Introduction to Artificial Intelligence

Ali Ridho Barakbah

Intelligent Knowledge Creation (iKnow) Research Group
Department of Information and Computer Engineering
Politeknik Elektronika Negeri Surabaya

Intelligence?

Intelligence (also called intellect) is an umbrella term used to describe a property of the mind that encompasses many related abilities, such as the capacities to reason, to plan, to solve problems, to think abstractly, to comprehend ideas, to use language, and to learn (Wikipedia)

The capacity to learn and solve problems (Websters dictionary)

Intelligence

- Ability to solve novel problems
- Ability to act rationally
- Ability to act like humans

What's involved in Intelligence?

- Ability to interact with the real world
  - to perceive, understand, and act
    e.g., speech recognition and understanding and synthesis
    e.g., image understanding
    e.g., ability to take actions, have an effect

- Reasoning and Planning
  - modeling the external world, given input
  - solving new problems, planning, and making decisions
  - ability to deal with unexpected problems, uncertainties

- Learning and Adaptation
  - we are continuously learning and adapting
  - our internal models are always being "updated"
    - e.g., a baby learning to categorize and recognize animals

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Artificial Intelligence?

- Building systems that think like humans
  - "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, … such as decision-making, problem solving, learning, …" (Bellman, 1978)

- Building systems that act like humans
  - "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)

Artificial Intelligence?

- Building systems that think rationally
  - "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)

- Building systems that act rationally
  - "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)

What is Artificial Intelligence?
(John McCarthy, Stanford University)

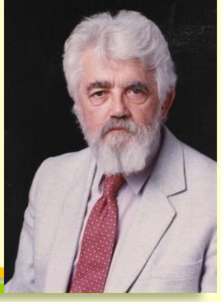

- **What is artificial intelligence?**
  It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

- **Yes, but what is intelligence?**
  Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.

- **Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?**
  Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.

- **More in: http://www-formal.stanford.edu/jmc/whatisai/node1.html**

History of AI

- 1943: early beginnings
  - McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing
  - Turing's "Computing Machinery and Intelligence“
- 1956: birth of AI
  - Dartmouth meeting: "Artificial Intelligence“ name adopted
- 1950s: initial promise
  - Early AI programs, including
  - Samuel's checkers program
  - Newell & Simon's Logic Theorist
- 1955-65: “great enthusiasm”
  - Newell and Simon: GPS, general problem solver
  - Gelertner: Geometry Theorem Prover
  - McCarthy: invention of LISP

History of AI

- 1966—73: Reality dawns
  - Realization that many AI problems are intractable
  - Limitations of existing neural network methods identified
    - Neural network research almost disappears
- 1969—85: Adding domain knowledge
  - Development of knowledge-based systems
  - Success of rule-based expert systems,
    - E.g., DENDRAL, MYCIN
    - But were brittle and did not scale well in practice
- 1986-- Rise of machine learning
  - Neural networks return to popularity
  - Major advances in machine learning algorithms and applications
- 1990-- Role of uncertainty
  - Bayesian networks as a knowledge representation framework
- 1995-- AI as Science
  - Integration of learning, reasoning, knowledge representation
  - AI methods used in vision, language, data mining, etc

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Success Stories

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- AI program proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- 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
- Robot driving: DARPA grand challenge 2003-2007
- 2006: face recognition software available in consumer cameras

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Different Types of Artificial Intelligence

1 Systems that think like humans	3 Systems that think rationally
2 Systems that act like humans	4 Systems that act rationally

1

Systems that think like humans

- How do humans think?
  - Requires scientific theories of internal brain activities (cognitive model):
    - Level of abstraction? (knowledge or circuitry?)
    - Validation?
      - Predicting and testing human behavior
      - Identification from neurological data
  - Cognitive Science vs. Cognitive neuroscience
- Both approaches are now distinct from AI
- Share that the available theories do not explain anything resembling human intelligence.
  - Three fields share a principal direction.

Source: Tom Lenaerts, Artificial Intelligence I: Introduction, Vrije Universiteit Brussel.

