دكاء اصطناعي

Artificial Intelligence

Recommended Texts

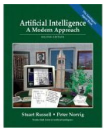

Artificial Intelligence :
A Modern Approach
Second Edition by Stuart J. Russell & Peter Norvig
(2003, Prentice Hall)

Artificial Intelligence :
Structures and Strategies for Complex Problem Solving
Fifth Edition by George F. Luger
(2005, Addison Wesley)

Syllabus

Heuristic Search, Problem solving Game Playing
Knowledge Representation, Logical Inference, Planning
Reasoning under uncertainty
Expert Systems
Learning, Perception
Language understanding

AI

G. F. Luger & W. F. Stubblefield (1993),
G. F. Luger (2005)

AI is the branch of computer science concerned with the automation of intelligent behavior.

The sources of studying AI

- Philosophy (ontology, epistemology, ...)
- Mathematics (logic, geometry, probability, decision theory, ...)
- Psychology
- Linguistics, psycholinguistics
- Computing (theory; engineering practice)

AI Languages

Programming languages best suited to AI tasks are Lisp (1960) and Prolog (1972).

There also have been specialized knowledge representation systems and languages, used to develop knowledge bases and knowledge-based systems. This includes expert systems, in which probability and beliefs play an important role.

Human vs Machine Intelligence

- Intelligence is not a unique and unshared capability of human. It is more an open collection of attributes than it is a single well-defined entity
- Humans embody many aspects of intelligence while animals typically embody a smaller number of intelligent characteristics, and usually at a much lower level
- The advent of digital computers made possible credible attempts to fulfill the AI dreams
- Computer based intelligence must be specialized to very restricted domains to be comparable to human performance

Human vs Machine Intelligence

- The two hemispheres of the human brain deal with problems in two distinct paradigms:
  - sequential (or logical) approach that considers only a small portion of the available data at a time
  - parallel processing looks at data on a global basis

Many tasks which we might reasonably think require **intelligence** are performed by computers **without even thinking** → Complex Arithmetic

Other tasks that people do without thinking are extremely difficult to automate → Recognizing a Face

Disciplines which form the core of AI inside
Fields which draw from these discipline outside.

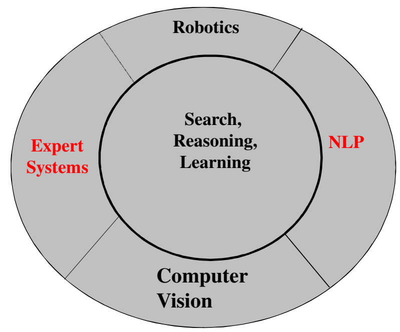

Allied Disciplines

Philosophy	Knowledge Rep., Logic, Foundation of AI
Math	Search, Analysis of search algorithms, logic
Economics	Expert Systems, Decision Theory, Principles of Rational Behavior
Psychology	Behavioristic insights into AI programs
Brain Science	Learning, Neural Nets
Physics	Learning, Information Theory & AI, Entropy, Robotics
Computer Sc. & Eng.	Systems for AI

AI Definitions

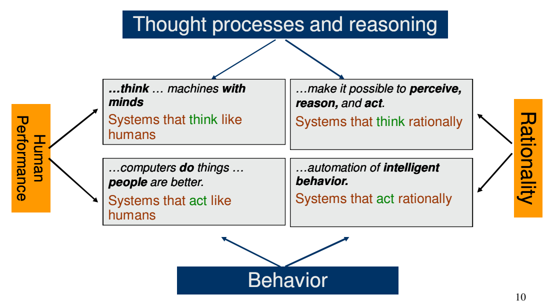

AI Definitions

For Defining AI, Books go in four directions:

- Think Like a human
- Act Like a human
- Think rationally
- Act rationally

Thinking humanly

- 1960s "cognitive revolution"
- How the computer performs tasks does matter
- The reasoning steps are important
- Requires scientific theories of the activities of the brain
- Need ability to manipulate symbols

Cognitive Science

Computer models from AI + Experimental techniques from psychology
→ Constructing working theories of human mind

- Approaches
- 1) Predicting and testing behavior of human subjects (top-down)--- Cognitive Science or
  2) Direct identification from neurological data (bottom-up) cognitive neuroscience

Both approaches are now distinct from AI Approaches

Thinking humanly

The Cognitive Modeling approach

To develop a program that think like human , the way the human think should be known

Knowing the precise theory of mind (**how human think**?)

-----> expressing the theory as a computer program

Ex:

GPS (General Problem Solver) [by Newell & Simon, 1961]

Were concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problem rather that *correctly solve problems*

Thinking rationally

Aristotle: what are correct arguments processes?

