Artificial Intelligence
Background and Overview

L. Manevitz

Requirements

- Two – three projects (obligatory)
- Final
- Final Grade – between 33% - 67% projects (probably 40 - 50%)
- Late projects will be penalized
- Attendance in Lectures is strongly recommended

Text and References

- Texts: (on reserve)
  - Elaine Rich Artificial Intelligence
  - Nils Nilsson Artificial Intelligence
  - Patrick Winston Artificial Intelligence
  - Newer Books:
  - Russell & Norvig Artificial Intelligence: A Modern Approach
  - Nils Nilsson Artificial Intelligence:A New Synthesis
  - Webber & Nilsson Readings in Artificial Intelligence

Artificial Intelligence-What is it?

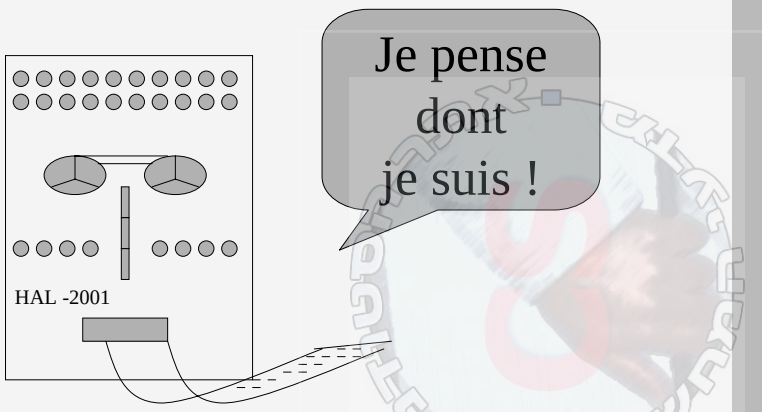

What is intelligence ?

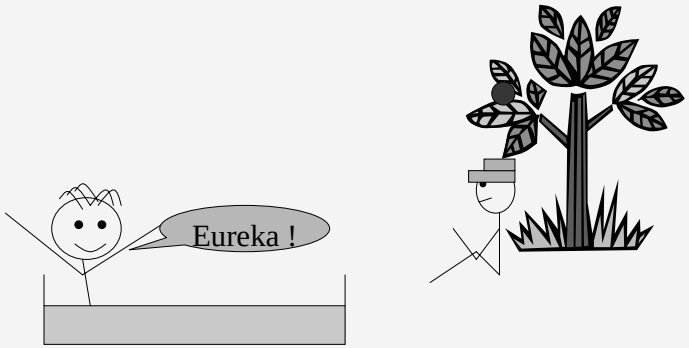

What is Artificial Intelligence?

- How do we define it?
- What is it good for?
- History ?
- Successes/ Failures?
- Future Outlook?

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

Possible Examples?

- Chess?
- Recognizing Credit Card Fraud?
- Automatic Train Driver?
- Automatic Car Driver?
- Recommender System?
- Agents in a Computer Game?
- Boy Robot in movie “AI”?
- Emotions?
- Solving Geometry Problems?
- Proving Theorems?
- Giving Directions?
- Answering Questions for Registering Students?

How do we define it?

- What is artificial intelligence?
- What is natural intelligence?
- How do we know if we achieve it?

Psychology

- How do humans and animals think and act?
- Cognitive Psychology and Cognitive Science
- Behaviorism

Neuroscience

- How does the Brain Work?
- Compare with Computer?
  - 10**11 neurons vs 1 CPU (10**8 gates)
  - 10 **-3 sec vs 10 **-10 sec
  - 10**14 bits/sec vs 10**10 bits/sec
  - Moore's Law (doubles every 1.5 years) CPU gate count will equal neurons in 2020.
  - Does this mean anything?

Mathematics

- Formal Rules to Draw Conclusions
- What can be Computed?
  - Godel's Theorem
  - NP completeness
- How to Reason with Uncertain Information?
- Probability
- Game Theory

Are the Means Important?

- Black Box – raw power versus "cleverness"
- Or White Box - somehow models how people work
- Or White Box – somehow cleverness sneaks in by some over-riding idea?

Is Adaptivity (learning) Crucial?

Artificial Intelligence

- Goals.
- Methods.
- Examples.

Goals in A.I.

- Understand thinking.
- Make computers perform tasks that require intelligence if performed by people.

