
Introduction to CMPE 540 Principles of AI Course

1

About the Course

- **Instructor:** Prof. Dr. H. Levent AKIN
- **Phone:** + 90 (212) 359 67 69
- **E-mail:** akin@boun.edu.tr
- **Course Web Site:**
<http://www.cmpe.boun.edu.tr/~akin/cmpe540.html>
- **Course Mailing List**
cmpe540@cmpe.boun.edu.tr

2

Books

- **Text Book:**
 - S. Russell and P. Norvig, "Artificial Intelligence A Modern Approach", 2/e, Prentice Hall, 2003

3

Grading

- | | |
|--------------|----------|
| ■ 2 Midterms | 20% each |
| ■ Final Exam | 30% |
| ■ Projects | 30% |

4

50 and Going Strong !



5

AI's death Highly Exaggerated!

"There's this stupid myth out there that A.I. has failed, but A.I. is everywhere around you every second of the day. People just don't notice it. You've got A.I. systems in cars, tuning the parameters of the fuel injection systems. When you land in an airplane, your gate gets chosen by an A.I. scheduling system. Every time you use a piece of Microsoft software, you've got an A.I. system trying to figure out what you're doing, like writing a letter, and it does a pretty damned good job. Every time you see a movie with computer-generated characters, they're all little A.I. characters behaving as a group. Every time you play a video game, you're playing against an A.I. system. "

Rodney Brooks, director of the MIT AI Lab



6

Why study AI?



Labor



Science



Search engines



Medicine/
Diagnosis



Appliances



7

Asimo



8

Robocup League




9

Rescue Simulation



10

What is "Intelligence" ?



11

Capabilities of Intelligent Beings

- Thinking and problem solving
- Learning and memory
- Language
- Intuition and creativity
- Consciousness
- Emotions
- Surviving in a complex world
- Perceptual and motor abilities

12

Hardware



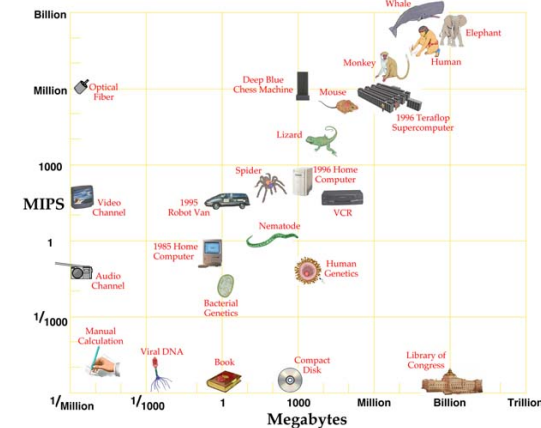
10^{11} neurons
 10^{14} synapses
 cycle time: 10^{-3} sec



10^7 transistors
 10^{10} bits of RAM
 cycle time: 10^{-9} sec


Computer vs. Brain

All Thinks, Great and Small



From ROBOT, Moravec, Oxford, 1998

Evolution of Computers



Moravec

- In near future we can have computers with as many processing elements as our brain, but:
 - far fewer interconnections (wires or synapses)
 - much faster updates.
- Fundamentally different hardware may require fundamentally different algorithms!
 - Very much an open question.
 - Artificial neural networks research.

What is AI?

<ul style="list-style-type: none"> ■ The exciting new effort to make computers think ... <i>machine with minds</i>, in the full and literal sense" (Haugeland 1985) 	<ul style="list-style-type: none"> ■ "The study of mental faculties through the use of computational models" (Charniak et al. 1985)
<ul style="list-style-type: none"> • "The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) 	<ul style="list-style-type: none"> • A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkol, 1990)


Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

17

Acting humanly: The Turing test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" ⇒ "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game



- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning

18

Turing Test

- The computer would need to possess the following capabilities:
 - **natural language processing** to enable it to communicate successfully in English (or some other human language);
 - **knowledge representation** to store information provided before or during the interrogation;
 - **automated reasoning** to use the stored information to answer questions and to draw new conclusions;
 - **machine learning** to adapt to new circumstances and to detect and extrapolate patterns.

19

Acting Humanly: The Full Turing Test

- Problem:
 1. Turing test is not reproducible, constructive, and amenable to mathematic analysis.
 2. What about physical interaction with interrogator and environment?
- **Total Turing Test:** Requires physical interaction and needs perception and actuation.


