

Artificial Intelligence

Engr. Madeha Mushtaq
Department of Computer Science
Iqra National University

What is intelligence?


- The ability to acquire and apply knowledge and skills.
- The ability to solve problems.
- The ability to perceive, interpret and understand.
- The ability to learn, emotional knowledge, reason, plan, creativity, critical thinking

Artificial Intelligence

- Historically four approaches to artificial intelligence have been followed, each by different people with different methods:
 - Acting Humanly
 - Thinking Humanly
 - Acting rationally
 - Thinking rationally

Acting humanly: Turing Test

- Modeling exactly how humans actually act.
- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game



- A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer.

Acting humanly: Turing Test

- The computer would need to possess the following capabilities:
 - natural language processing to enable it to communicate successfully in English;
 - knowledge representation to store what it knows or hears;
 - 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.

Thinking Humanly

- “The exciting new effort to make computers think ...machines with minds, in the full and literal sense.
- [The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning.
- Modeling exactly how humans actually think, cognitive models of human reasoning.

Acting rationally: rational agent

- Rational behavior: Doing that which is expected to maximize one's "utility function" in this world.
- An agent is an entity that perceives and acts. A rational agent acts rationally.
- This course is about designing rational agents
- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance.

Thinking Rationally

- “The study of mental faculties through the use of computational models.”
- “The study of the computations that make it possible to perceive, reason, and act.”
- Modeling how ideal agents “should think”
 - models of “rational” thought (formal logic)
 - note: humans are often not rational!

History of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1952—69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity
Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans
- Best vehicle in Darpa challenge made it 7 miles into the desert...

Consider what might be involved in building a “smart” computer....

- What are the “components” that might be useful?
 - Fast hardware?
 - Foolproof software?
 - Chess-playing at grandmaster level?
 - Speech interaction?
 - speech synthesis
 - speech recognition
 - speech understanding
 - Image recognition and understanding ?
 - Learning?
 - Planning and decision-making?

Can we build hardware as complex as the brain?

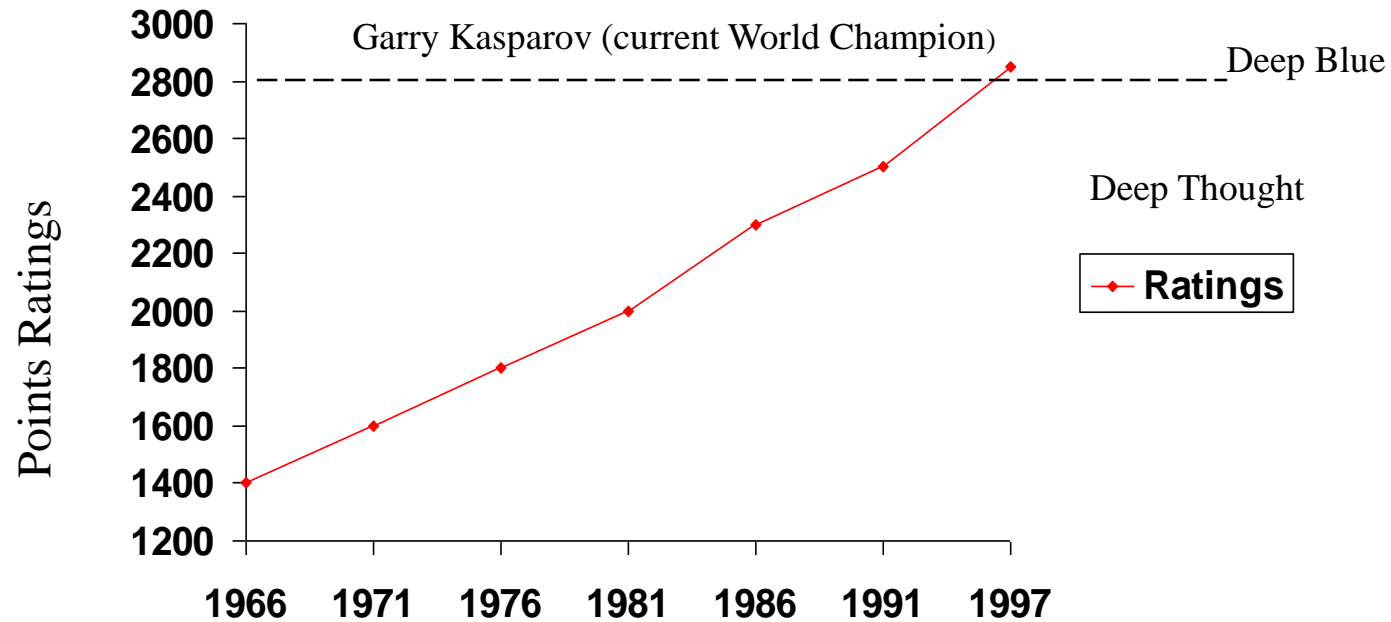
- How complicated is our brain?
 - a neuron, or nerve cell, is the basic information processing unit
 - estimated to be on the order of 10^{12} neurons in a human brain
 - many more synapses (10^{14}) connecting these neurons
 - cycle time: 10^{-3} seconds (1 millisecond)
- How complex can we make computers?
 - 10^6 or more transistors per CPU
 - supercomputer: hundreds of CPUs, 10^9 bits of RAM
 - cycle times: order of 10^{-8} seconds
- Conclusion
 - **YES:** in the near future we can have computers with as many basic processing elements as our brain, but with
 - far fewer interconnections (wires or synapses) than the brain
 - much faster updates than the brain
 - **but** building hardware is very different from making a computer behave like a brain!