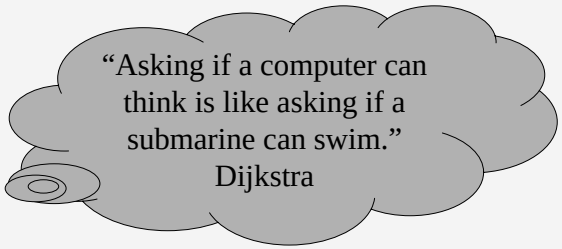

Goals in A.I.

Main Goal :

Artificial intelligence.

Sub Goal :

Understanding human intelligence and how it is possible.

Related subjects :

Neurophysiology
Cognitive Psychology
(Artificial Neural Networks)

AI Supplements Philosophy, Psychology, Linquistics, etc.

1) Use of computer metaphors has led to rich language for talking and thinking about thinking.
2) Computer models force precision.
3) Computer implementations quantify task requirements.
4) Computer programs can be experimented on in ways that animal brains can not.

Aspects

- Engineering :
  Solve "real world" problems using ideas about knowledge representation and handling.
- Scientific :
  Discover ideas about knowledge that helps explain various orts of intelligence.

State of the Art

- NASA Mars Robots (planning program)
- Game Playing: Deep Blue & Junior, etc
- Autonomous Control and Learning
  - Alvinn (CMU) drove across the USA (98% of the time)

  Medical Diagnosis Programs –state of art

  AI Planning handles US logistics

  Robotics Assistant Surgery

  Language Understanding and Problem Solvers

  Recognizing Cognitive Actions by Looking at Brain Scans!

Applications

- Farming robots, Manufacturing, Medical, Household.
- Data mining, Scheduling, Risk Management Control, Agents.
- Internet + AI : natural laboratory

Underlying Hypothesis

- Physical Symbol Hypothesis (Newell) :
  A physical symbol system has the necessary and sufficient means for intelligent action.

This hypothesis means that we can hope to implement this in the computer.

- Note : Use of term "intelligent action" not "intelligence". Compare with Searle "Chinese Room".

Definitions of AI

- "Science of making machines do things that would require intelligence if done by man."
  M. Minsky
- "… make computers more useful and to understand principles that make intelligence possible"
  P. Winston

Definitions of AI cont.

- "… main tenet that there are common processes underlie thinking … these can be understood and studied scientifically … unimportant who is doing thinking – man or computer. This is an implementation detail."
  N. Nilsson

Some History

- Literature: Golem, Frankenstein, Odysseus, Asimov, …
- Rules of Thought: Greeks (Aristotle, correctness of proofs), Formal Systems (Aristotle, Saadia Gaon), Leibniz, Boole, Godel, Turing.

Note two aspects: Physical and Mental (corresponds to Robotics and AI today)

Some History

- Babylonians
- Greeks – Plato, Aristotle, Greek Mythology
- Arabic Culture – Saadia Gaon, AlKhwarzi
- Frankenstein, Golem
- Analytical Engine, Babbage, Lovelace
- Mechanical Calculation and Mechanical Proof - Pascal, Leibniz, Hilbert
- WWII : Turing, von Neumann, Godel, ACE , Einiac
- Artificial Neuron
- Dartmouth (Modern AI)
- Distributed Agents, Internet
- Machine Learning

- 1940s Turing, Shannon, Von Neumann.
- 1950s –1960s Learning Machines; Naïve Translators,Naïve Chess Programs, (Simon's 10 year prediction) Perceptrons.
- 1960s – 1970s MIT, Stanford,Carnegie-Mellon,(Minsky, McCarthy, Simon) General Purpose Algorithms.
- 1980s Multi-level Perceptrons,Expert Systems, Knowledge Based Systems, Logical AI, Uncertainty Reasoning.
- 1990s NN applications, Theory of Learning, Agent Paradigm,Internet Applications,Space Robots,Computer Chess Champion.
- 2000s SVM and Kernel Learning, Mixed applications

Spin-offs of A.I.

- The Computer.
- Formal Mathematical Logic.
- Much of Mathematics.
- Time Sharing.
- Computer Languages.
- Computer Vision.
- Expert Systems
- Theory of Learning

Spin-offs of A.I. cont.

- Expert Systems
- Data Mining.
- Soft bots.
- Robotics.
- Video Display.
- Information retrieval.

Problems addressed by A.I.

- Game playing.
- Theorem proving.
- Perception (Vision, Speech).
- Natural Language Understanding.