- Logic and reasoning
- What is the correct argument
- Requires reasoning structures
- Approaches: Machine Learning, NN, Reasoning

Two main obstacles:

- Not all things can be formally represented in logic notation, particularly if there is any uncertainty
- It is usually the case that even small scale problems can exhaust the computational power of any computer unless heuristics are used

Acting humanly

- Stared very early : Alan Turing 1950 designed a test for intelligent behaviour, i.e.
  - ability to achieve human-level performance in all cognitive tasks
- Take a task that people normally do : playing chess, diagnosis a disease, navigating a building … etc

Then build a computer system does is automatically

To pass the Turing Test, the computer would need:

1. Natural Language Processing → Communication
2. Knowledge Representation→store info before and during interrogation
3. Automated Reasoning→ answer questions and draw new conclusions
4. Machine learning→ adapt to new circumstances

Acting rationally

- Rational behaviour: doing the right / best thing
- The right thing: maximize goal achievement
- Does not necessarily involve thinking
- Act rationally = reason logically to the conclusion and act on that conclusion

The Rational Agent Approach

- An agent is something that perceives and acts
- A rational agent is one that acts so as to achieve the best outcome

AI Languages

Programming languages best suited to AI tasks are Lisp (1960) and Prolog (1972).

- PROgramming in LOGic
- Emphasis on *what* rather than *how*

There also have been specialized knowledge representation systems and languages, used to develop knowledge bases and knowledge-based systems. This includes expert systems, in which probability and beliefs play an important role.

The role of search

- There is a hypothesis: intelligent action can be reduced to search.
- While intelligence is certainly more than search, the hypothesis is attractive: search is well understood, easily mechanized, manageable, and so on.
- Formally represented knowledge can also be easily used in search.
- Search means systematic traversal of a space of possible solutions of a problem.
- A search space is usually a graph (often simple as a tree).
  - A node represents a partial solution.
  - An edge represents a step in the construction of a solution.

The role of search

The purpose of search may be:

- to find a path in the graph from a start node to a goal node (that is, from an initial to a final situation),
- to find a goal node.

Search Ex1: Planning

- (a) which block to *pick*, (b) which to *stack*, (c) which to *unstack*, (d) whether to *stack* a block or (e) whether to *unstack* an already stacked block. These options have to be searched in order to arrive at the right sequence of actions.

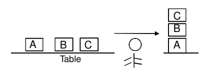

Search Ex2: Vision

- A search needs to be carried out to find which point in the image of *L* corresponds to which point in *R*. Naively carried out, this can become an *O(n2)* process where *n* is the number of points in the image.

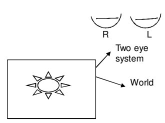

Search Ex3: Robot Path Planning

- searching amongst the options of moving ***L***eft, ***R***ight, ***U***p or ***D***own. Additionally, each movement has an associated cost representing the relative difficulty of each movement. The search then will have to find the *optimal*, *i.e.*, the *least cost* path.

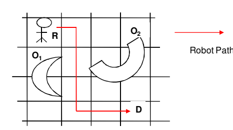

Search Ex4: Natural Language Processing

- search among many combinations of parts of speech on the way to deciphering the meaning. This applies to every level of processing- *syntax, semantics, pragmatics* and *discourse*.

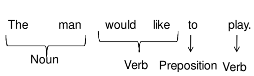

Search Ex6: Game Playing

Path to goal isn't quite right.

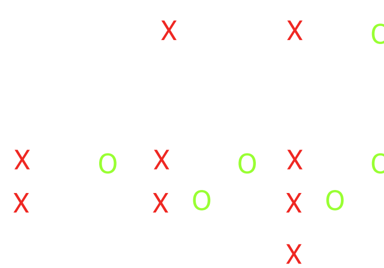

8-puzzle

1		2
4	5	3
7	8	6

Initial state

1	2	3
4	5	6
7	8	

goal state

8-puzzle

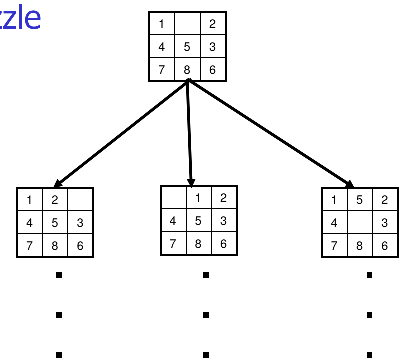

want to visit all vertices on the graph

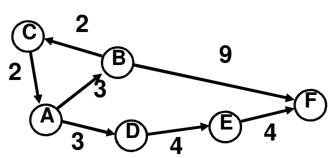

Full search tree

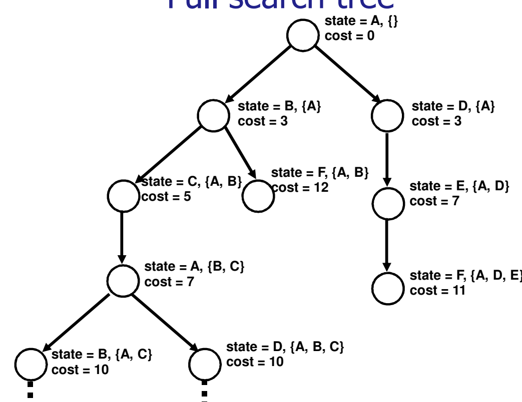

Uninformed search

Given a state, we only know whether it is a goal state or not

Cannot say one nongoal state looks better than another nongoal state

Can only traverse state space blindly in hope of somehow hitting a goal state at some point

