CSE 473:

Introduction to Artificial Intelligence

Spring 2001

Dieter Fox

Organization

- Team:
  - Instructor: Dieter Fox
  - TA: Bill Pentney
  - TA: Chris Waterman
- Web page: `http://www.cs.washington.edu/473`
- Mailing list (see web page for further information)

Readings

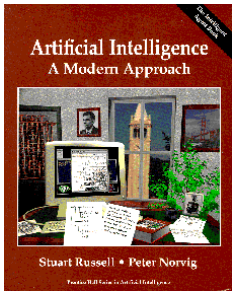

Artificial Intelligence – A Modern Approach
Stuart Russel - Peter Norvig

Papers will be posted on the web

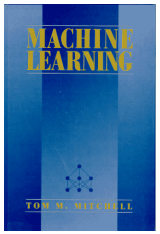

Machine Learning
Tom Mitchell

Assignments

- Four homeworks
- Two group projects
- Midterm
- Final

What is Artificial Intelligence?

- The attempt to make computers “more intelligent”.
- To understand the nature of human intelligence.

$\rightsquigarrow$ 4 approaches:

  – Is it about thinking . . .
  – . . . or acting?
  – Oriented at human model (including all its weaknesses) . . .
  – . . . or normative (how should a rational agent think/act)?

Some definitions . . .

"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

Systems that act like humans

Turing Test

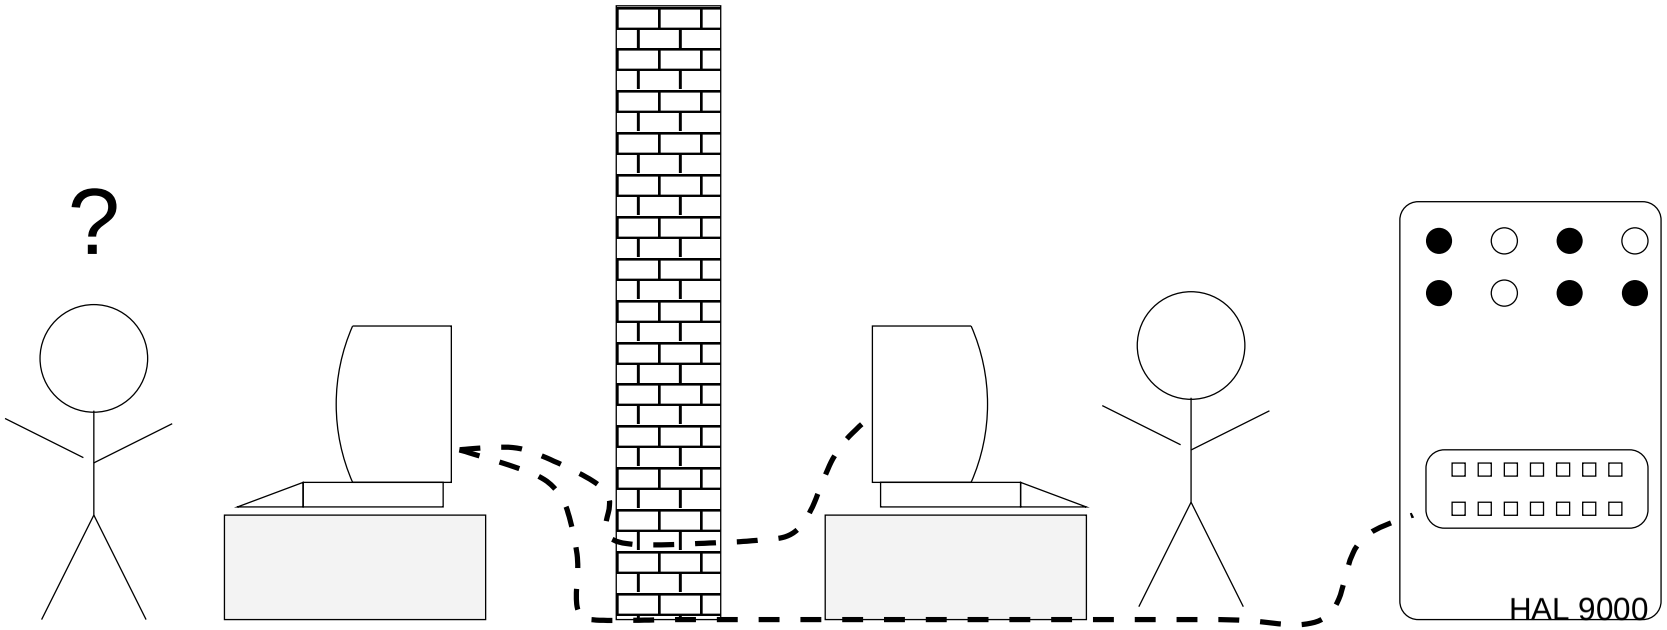

Topics

- Problem solving and search
- Logic and knowledge representation
- Planning
- Uncertainty
- Learning
- Applications in robotics

45 Years AI (1)

50's: Neurons, logic, games

1956: Dartmouth Workshop – McCarthy introduces the term *Artificial Intelligence* – and early enthusiasm:

> 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 the visible future*—the range of problems they can handle will be coextensive with the range to which human mind has been applied. [Simon, 1957]

60's: Intelligent problem solving in microworlds (e.g. blocks world), perceptron learning

45 Years AI (2)

70's: Problems:

- Systems didn't scale $\leadsto$ "real world" applications
- "Intelligent behavior" requires background knowledge
  $\leadsto$ knowledge-based systems

80's: AI becomes an industry, neural networks return

End of the 80's: Disillusionment of expert systems, "AI winter"

Since 90's: Probabilistic methods, agent based view, learning, web, realistic applications