Problems addressed by A.I. cont.

- Expert Problem Solving :
  - Symbolic Mathematics.
  - Medical Diagnosis.
  - Chemical Analysis.
  - Engineering Design.
- Intelligent Agents.
- Automated Negotiation.
- Data Mining.
- Web Search.

Artificial Intelligence

- Expert Systems.
- Vision.
- Speech.
- Language.
- Games.
- Planning and Action.
- Theories of Knowledge.
- Neural Networks.

Fields of Application

- Expert Systems.
- Language Understanding.
- Robotics.
- Automated Negotiation.
- Internet Retrieval.
- Educational Tools.
- Software Assistants.

Give examples of Expert Systems

- Xcon
- System Pressure
- Air Pressure
- Differential Equations

Methodologies

- Algorithmic (e.g. vision …).
- Heuristic (games, expert systems).
- Linguistic, Semantics (speech, language understanding).
- Symbolic Manipulation (most subjects).
- Logical Systems (formal)
- Game Theoretic.

Methodologies cont.

- Truth Maintenance Systems.
- Fuzzy Logic (knowledge representation, expert systems).
- "Knowledge Engineering" (expert systems).
- Neural Networks (learning-non-symbolic representations).
- Baysean Analysis + related (uncertainty processing)
- Learning Systems and learning theory.

Types of programs in A.I.

- Theoretically all programming languages and computers are equivalent.
- Practically there are huge differences in efficiency, even possibility NP-complete problems.
- Languages:
  - LISP – very flexible.
  - PROLOG – designed to fit "back tacking", "resolution", expert systems.
- Methodology :
  - Heuristic vs. Algorithmic.

- Algorithm :

  A recipe (set of instructions) which when followed always gives the correct solution.

- Feasible Algorithm :

  An algorithm which can in fact be implemented in such a way that the solution can be found in a reasonable time.

- Heuristic :

  a set of instructions which one has reason to believe will often give reasonably correct answers.

Heuristic vs. Algorithmic

Heuristic :

- Rules of the thumb.
- No guarantee.

Algorithmic :

- Guarantee correct results.

Why use heuristic ?

- Many problems either :
  - Can be proven to have no algorithm :
    - Theorem proving.
    - Halting problem.
  - Can be proven to have no feasible algorithm :
    - NP-complete.
    - Traveling salesman.
    - Scheduler.
    - Packing.
  - No algorithm is known although one exists :
    - Chess.

Examples of heuristics functions

- Chess :
  - No. of my pieces – No. of opponents.
  - Weighted values of pieces.
  - Positional ideas.
- Traveling Salesman :
  - Comparison with neighbors.

Algorithmic Aspects

- Undecidable Problems.
- Infeasible Problems.
- People versus NP-complete.

Other Techniques

- General Search and Matching Algorithms
- Representations of Knowledge
  - (See book by E.Davis : Naïve Physics,
  - Conceptual Dependencies (Schank),
  - Object Oriented
  - Models of Memory: Kanerva, Anderson, Grossberg
- Logic (See LICS conferences)
  - Automatic Theorem Proving
  - Non-Traditional Logics
    - Non-monotonic Logics
    - Circumscription
    - Closed World Assumption
    - Prolog

Other Techniques (cont)

- Dealing with Time
  - Logics: Temporal and Modal
  - Noise in Neural Networks
- Uncertainty
  - Baysean Networks (Pearl) (Microsoft assistant)
  - Combination Formulas
    - Dempster-Shafer
    - Fuzzy Logics
    - Hummel-Landy-Manevitz
    - Mycin, etc
  - Independence Assumptions and Weakenings
- Natural Language
- Speech

Philosophy

- Is there any possibility of AI?(Searle, Dreyfuss, Symbol Manipulation Assumption, Godel's Theorem, Consciousness).
- What would it mean to have AI?
- Turing Test: Makes Sense or Not?

Turing Test

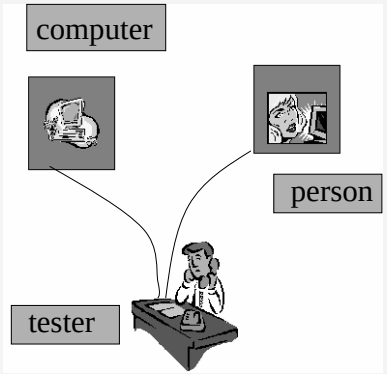

Searle's Chinese Box