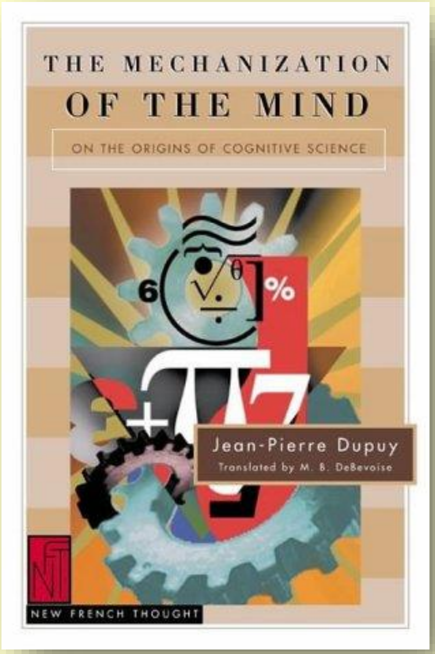
THE MECHANIZATION
OF THE MIND
ON THE ORIGINS OF COGNITIVE SCIENCE
Jean-Pierre Dupuy
Translated by M. B. DeBevoise
NEW FRENCH THOUGHT

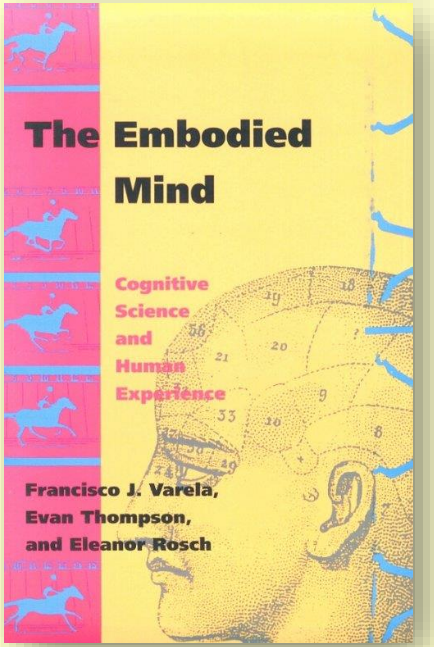
The Embodied
Mind
Cognitive
Science
and
Human
Experience
Francisco J. Varela,
Evan Thompson,
and Eleanor Rosch

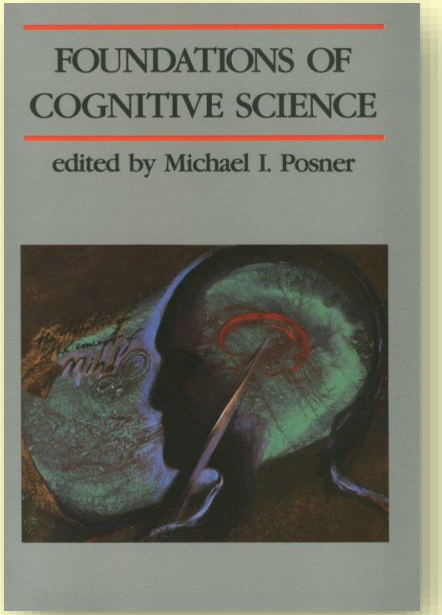
FOUNDATIONS OF
COGNITIVE SCIENCE
edited by Michael I. Posner

Thinking humanly

- Cognitive Science approach
  - Try to get "inside" our minds
  - E.g., conduct experiments with people to try to "reverse-engineer" how we reason, learning, remember, predict

- Problems
  - Humans don't behave rationally
    - e.g., insurance
  - The reverse engineering is very hard to do
  - The brain's hardware is very different to a computer program

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

2 Systems that act like humans

- When does a system behave intelligently?
  - Turing (1950) *Computing Machinery and Intelligence*
  - Operational test of intelligence: imitation game

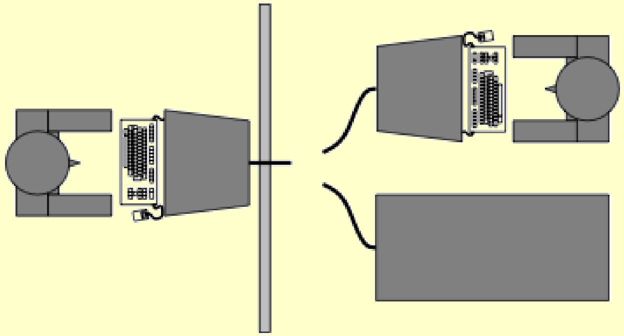

  - Test still relevant now, yet might be the wrong question.
  - Requires the collaboration of major components of AI: knowledge, reasoning, language understanding, learning, …

Source: Tom Lenaerts, Artificial Intelligence I: Introduction, Vrije Universiteit Brussel.

