

Sample topics

- History of AI
- Robotics
  - Hardware: sensors and effectors
  - Control: planned vs reflexive
- Vision
- Natural language
- Problem solving and reasoning (including games)
- Learning
- Intelligence -- what is it? How can it be evaluated?
- Ethics
- Creativity

Lab

Not this

Lab

Or this

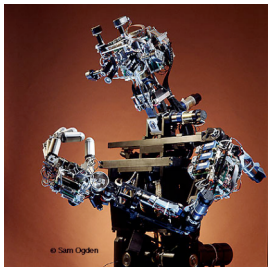

Lab

But definitely

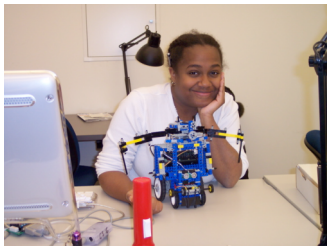

Syllabus

The syllabus and much more can be found at
http://www.cs.williams.edu/~andrea/cs108

- Contact information
  - "by appointment" means that and (almost) "any time you can find me"
- We have a TA
  - Bill Jannen '09

Syllabus

The syllabus and much more can be found at
http://www.cs.williams.edu/~andrea/cs108

- Text and Reading Packet
  - Get Reading Packet from Lorraine Robinson
  - Will order more; spare copy in CS Common Room for now
  - Everyone should have access to the readings

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Lectures and Discussions
  - Lectures should be interactive
  - Might consider other meeting places for discussions

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Labs
  - Attendance is mandatory
  - Default model: work in small groups
  - Won't always finish during lab time

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Assignments
  - A good thing, not a burden
  - Papers and problem sets won't overlap

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Schedule
  - Might see changes along the way
  - "When should I do the reading?"

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Honor Code
  - Honor Code for Courses in Computer Science
  - User Policy and Account Agreement
  - Sign in lab
  - *Please* follow it

Syllabus

The syllabus and much more can be found at http://www.cs.williams.edu/~andrea/cs108

- Grading, Responsibilities
  - Come to class
  - A great course depends on the prof and the students

Today and Monday

What is Artificial Intelligence (AI) anyway?

(Brief) history

To gain perspective on why the field is what it is; why we focus on certain topics and ignore others.

AI: the tumultuous history of the search for artificial intelligence, Daniel Crevier, Basic Books, 1993.

Goals of AI

- **Engineering goal:** To solve real-world problems. Build systems that exhibit intelligent behavior.
- **Scientific goal:** To discover and understand the computational mechanisms needed for modeling intelligent behavior.
- Interdisciplinary roots:
  - Computer Science and Engineering
  - Philosophy, Psychology, Cognitive Science
  - Mathematics, Physics, Economics, Statistics

Different Approaches

- Cognitive approach
  Building models of human cognition
- Logical agent approach
  Emphasis is on "correct" inference
- Rational agent approach
  Emphasis on developing methods to match or exceed human performance, possibly by very different means

What do you know about AI?

For many, understanding of AI comes from literature and film.

From 2001: A Space Odyssey

Probably no one would ever know this; it did not matter. In the 1980s, Minsky and Good had shown how neural networks could be generated automatically -- self-replicated -- in accordance with any arbitrary learning program. Artificial brains cold be grown…
Whatever way it worked, the final result was a machine intelligence that could reproduce -- some philosophers still preferred to use the word "mimic" -- most of the activities of the human brain, and with far greater speed and reliability.

"I am a HAL Nine Thousand computer Production Number 3. I became operational at the Hal Plant in Urbana, Illinois, on January 12, 1997."

Long ago… before AI was AI

- But not too long ago
  - 1943: Warren McCulloch and Walter Pitts propose a model of artificial neurons; show that any function can be computed by some network of connected neurons.
  - 1949: Donald Hebb demonstrates a simple updating process that allows neural networks to learn.
  - 1951: Marvin Minsky and Dean Edmonds build first neural network computer, the SNARC.

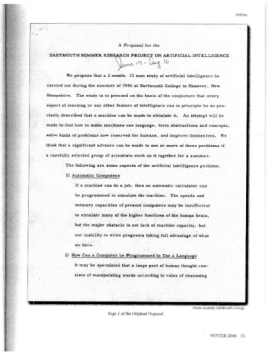

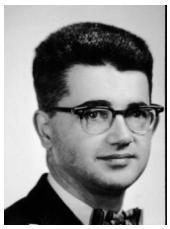

The Dartmouth Conference (1956)

- Introduced the major figures to each other:
  - Marvin Minsky, Nathaniel Rochester, Arthur Samuel, Oliver Selfridge, Herb Simon, Claude Shannon, Trenchard More, Ray Solomonoff, Allen Newell
- Most lasting contribution of the conference: agreement to adopt McCarthy's new name for the field, *Artificial Intelligence*

AI at 50: Dartmouth 2006

Trenchard More, John McCarthy, Marvin Minsky, Oliver Selfridge, Ray Solomonoff

Predictions of the founders

- "It is not my aim to surprise or shock you - but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until - in a visible future - the range of problems they can handle will be coextensive with the range to which human mind has been applied." (Herb Simon, 1957)

Predictions of the founders

- In 1958, Herb Simon predicted that within 10 years:
  - A computer would be chess champion
  - An important new mathematical theorem would be proved by machine
  - Most theories in psychology would take the form of computer programs, or of qualitative statements about the characteristics of computer programs

What became of the predictions?

