

Agents and Rationality

An agent is something that acts
operate autonomously, perceive environment,
persist and adapt

Rational agent acts to achieve the
best (expected) outcome

Agents perceive environment through
sensors and acts upon through actuators

Agent function is abstract mathematical
description, agent program is concrete
implementation

Performance measure determines what
behavior is rational, should design with
desired environment in mind

Omniscience is knowing the actual outcome
Agent should learn, autonomy measures
reliance on designer prior knowledge

Environments

Task environment consists of performance
measure, environment, actuators and sensors

Fully/partially observable - agent's sensors
give it full/partial knowledge of environment

Single/multiagent - other agent B may try to
maximize own performance measure based on A,
see if you want to treat B as agent or object

Deterministic/stochastic - if next state is
completely determined by current state and
agent action, stochastic \rightarrow probabilities,
partially observable may appear to be stochastic,
nondeterministic \rightarrow no probabilities

Episode/sequential - current decision does not
affect later decision (part checking vs. chess/drawing)

Static/dynamic - if environment changes while
agent deliberates \rightarrow dynamic (taxi), doesn't change
but performance measure does \rightarrow semi-dynamic (class),
otherwise static (crossword)

Discrete/continuous - time handling (chess \rightarrow discrete),
taxi/input from digital source \rightarrow continuous

Known/unknown - known if outcome/probabilities all known,
unknown \rightarrow do not know how environment responds

Types of Agents

Simple Reflex Agents

Base actions off of current percept, no history
Use percept in lookup table, assumes fully obs.

Infinite loops may happen, can randomize actions

Model-based Reflex Agents

Maintain internal state based on percept history
Use model of the world, update state based on percept/
Creates a "best-guess" of world state action

Goal-based agents

Make decision based on state as well as goal
May make some decisions to reach goal spatially,
don't crash if you want to move it

Goal-based more flexible than reflex-based

Utility-based agents

Generalization of goal-based agents for optimality
rational utility-based agent maximizes expected utility
Can make proper tradeoffs, but we must still
define

Search Problems

Generally, environment is observable, discrete,
known and deterministic

"Formulate, search, execute" - agent carries out
solution with eyes closed after searching

Problem (in parts)

① Initial state - where/when agent begins

② Actions - description of possible available actions

③ Transition Model - description of what actions
do, i.e. $Result(s, a) = s'$

Forms graph where nodes are states and links=actions

④ Goal Test - check if s matches goal state

⑤ Path cost function - assign numeric cost to path
step-cost $\Rightarrow C(s, a, s')$

Optimal solutions entail lowest path cost

Tree-search + explored set \Rightarrow Graph-search

In each node n in search tree, remember:
 n . State - state in state space, n . parent - parent

n . Action - action applied from parent

n . Path-Cost - path cost so far by parent pointers

Use different types of queues for fringe

Completeness - will the algorithm find a solution if
there is one?
write complexity in terms of branching factor,
depth d and m , max length of path in state space

Uniformed Search Strategies

Breadth-first search (queue as fringe)

Goal Test applied when node generated, vs. expanded
Complete and will find shallowest node d , not
always optimal

Time + space complexity - $O(b^d)$,
dominated by size of the frontier

Uniform-cost search (priority queue as fringe)

Expands the node with lowest path cost $g(n)$

When node n expanded, optimal path to $n = \text{known}$

Complete if cost of every step exceeds small, fixed ϵ

Time/space complexity = $O(b^{Ll + C^*/\epsilon})$

Depth-first search (stack as fringe)

Graph search in finite state space complete,
tree search is not (infinite loops)

Definitely not optimal by search nature

Time Comp - $O(b^m)$, all nodes

Space - $O(b^m) = \text{good for memory}$

Agent Program Representations

Atomic "Black Box"

Each state of the world is indivisible
and has no internal structure.