Acting humanly

- "Can machines think?" → "Can machines behave intelligently?“
- Suggests major components required for AI:
  - knowledge representation
  - reasoning,
  - language/image understanding,
  - learning

* Question: is it important that an intelligent system act like a human?

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

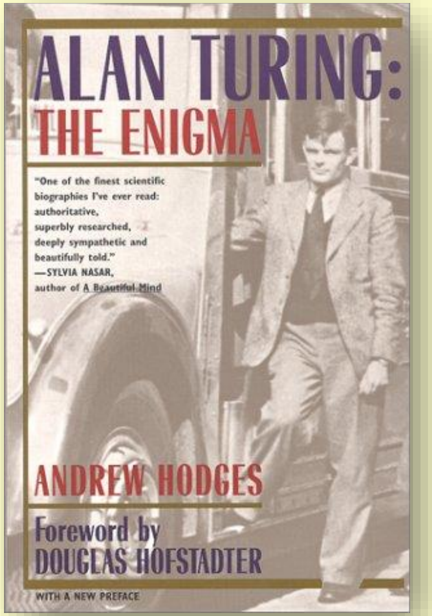
ALAN TURING:
THE ENIGMA
"One of the finest scientific biographies I've ever read: authoritative, superbly researched, deeply sympathetic and beautifully told."
—SYLVIA NASAR, author of A Beautiful Mind
ANDREW HODGES
Foreword by
DOUGLAS HOFSTADTER
WITH A NEW PREFACE

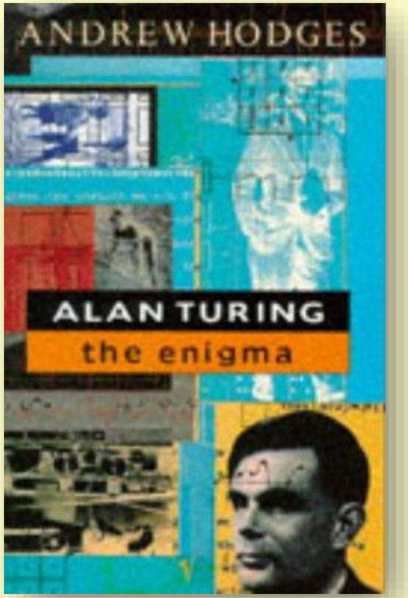
ANDREW HODGES
ALAN TURING
the enigma

3 Systems that think rationally

- Capturing the laws of thought
  - Aristotle: What are 'correct' argument and thought processes?
    - Correctness depends on irrefutability of reasoning processes.
  - This study initiated the field of logic.
    - The logicist tradition in AI hopes to create intelligent systems using logic programming.
  - Problems:
    - Not all intelligence is mediated by logic behavior
    - What is the purpose of thinking? What thought should one have?

Source: Tom Lenaerts, Artificial Intelligence I: Introduction, Vrije Universiteit Brussel.

