Definition of AI

Acting humanly: The Turing Test approach
Thinking humanly: The cognitive modeling approach
Thinking rationally: The “laws of thought” approach(Logic)
Acting rationally: The rational agent approach

Turing test

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.
total Turing test
computer vision to perceive objects;
robotics to manipulate objects and move about.

Agents and Environments

An **agent** is anything that can be viewed as perceiving its **environment** through **sensors** and acting upon that environment through **actuators**. We use the term **percept** to refer to the agent’s perceptual inputs at any given instant. An agent’s **percept sequence** is the complete history of everything the agent has ever perceive.

PEAS

Performance measure, Environment, Actuators, Sensors

Knowledge Base

Conceptually, a set of sentences;
.Defined by TELL/ASK interface

Declarative Approach

Program an agent by TELLing it things
Equivalently, construct and install a KB
Advantages
1.Flexibility: knowledge independent of how it would be used
2.Transparency to humans
3.Agent behavior malleable via language

Knowledge Representation

Knowledge representation language: notation for expressing a KB
Consist of
Syntax: defines the legal sentences
Semantics: facts in the world to which sentences correspond, an interpretation for symbols in the logic
Logic: KR language with well-defined syntax and semantics
Model: Specifies truth or falsity of sentences

Properties of Sentences

Valid: True in all models
Satisfiable: True in some model

Rational agent

For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

Autonomy

A system is autonomous that its behavior is determined by its own percepts, rather than the prior knowledge of its designer.

Coherence

Actions must be consistent with beliefs and desires. Necessary for knowledge-level understanding of computer programs.

Agent Designs

Table Lookup, Simple Reflex, Model-based, Goal-based, Utility-based, Learning

Environment Properties

fully/partially observable, single/multi, deterministic/stochastic, episodic/sequential, static/dynamic, discrete/continuous, known/unknown

Admissibility and Consistency

Admissibility: $h(n) \leq c(n)$, where $c(n)$ is the true cost of a solution along the current path through n.
Consistency: $h(n) \leq c(n,a,n')+h(n')$.
Optimality: If h is admissible and consistent, A* is optimal.
A* is optimal, complete and optimally efficient among all algorithms that extend search paths from root.

Entailment and Inference

Entailment: $\alpha \models \beta$ if and only if $M(\alpha) \subseteq M(\beta)$. Truth of β is contained in α. $\alpha \models \beta$ if and only if the sentence $(\alpha \Rightarrow \beta)$ is valid. $\alpha \equiv \beta$ if and only if $\alpha \models \beta$ and $\beta \models \alpha$.
Inference: $\alpha \vdash \beta$ means β can be derived from α.
An inference algorithm that derives only entailed sentences is called **sound**.
An inference algorithm is **complete** if it can derive any sentence that is entailed

FOL to CNF

1.Translate bidirectionals to implications
2.Translate implications to disjunctions
3.Move negations inward (Use De Morgan's laws until only atoms are negated)
4.Standardize variables
5.Eliminate existentials via skolemization
6.Drop universal quantifiers
7.Distribute and associate into CNF

Uninformed search strategies

Breadth-first, Uniform-cost, Depth-first, Depth-limited, Iterative-deepening, Bidirectional, Graph search

Analysis

Criterion	BFS	UC
Complete	Yes	Yes
Optimal	Yes	Yes
Time	$O(b^{d+1})$	$O(b^{1+floor(C^*/e)})$
Space	$O(b^{d+1})$	$O(b^{1+floor(C^*/e)})$

Analysis

Criterion	DFS	DLDFS	IDDFS
Complete	No	No	Yes
Optimal	No	No	Yes
Time	$O(b^m)$	$O(b^L)$	$O(b^d)$
Space	O(bm)	O(bL)	O(bd)

Informed Search

Priority-first(PFS), Greedy, A*, Iterative-deepening A*, heuristics and admissibility/consistency

Flavors of PFS

Number of edges from origin: BFS
g(n), the path cost from the initial state(node) to node n: UCS(Dijkstra)
h(n), the estimated path cost from node n to goal: Greedy Search
f(n)=g(n)+h(n), the estimated cost of a path passing through n: A* search

Memory-bounded heuristic search

iterative-deepening A* (IDA*)
Recursive best-first search (RBFS): best first search with linear space

Local search strategies

Hill-climbing/stochastic, beam, simulated annealing, genetic algorithms

Local search summary

Key advantages
1.Very little memory
2.Can often find reasonable solutions in large or infinite state spaces where systematic approaches are unsuitable
3.Better answers the more time spent
Disadvantages
Usually incomplete and not optimal

Constraint propagation

Definition: Propagating the implications of a constraint on one variable onto other variables.
k-consistent: for any consistent assignment of k-1 variables, exists consistent value of any kth(Node/Arc/Path)
Strongly k-consistent: j-consistent for all j less than or equal to k

Constrain Satisfaction Search

Backtracking
Minimum remaining values heuristic: choose variable with fewest legal values remaining
Degree heuristic: choose variable with largest number of constraints
Forward checking: delete inconsistent values ahead
Backjumping
Basic: backtrack to most recent decision
Conflict-Directed: backtrack to most recent variable in conflict with