- In 1996, Kasparov was still beating Deep Blue, but that changed in 1997
- Robbins' problem in finite algebra, a 60-year open problem: first "creative" proof by a computer (November 1996)
- Psychology has not changed in the ways Simon predicted (though cognitive psychologists do make use of computational models of intelligence and do so extensively)

Historical snapshots 1952-1969 Enthusiasm

- Newell and Simon's General Problem Solver (GPS)
- Arthur Samuel's checkers player
- John McCarthy moves to MIT - invents Lisp, time sharing, and moves to a logic approach, which he takes to Stanford
- Minsky moves to MIT and adopts an anti-logical outlook
- Shakey the robot - project starts at the new SRI
- Early work on neural nets is flourishing

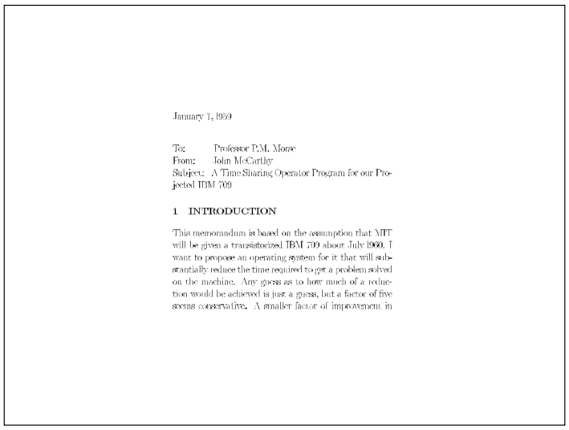

January 1, 1959

To: Professor P.M. Morse
From: John McCarthy
Subject: A Time Sharing Operator Program for our Projected IBM 709

1 INTRODUCTION

This memorandum is based on the assumption that MIT will be given a transistorized IBM 709 about July 1960. I want to propose an operating system for it that will substantially reduce the time required to get a problem solved on the machine. Any guess as to how much of a reduction would be achieved is just a guess, but a factor of five seems conservative. A smaller factor of improvement in

1966-1974 Reality

- Difficulties due to the intractability of AI problems
  - Theoretical
  - Practical (time and memory requirements of AI algorithms; imprecision of devices used in robotics)
- Microworlds don't scale up
- Lighthill report kills AI research in British universities
- Neural net research dies (Minsky and Papert's *Perceptrons*)

1969-1979 Paradigm Shift

- 1969-1979 Don't give up: use *knowledge and apply expert reasoning*
  - Birth of expert systems
  - Importance of domain knowledge and abstraction in all fields
- 1980-1988 The AI Industry
  - Companies emerge that offer AI tools
  - Proliferation of expert systems

Late 80s

- Companies realize that knowledge engineering is hard to do
  - The need for learning, probabilistic reasoning
- Re-invention of backpropagation learning algorithm
- Attempts to combine different techniques in principled ways*

1990s to the Present

- Increases in computational power and availability of data have played a large role
  - Probabilistic and data-intensive approaches viable
  - Real applications
- Renewed interest in knowledge-based approaches (or at least the integration of knowledge with other approaches)

Where we are today Example 1: Chess*

- February 1996: Kasparov vs Deep Blue
  - Kasparov victorious: 3 wins, 2 draws, 1 loss
- March 1997: Kasparov vs Deeper Blue
  - First match won against world champion
  - 512 processors; 200 million chess positions per second
  - "intelligent and creative play"

*Reasoning in the presence of an adversary.

Where we are today
Example 2: Other Games

- Backgammon: TD-Gammon (Tesauro 93, 95)
  - World-champion level
  - Learns by playing millions of games against itself
  - Has changed human play
- Checkers is solved!
  - (Jonathan Schaeffer et al, Science 07)

Where we are today
Example 3: Data Mining*

- Applied in many areas
  - Fraud detection
  - Analysis of scientific data (computational biology)
  - Analysis of demographic data
  - Homeland security

* Learning

Example 4: Robot Vehicles*

- No Hands Across America - July 1995
  - 2797/2849 miles (98.2%) Pittsburgh to San Diego
  - 1990 Pontiac TransSport
  - Portable computer, windshield mounted camera, GPS receiver
  - The RALPH computer program
  - Dean Pomerleau, Williams '87

*Vision, learning, reasoning, control

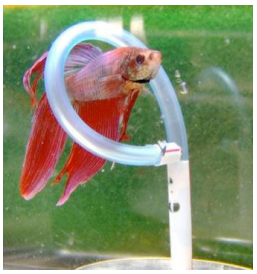

Example 4: Robot Vehicles

- DARPA Grand Challenge
  - 2005 winner: "Stanley" (Stanford)
  - Machine learning a critical component
  - 132 miles in less than 7 hours

Urban Challenge

- Finals: November 3, 2007
- Winner: Tartan Racing (CMU, GM, et al)

Example 5: Natural language and speech systems

- Speech replacing touch-tone interfaces
- "You talk, it types"
- Automatic translators

Example 6: Autonomous Bidding Agents*

- Complex market scenarios
- Empirical and game theoretic approaches
- Agents to help inform AT&T's bidding in FCC spectrum auction of Dec 2000, which brought in over $16 billion dollars (P.Stone)

*Reasoning

So, what is AI?

A moving target…

So, what is AI?

"the study of how to do things, which, at the moment, people do better" (Rich and Knight)

So, what is AI?

"the study of how to do things, which, at the moment, people do better" (Rich and Knight)

"the design and study of computer programs that behave intelligently" (Dean, Allen, and Aloimonos)

So, what is AI?

"the study of how to do things, which, at the moment, people do better" (Rich and Knight)

"the design and study of computer programs that behave intelligently" (Dean, Allen, and Aloimonos)

"the study of [rational] agents that exist in an environment and perceive and act" (Russell and Norvig)