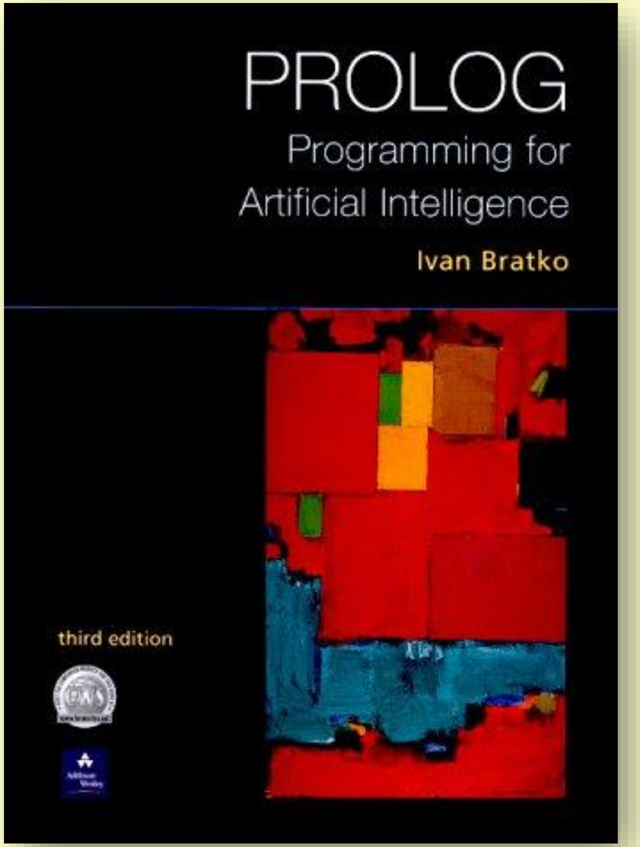
PROLOG
Programming for
Artificial Intelligence
Ivan Bratko
third edition

Thinking rationally

- Represent facts about the world via logic
- Use logical inference as a basis for reasoning about these facts
- Can be a very useful approach to AI
  - E.g., theorem-provers

- Limitations
  - Does not account for an agent's uncertainty about the world
    - E.g., difficult to couple to vision or speech systems
  - Has no way to represent goals, costs, etc (important aspects of real-world environments)

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

4 Systems that act rationally

- Rational behavior: “doing the right thing”
- The “Right thing” is that what is expected to *maximize goal achievement given the available information*.
- Can include thinking, yet in service of rational action.
    - Action without thinking: e.g. reflexes.

- Two advantages over previous approaches:
    - More general than law of thoughts approach
    - More amenable to scientific development.
- Yet rationality is only applicable in *ideal* environments.
- Moreover rationality is not a very good model of reality.

Source: Tom Lenaerts, Artificial Intelligence I: Introduction, Vrije Universiteit Brussel.

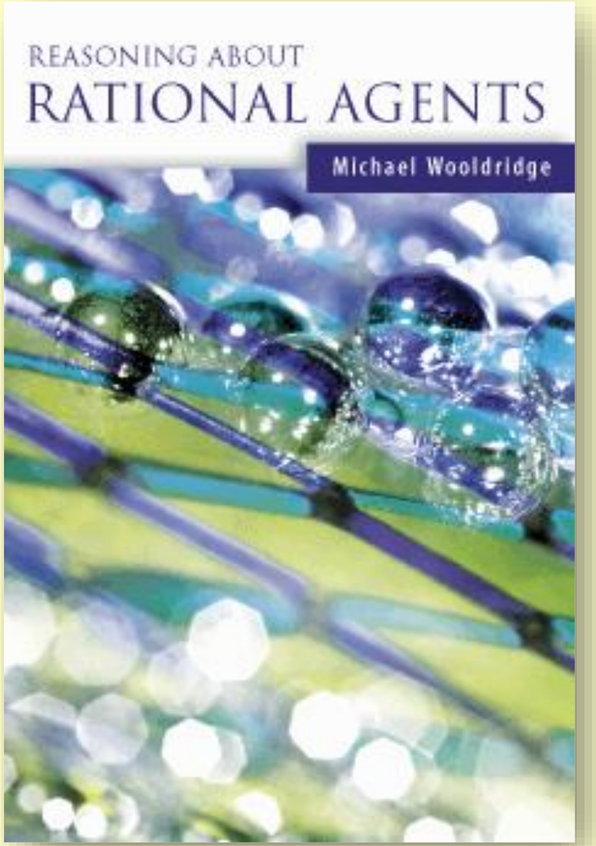
REASONING ABOUT
RATIONAL AGENTS
Michael Wooldridge

Acting rationally

- Decision theory/Economics
  - Set of future states of the world
  - Set of possible actions an agent can take
  - Utility = gain to an agent for each action/state pair
  - An agent acts rationally if it selects the action that maximizes its “utility”
    - Or expected utility if there is uncertainty
- Emphasis is on autonomous agents that behave rationally (make the best predictions, take the best actions)
  - on average over time
  - within computational limitations (“bounded rationality”)

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Academic Disciplines relevant to AI

- 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/Statistics → modeling uncertainty, learning from data
- Economics → utility, decision theory, rational economic agents
- Neuroscience → neurons as information processing units.
- Psychology/Cognitive Science → how do people behave, perceive, process cognitive

  information, represent knowledge
- Computer engineering → building fast computers

  Control theory → design systems that maximize an objective function over time
- Linguistics → knowledge representation, grammars

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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?
  - 10^8 or more transistors per CPU
  - supercomputer: hundreds of CPUs, 10^{12} bits of RAM
  - cycle times: order of 10^{-9} 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!