Must an Intelligent System be Foolproof?

- A “foolproof” system is one that never makes an error:
 - Types of possible computer errors
 - hardware errors, e.g., memory errors
 - software errors, e.g., coding bugs
 - “human-like” errors
 - Clearly, hardware and software errors are possible in practice
 - what about “human-like” errors?
- An intelligent system can make errors and still be intelligent
 - humans are not right all of the time
 - we learn and adapt from making mistakes
 - e.g., consider learning to surf or ski
 - we improve by taking risks and falling
 - an intelligent system can learn in the same way
- Conclusion:
 - **NO**: intelligent systems will not (and need not) be foolproof

Can Computers play Humans at Chess?

- Chess Playing is a classic AI problem
 - well-defined problem
 - very complex: difficult for humans to play well



- Conclusion: YES: today's computers can beat even the best human

Can Computers Talk?

- This is known as “speech synthesis”
 - translate text to phonetic form
 - e.g., “fictitious” -> fik-tish-es
 - use pronunciation rules to map phonemes to actual sound
 - e.g., “tish” -> sequence of basic audio sounds
- Difficulties
 - sounds made by this “lookup” approach sound unnatural
 - sounds are not independent
 - e.g., “act” and “action”
 - modern systems (e.g., at AT&T) can handle this pretty well
 - a harder problem is emphasis, emotion, etc
 - humans understand what they are saying
 - machines don't: so they sound unnatural
- Conclusion: NO, for complete sentences, but YES for individual words

Can Computers Recognize Speech?

- Speech Recognition:
 - mapping sounds from a microphone into a list of words.
 - Hard problem: noise, more than one person talking, occlusion, speech variability,..
 - Even if we recognize each word, we may not understand its meaning.
- Recognizing single words from a small vocabulary
 - systems can do this with high accuracy (order of 99%)
 - e.g., directory inquiries
 - limited vocabulary (area codes, city names)
 - computer tries to recognize you first, if unsuccessful hands you over to a human operator
 - saves millions of dollars a year for the phone companies

Recognizing human speech (ctd.)

- Recognizing normal speech is much more difficult
 - speech is continuous: where are the boundaries between words?
 - e.g., “John’s car has a flat tire”
 - large vocabularies
 - can be many thousands of possible words
 - we can use **context** to help figure out what someone said
 - try telling a waiter in a restaurant:
“I would like some cream and sugar in my coffee”
 - background noise, other speakers, accents, colds, etc
 - on normal speech, modern systems are only about 60% accurate
- Conclusion: NO, normal speech is too complex to accurately recognize, but YES for restricted problems
 - (e.g., recent software for PC use by IBM, Dragon systems, etc)

Can Computers Understand speech?


- Understanding is different to recognition:
 - “Time flies like an arrow”
 - assume the computer can recognize all the words
 - but how could it understand it?
 - 1. time passes quickly like an arrow?
 - 2. command: time the flies the way an arrow times the flies
 - 3. command: only time those flies which are like an arrow
 - 4. “time-flies” are fond of arrows
 - only 1. makes any sense, but how could a computer figure this out?
 - clearly humans use a lot of implicit commonsense knowledge in communication
- Conclusion: NO, much of what we say is beyond the capabilities of a computer to understand at present

Can Computers Learn and Adapt ?