Factored Representation

Split state into set of variables and
attributes, each has a value

Can use booleans and continuous

Able to represent uncertainty (leave blank)

Structural Representation

Model how different parts of the
world are related to each other.

Learning Agents

Learning element - makes improvements

Performance element - makes decisions, fed by learning
element via critic

Learning elem uses feedback from critic

Problem generator - suggests to performance element
things that will lead to informative experiences

Depth-Limited Search

Dfs to predetermined depth limit l

Possible incompleteness if choose $l < d$
Pretends that no children after l

Iterative deepening DFS $O(b^d)$ time, $O(b^l)$ space

Gradually increases depth-limit, expands (first)

Iterative lengthening search

Increasing path-cost vs. depth-limits

Bidirectional Search

Run BFS/UCS from start and goal,
so $b^{d/2} + b^{d/2} < b^d$

Must be able to compute predecessors

Informed (Heuristic) Search

Make use of a combination of path cost $g(n)$
and estimated remaining cost $h(n)$

Greedy best-first search

$f(n) = h(n)$, tries to expand node closest to goal

Tree-search infinite loop, graph-search is complete

A* Search

$f(n) = g(n) + h(n)$

Complete + optimal, if $h(n)$ is admissible
tree search is optimal, if consistent, graph
search is optimal

Consistent if: $h(n) \leq C(n, a, n') + h(n')$

Whenever node expanded, optimal path found
because (assuming heuristic is admissible),
if n' on way to n , then would've been
expanded first

Expands all nodes $f(n) < C^*$, may

expand nodes on goal contour

Optimally efficient \exists no other optimal

algorithm guaranteed to expand fewer nodes

Iterative Deepening A*

round is f -cost, better than memory

Recursive best-first search

Substitutes to maintain state, does not store

Heuristic Functions

Effective branching factor

$N+1 = 1 + b^* + (b^*)^2 + (b^*)^3 \dots$

h_2 dominates $h_1 \Rightarrow A^*$ never expands

more nodes using h_2 than h_1

Relaxed problems can be used to generate

admissible consistent heuristics

Learning + linear combination

$h(n) = c_1 x_1(n) + c_2 x_2(n)$,

not always admissible + consistent

Local Search

Evaluate and modify current state, path cost is not relevant

- ① Use small amounts of memory
 - ② can generally find reasonable solutions in large/infinite spaces
- Complete - always finds a solution
optimal - always finds global min/max

Hill-climbing Search

Steepest-ascend/greedy local search
Complete - state formulation - all variables are present at one time

- Ⓐ Local maxima - gets stuck by steepest ascent
- Ⓑ ridges - sequence of local maxima
- Ⓒ Plateau - flat area or shoulder

sideways move to escape from shoulder
Stochastic hill-climbing - random from uphill moves w/ weighted probability

First-choice hill-climbing - generate random successors until one is better
random-restart hill-climbing - conduct search from random states, sometimes restart

Simulated-annealing

Refine state by picking random move and if better, take it, else accept with probability α , which \downarrow
Shake out of local min/max

Local-beam Search

Start with k random states
Generate all successors of each state and only keep top k
Parallel search threads, pass info
Stochastic beam-search - choose k successors at random α how "good"

Genetic algorithm

Similar to stochastic beam search
Successors generated by combining two parent states
Objective function as fitness function
crossover point and mutation

Adversarial Search

Definition of game as search problem

S_0 : Initial state, how game set up at start

Player(s): which player has move in white.

