Introduction to Artificial Intelligence

CSCE 476-876, Spring 2002

URL: www.cse.unl.edu/~choueiry/S02-476-876

Berthe Y. Choueiry (Shu-we-ri)
choueiry@cse.unl.edu
Tel: (402)476-5444

Outline

- Overview and administrative trivia: look up syllabus
- What is AI?

Pretest

- Scheduled for Friday, Jan 18, 2002
- One part to be completed in the class: crib sheet policy
- One part to be completed at home: collaboration, discussion strictly forbidden
- Content: basic knowledge of mathematics, logic, algorithm, data structure, complexity

Topics

- Lisp (quickly)
- Intelligent agents
- Search
- Constraint satisfaction
- Logical systems

If time allows:

- Planning systems
- Uncertainty: probability and decision theory
- Learning, perception, robotics, NLP, philosophical issues

Goal of AI

- Understand intelligent entities (reasoning mechanisms)
- Build intelligent entities (systems)
  contrast with cognitive science and philosophy

→ Build computers with human-level intelligence.. or better
(human reasoning exhibits systematic errors)

Using: slow, tiny brain, biological or electronic

In order to: perceive, understand, predict and manipulate a far more complex world

Proof of feasibility: human beings
just look in the mirror :–)

New discipline, old topic

AI is a new discipline (vs. physics):
- term coined in 1956 by John McCarthy
- task is enormous, opportunities are wide, easy to make a difference
- Einstein is (probably) yet to come

Study of Intelligence is an old topic. Philosophy: learned but speculative

Advent of computers introduced a new experimental and theoretical discipline: theories can now be tested
⟶ out of the armchair, into the fire

Early Systems were naive (rule-based, etc.)
Paradigms are getting more difficult, elaborate, richer, more subtle

Focus and fields

General:

- perception
- logical reasoning

Specific: (task oriented)

- chess
- proving mathematical theorems
- pun writing
- diagnosing diseases
- planning/scheduling tasks of building construction

A truly universal field

Often scientists/engineers become AI researchers:
want to formalize, systematize, automate the intellectual tasks they are trained to carry out (electrical engineers, civil engineers, medical doctors)

Sometimes, AI researchers delve into specific fields to apply their methods (biology, power systems)

"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)

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Dimensions for classification:

Vertical: concern, focus of efforts

→ thought process and reasoning

→ behavior and action

Horizontal: evaluation of success

→ against human performance

→ against ideal concepts of intelligence

Rationality = do the right think

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

Classification contrast human & rationality:
- Human: empirical science, hypothesis and experimental
- Rationality: mathematics + engineering

No right/wrong, all four approaches are valuable