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Can Computers beat Humans at Chess?

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

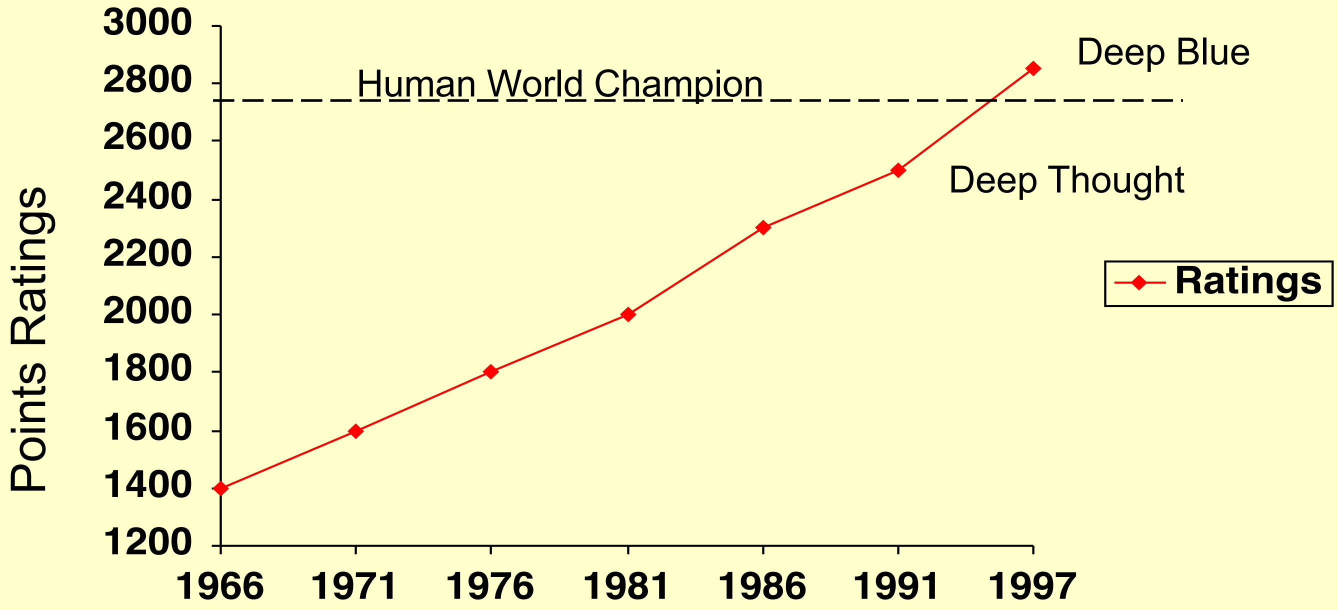

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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
  - YES, for individual words

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Can Computers Recognize Speech?

- Speech Recognition:
  - mapping sounds from a microphone into a list of words
  - classic problem in AI, very difficult
    - "Lets talk about how to wreck a nice beach"
    - (I really said "______________________________")
- 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

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Can Computers Understand speech?

- Understanding is different to recognition:
  - "Time flies like an arrow"
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 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

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Can Computers Learn and Adapt ?

- Learning and Adaptation
  - consider a computer learning to drive on the freeway
  - we could teach it 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
  - many successful applications:
    - requires some "set-up": does not mean your PC can learn to forecast the stock market or become a brain surgeon

- Conclusion: YES, computers can learn and adapt, when presented with information in the appropriate way

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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?