20

What would a computer need to pass the Total Turing test?

- **Vision:** to recognize the examiner's actions and various objects presented by the examiner.
- **Motor control:** to act upon objects as requested.
- **Other senses:** such as audition, smell, touch, etc.

21

Total Turing Test ?




Repliee Q1

<http://www.androidscience.com/links.html>

22


The Chinese Room



23

The Chinese Room

■ The Chinese Room



24

Thinking humanly: Cognitive Science

- 1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism
- Requires scientific theories of internal activities of the brain
 - What level of abstraction? "Knowledge" or "circuits"?
 - How to validate? Requires
 - 1) Predicting and testing behavior of human subjects (top-down) or
 - 2) Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI
- Both share with AI the following characteristic:
 - the available theories do not explain (or engender) anything resembling human-level general intelligence

25

Thinking Rationally: Laws of Thought

- Aristotle (~ 450 B.C.) attempted to codify "right thinking"
What are correct arguments/thought processes?
- e.g., "Socrates is a man, all men are mortal; therefore Socrates is mortal"
- Several Greek schools developed various forms of logic:
notation plus rules of derivation for thoughts.

26

Thinking Rationally: Laws of Thought

■ Problems:

- 1) Uncertainty: Not all facts are certain (e.g., *the flight might be delayed*).
- 2) Resource limitations:
 - Not enough time to compute/process
 - Insufficient memory/disk/etc
 - Etc.

27

Acting Rationally: The Rational Agent

- Rational behavior: Doing the right thing!
- The right thing: That which is expected to maximize the expected return
- Provides the most general view of AI because it includes:
 - Correct inference ("Laws of thought")
 - Uncertainty handling
 - Resource limitation considerations (e.g., reflex vs. deliberation)
 - Cognitive skills (NLP, AR, knowledge representation, ML, etc.)
- Advantages:
 - More general
 - Its goal of rationality is well defined

28

Rational agents

- An **agent** is just something that acts (*agent* comes from the Latin *agere*, to do).
- Attributes that distinguish **agents** from mere “programs” include:
 - operating under autonomous control,
 - perceiving their environment,
 - persisting over a prolonged time period,
 - adapting to change, and
 - being capable of taking on another’s goals.
- This course is about designing rational agents.
- A **rational agent** is one that acts so as to achieve the best outcome, or when there is uncertainty, the best expected outcome.

29

AI Prehistory

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

30

AI History

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing’s “Computing Machinery and Intelligence”
- 1952–69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel’s checkers program, Newell & Simon’s Logic Theorist, Gelernter’s Geometry Engine
- 1956 Dartmouth meeting: “Artificial Intelligence” adopted
- 1965 Robinson’s complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: “AI Winter”
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth
“Nouvelle AI”: ALife, GAs, soft computing
- 1995– Agents, agents, everywhere . . .
- 2003– Human-level AI back on the agenda

31

AI State of the art

- Have the following been achieved by AI?
 - Play a decent game of table tennis
 - Drive safely along a curving mountain road
 - Drive safely along Hisarüstü ?
 - Buy a week’s worth of groceries on the web
 - Buy a week’s worth of groceries at a hipermarket
 - Play a decent game of bridge
 - Discover and prove a new mathematical theorem
 - Design and execute a research program in molecular biology
 - Write an intentionally funny story
 - Give competent legal advice in a specialized area of law
 - Translate spoken English into spoken Turkish in real time
 - Converse successfully with another person for an hour
 - Perform a complex surgical operation
 - Unload any dishwasher and put everything away

32

Recurrent Themes

■ Neural nets vs AI

- McCulloch & Pitts 1943
- Died out in 1960's, revived in 1980's
 - Neural nets vastly simplified model of real neurons, but still useful & practical – massive parallelism
 - particular family of learning and representation techniques

■ Logic vs Probability

- In 1950's logic seemed more computationally & expressively attractive (McCarthy, Newell)
 - attempts to extend logic “just a little” to deal with the fact that the world is uncertain!
- 1988 – Judea Pearl's work on Bayes nets
 - provided efficient computational framework
- Today – no longer rivals
 - hot topic: combining probability & first-order logic