Actions: Returns set of legal moves

Result(s, a): Transition model - result of move

Terminal test(s): check if game over

Utility(s, p): objective/payoff function

Zen-sum \Rightarrow total payoff same in all games

Initial state, Actions, Result define game tree

Minimax

Alternation of maximizing (you) and minimizing (opponent) plies, accomplished by DFS to terminal states

Multiplayer games may use a vector (a, b, c) for combined values and may have different plies

Alpha-beta Pruning

As we search, keep values α (best max value) and β (best min value), cut off exploration if value worse than β or

May use transposition table of known states to cut down

Cutoff search by evaluation to make imperfect roll-time

Can modify α - β to use eval beyond a certain ^{decision} point

Good for quiescent positions, horizon effect is trickier

Searching & Nondeterministic Actions

Percepts help in partially observable or nondeterministic environment

contingency plan (strategy) - for all possible sequence of percepts

Solutions look like trees (if-then-else) rather than sequences

And-or search trees alternate OR (you) and AND (environment), solution must

be able to handle at least one OR options and all AND options

Some nondeterministic environments may require cyclical solutions

Searching & Partial Observation

Belief state is the set of all possible physical states that an agent could be in.

No Observation

Belief-states: up to 2^N unique sets of belief states \Rightarrow belief-state space, set of possible physical

Initial state: set of all states in P initially

Actions: If illegal actions have no effect,

$\bigcup_{s \in b} \text{Actions}(s)$, else $\bigcap_{s \in b} \text{Actions}(s) \Rightarrow$ safer

Transition Model: deterministic \Rightarrow

$b' = \text{Result}(b, a) = \{s' : s' = \text{Result}(s, a) \text{ and } s \in b\}$

else $\bigcup_{s \in b} \text{Results}(s, a)$

Goal Test: check that all belief states satisfy

Path Cost: May be different between states

with observation

prediction: $\hat{b} = \text{Predict}(b, a)$

Observation prediction: determine set of percepts that could be observed

Update: Partitions and returns set of belief states that correspond to intake

percept $o \Rightarrow b = \text{update}(\hat{b}, o)$

Nondeterminism will just enlarge the number of belief states

Stochastic Games

Introduce chance nodes, value equivalent to expected value of outcomes

minimax \Rightarrow expectiminimax

Alternating a-O-b-O... pattern

Constraint Satisfaction Problems

Utilize a factored representation where variable assignment leads to simultaneous satisfaction of all constraints

X - set of variables $\{x_1, \dots, x_n\}$, D - set of domains, one for each x_i

C - set of constraints \Rightarrow allowable combinations of values

Constraint graphs to visualize systems of binary constraints

Global constraint - arbitrary # of variables

node consistency \equiv entire unary, arc - entire binary
path consistency \equiv triples, $\{x_i, x_j\} \wedge x_k$ if $\{x_i, x_j\}, \{x_i, x_k\}, \{x_j, x_k\}$

Solutions are commutative
Backtracking Search (DFS one var at a time, assign, backtrack if there is no domain), flexible

Heuristics: MRV - minimum remaining value

Degree - prioritize those w/ higher degree

LCV - least constraining value

Forward checking - choose variable, establish core consistency from neighbors and continue, backtrack if no domain

MFC (Maintaining the consistency) - call AC3 instead of just assigning from neighbors

chronological backtracking + backjumping

min-conflicts (pick random var and minimize its conflicts, repeat)

cutset conditioning

choose subset of graph to form a tree

Loop through cutset and run tree CSP solver

Enforces bounded consistency then runs forward checking

Logical Agents

Knowledge base \equiv set of sentences in knowledge representation language that describe the world

TELL and Ask operations, both may involve inference, denoting new sentence from old

Logic

entailment implies β follows logically from α

logical inference / model checking

sound - can derive entailed sentences,

complete - can derive any sentence that is entailed

Propositional Logic

$\neg, \wedge, \vee, \Rightarrow, \Leftrightarrow$

truth tables to express truthfulness of sentences

Model checking, Theorem-proving

Valid - true in all models, satisfiable - true in some model

True β found by unsatisfiability of $\alpha \wedge \neg \beta$, proof by contradiction

Initial state: KB; Actions: All inference rules applied

result: add sentence in inference rule

back half goal: state we try to prove

Conjunctive normal form

DNF - pure symbol, unit clause \rightarrow early termination