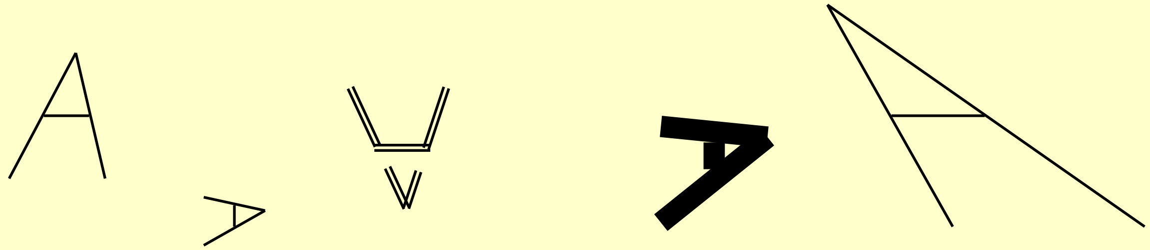

- Conclusion:
  - mostly NO: computers can only "see" certain types of objects under limited circumstances
  - YES for certain constrained problems (e.g., face recognition)

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Can computers plan and make optimal decisions?

- Intelligence
  - involves solving problems and making decisions and plans
  - e.g., you want to take a holiday in Brazil
    - 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

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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
- Customer Service
  - automatic voice recognition
- The Web
  - Identifying your age, gender, location, from your Web surfing
  - Automated fraud detection
- Digital Cameras
  - Automated face detection and focusing
- Computer Games
  - Intelligent characters/agents

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

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!
  - is this problem "AI-complete"?
- Nonetheless....
  - commercial systems can do a lot of the work very well (e.g.,restricted vocabularies in software documentation)
  - algorithms which combine dictionaries, grammar models, etc.
  - Recent progress using "black-box" machine learning techniques

Source: Prof. Padhraic Smyth, Introduction to Artificial Intelligence, Lecture note, Fall 2007, University of California, Irvine.

Learning process

- Supervised learning
- Unsupervised learning
- Reinforcement learning

Supervised Learning

- Supervised learning is the machine learning task of inferring a function from labeled training data.
- The training data consist of a set of training examples.
- In supervised learning, each example is a pair consisting of an input object (typically a vector) and a desired output value (also called the supervisory signal).
- A supervised learning algorithm analyzes the training data and produces an inferred function, which can be used for mapping new examples.
- An optimal scenario will allow for the algorithm to correctly determine the class labels for unseen instances.
- This requires the learning algorithm to generalize from the training data to unseen situations in a "reasonable" way.

Source: Wikipedia

Unsupervised Learning

- In machine learning, the problem of unsupervised learning is that of trying to find hidden structure in unlabeled data.
- Since the examples given to the learner are unlabeled, there is no error or reward signal to evaluate a potential solution.
- Unsupervised learning is closely related to the problem of density estimation in statistics.
- However unsupervised learning also encompasses many other techniques that seek to summarize and explain key features of the data.

Source: Wikipedia

Reinforcement Learning

- Reinforcement learning is an area of machine learning inspired by behaviorist psychology, concerned with how software agents ought to take actions in an environment so as to maximize some notion of cumulative reward.
- Reinforcement learning differs from standard supervised learning in that correct input/output pairs are never presented, nor sub-optimal actions explicitly corrected.
- There is a focus on on-line performance, which involves finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge).

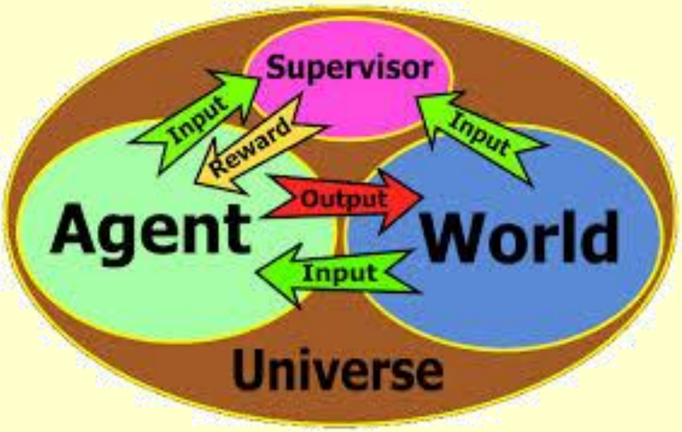

Source: Wikipedia