Also called blind search

Breadth-first search

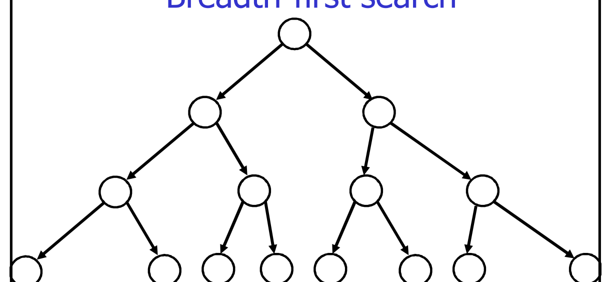

Properties of breadth-first search

Nodes are expanded in the same order in which they are generated

- Fringe (set of nodes generated but not expanded) can be maintained as a First-In-First-Out (FIFO) queue

BFS is complete: if a solution exists, one will be found

BFS finds a shortest solution

- Not necessarily an optimal solution

If every node has b successors (the branching factor), first solution is at depth d, then fringe size will be at least b^d at some point

- This much space (and time) required

Depth-first search

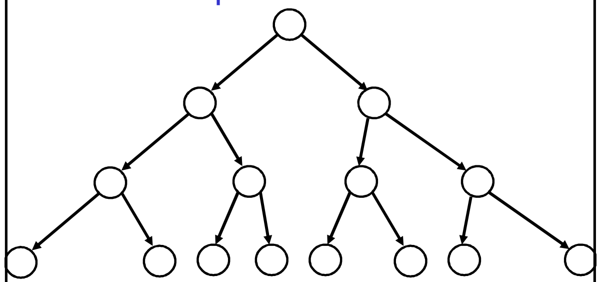

Properties of depth-first search

Fringe can be maintained as a Last-In-First-Out (LIFO) queue

Not complete (might cycle through nongoal states)

If solution found, generally not optimal/shortest

If every node has b successors (the branching factor), and we search to at most depth m, fringe is at most bm

- Much better space requirement
- Actually, generally don't even need to store all of fringe

Time: still need to look at every node

- $b^m + b^{m-1} + ... + 1$ (for b>1, $O(b^m)$)

A list of AI Search Algorithms

Systematic Search algorithms

- BFS, DFS,...
- A*
  - AO*
  - IDA* (Iterative Deepening)

Local Search Algorithms

- Minimax Search on Game Trees
- Viterbi Search on Probabilistic FSA
- Hill Climbing
- Simulated Annealing
- Gradient Descent
- Stack Based Search
- Genetic Algorithms
- Memetic Algorithms

Heuristic Search

Heuristics (Greek *heuriskein* = find, discover): "the study of the methods and rules of discovery and invention".

We use our knowledge of the problem to consider some (not all) successors of the current state (preferably just one). This means pruning the state space, gaining speed, but perhaps missing the solution!

In chess: consider one (apparently best) move, maybe a few -- but not all possible legal moves.

In the traveling salesman problem: select one nearest city, give up complete search (the greedy technique). This gives us, in polynomial time, an approximate solution of the inherently exponential problem; it can be proven that the approximation error is bounded.

Heuristic Search

For heuristic search to work, we must be able to rank the children of a node.

A heuristic function takes a state and returns a numeric value -- a composite assessment of this state. We then choose a child with the best score (this could be a maximum or minimum).

A heuristic function can help gain or lose a lot, but finding the right function is not always easy.

- The 8-puzzle: how many misplaced tiles? how many slots away from the correct place? and so on. - Chess: no simple counting of pieces is adequate.

Heuristic Search

The principal gain of using heuristic functions is the reduction of the state space.

For example:

- the full tree for Tic-Tac-Toe has 9! (362880) leaves.
- If we consider symmetries, the tree becomes six times smaller(60480), but it is still quite large.
- With a fairly simple heuristic function we can get the tree down to 40 states.
- Heuristics can also help speed up exhaustive, blind search, such as depth-first and breadth-first search.

Best-first search

The algorithm

select a heuristic function (e.g., distance to the goal);
put the initial node(s) on the open list;
repeat
select N, the best node on the open list;
succeed if N is a goal node; otherwise put N on the closed list and add N's children to the open list;
until
we succeed or the open list becomes empty (we fail);
Remark:
A closed node reached on a different path is made open.

NOTE: "the best" only means "currently appearing the best"...

Read Examples in books