- Learning and Adaptation
 - consider a computer learning to drive on the freeway
 - we could code lots of rules about what to do
 - or we could let it drive and steer it back on course when it heads for the embankment
 - systems like this are under development (e.g., Daimler Benz)
 - e.g., RALPH at CMU
 - in mid 90's it drove 98% of the way from Pittsburgh to San Diego without any human assistance
 - Machine learning allows computers to learn to do things without explicit programming
- Conclusion: YES, computers can learn and adapt, when presented with information in the appropriate way

Can Computers “see”?

- Recognition v. Understanding (like Speech)
 - Recognition and Understanding of Objects in a scene
 - look around this room
 - you can effortlessly recognize objects
 - human brain can map 2d visual image to 3d “map”
- Why is visual recognition a hard problem?



- Conclusion: mostly NO: computers can only “see” certain types of objects under limited circumstances: but YES for certain constrained problems (e.g., face recognition)

Can Computers plan and make decisions?

- Intelligence
 - involves solving problems and making decisions and plans
 - e.g., you want to visit your cousin in Boston
 - you need to decide on dates, flights
 - you need to get to the airport, etc
 - involves a sequence of decisions, plans, and actions
- What makes planning hard?
 - the world is not predictable:
 - your flight is canceled or there's a backup on the 405
 - there are a potentially huge number of details
 - do you consider all flights? all dates?
 - no: commonsense constrains your solutions
 - AI systems are only successful in constrained planning problems
- Conclusion: NO, real-world planning and decision-making is still beyond the capabilities of modern computers
 - exception: very well-defined, constrained problems: mission planning for satellites.

Summary of State of AI Systems in Practice

- Speech synthesis, recognition and understanding
 - very useful for limited vocabulary applications
 - unconstrained speech understanding is still too hard
- Computer vision
 - works for constrained problems (hand-written zip-codes)
 - understanding real-world, natural scenes is still too hard
- Learning
 - adaptive systems are used in many applications: have their limits
- Planning and Reasoning
 - only works for constrained problems: e.g., chess
 - real-world is too complex for general systems
- Overall:
 - many components of intelligent systems are “doable”
 - there are many interesting research problems remaining

Intelligent Systems in Your Everyday Life

- Post Office
 - automatic address recognition and sorting of mail
- Banks
 - automatic check readers, signature verification systems
 - automated loan application classification
- Telephone Companies
 - automatic voice recognition for directory inquiries
 - automatic fraud detection,
 - classification of phone numbers into groups
- Credit Card Companies
 - automated fraud detection, automated screening of applications
- Computer Companies
 - automated diagnosis for help-desk applications


AI Applications: Consumer Marketing

- Have you ever used any kind of credit/ATM/store card while shopping?
 - if so, you have very likely been “input” to an AI algorithm
- All of this information is recorded digitally
- Companies like Nielsen gather this information weekly and search for patterns
 - general changes in consumer behavior
 - tracking responses to new products
 - identifying customer segments: targeted marketing, e.g., they find out that consumers with sports cars who buy textbooks respond well to offers of new credit cards.
 - Currently a very hot area in marketing
- How do they do this?
 - Algorithms (“data mining”) search data for patterns
 - based on mathematical theories of learning
 - completely impractical to do manually

AI Applications: Identification Technologies

- ID cards
 - e.g., ATM cards
 - can be a nuisance and security risk:
 - cards can be lost, stolen, passwords forgotten, etc
- Biometric Identification
 - walk up to a locked door
 - camera
 - fingerprint device
 - microphone
 - computer uses your biometric signature for identification
 - face, eyes, fingerprints, voice pattern

AI Applications: Predicting the Stock Market



- The Prediction Problem
 - given the past, predict the future
 - very difficult problem!
 - we can use learning algorithms to learn a predictive model from historical data
 - $\text{prob}(\text{increase at day } t+1 \mid \text{values at day } t, t-1, t-2, \dots, t-k)$
 - such models are routinely used by banks and financial traders to manage portfolios worth millions of dollars

AI-Applications: Machine Translation

- Language problems in international business
 - e.g., at a meeting of Japanese, Korean, Vietnamese and Swedish investors, no common language
 - or: you are shipping your software manuals to 127 countries
 - solution; hire translators to translate
 - would be much cheaper if a machine could do this!
- How hard is automated translation
 - very difficult!
 - e.g., English to Russian
 - “The spirit is willing but the flesh is weak” (English)
 - “the vodka is good but the meat is rotten” (Russian)
 - not only must the words be translated, but their meaning also!
- Nonetheless....
 - commercial systems can do alot of the work very well (e.g.,restricted vocabularies in software documentation)
 - algorithms which combine dictionaries, grammar models, etc.
 - see for example babelfish.altavista.com

End of Slides