PROBLEM SOLVING AND SEARCH

Chapter 3

Outline

- \Diamond Problem-solving agents
- \diamondsuit Problem types
- \diamondsuit Problem formulation
- \diamond Example problems
- \diamond Basic search algorithms

Problem-solving agents

Restricted form of general agent:

```
function SIMPLE-PROBLEM-SOLVING-AGENT( percept) returns an action

static: seq, an action sequence, initially empty

state, some description of the current world state

goal, a goal, initially null

problem, a problem formulation

state \leftarrow UPDATE-STATE(state, percept)

if seq is empty then

goal \leftarrow FORMULATE-GOAL(state)

problem \leftarrow FORMULATE-PROBLEM(state, goal)

seq \leftarrow SEARCH( problem)

action \leftarrow RECOMMENDATION(seq, state)

seq \leftarrow REMAINDER(seq, state)

return action
```

Note: this is offline problem solving; solution executed "eyes closed." Online problem solving involves acting without complete knowledge.

Example: Romania

On holiday in Romania; currently in Arad. Flight leaves tomorrow from Bucharest

Formulate goal:

be in Bucharest

Formulate problem:

states: various cities actions: drive between cities

Find solution:

sequence of cities, e.g., Arad, Sibiu, Fagaras, Bucharest



Single-state problem formulation

A problem is defined by four items:

initial state e.g., "at Arad"

successor function S(x) = set of action-state pairse.g., $S(Arad) = \{\langle Arad \rightarrow Zerind, Zerind \rangle, \ldots \}$

goal test, can be explicit, e.g., x = "at Bucharest" implicit, e.g., NoDirt(x)

```
path cost (additive)
e.g., sum of distances, number of actions executed, etc.
c(x, a, y) is the step cost, assumed to be \geq 0
```

A solution is a sequence of actions leading from the initial state to a goal state

Example: The 8-puzzle





Start State

Goal State

<u>states</u>?? <u>actions</u>?? <u>goal test</u>?? <u>path cost</u>??



states??: integer locations of tiles (ignore intermediate positions)
actions??
goal test??
path cost??



states??: integer locations of tiles (ignore intermediate positions)
actions??: move blank left, right, up, down (ignore unjamming etc.)
goal test??
path cost??



states??: integer locations of tiles (ignore intermediate positions)
actions??: move blank left, right, up, down (ignore unjamming etc.)
goal test??: = goal state (given)
path cost??



states??: integer locations of tiles (ignore intermediate positions)
actions??: move blank left, right, up, down (ignore unjamming etc.)
goal test??: = goal state (given)
path cost??: 1 per move

[Note: optimal solution of *n*-Puzzle family is NP-hard]

Tree search algorithms

Basic idea:

offline, simulated exploration of state space by generating successors of already-explored states (a.k.a. expanding states)

function TREE-SEARCH(problem, strategy) returns a solution, or failure
initialize the search tree using the initial state of problem
loop do
 if there are no candidates for expansion then return failure
 choose a leaf node for expansion according to strategy
 if the node contains a goal state then return the corresponding solution
 else expand the node and add the resulting nodes to the search tree
end

Tree search example



Tree search example



Tree search example



Implementation: states vs. nodes



The Expand function creates new nodes, filling in the various fields and using the SUCCESSORFN of the problem to create the corresponding states.

Implementation: general tree search

```
function TREE-SEARCH (problem, fringe) returns a solution, or failure
   fringe \leftarrow \text{INSERT}(\text{MAKE-NODE}(\text{INITIAL-STATE}[problem]), fringe)
   loop do
        if fringe is empty then return failure
        node \leftarrow \text{REMOVE-FRONT}(fringe)
        if GOAL-TEST(problem, STATE(node)) then return node
        fringe \leftarrow \text{INSERTALL}(\text{EXPAND}(node, problem), fringe)
function EXPAND(node, problem) returns a set of nodes
   successors \leftarrow \text{the empty set}
   for each action, result in SUCCESSOR-FN(problem, STATE[node]) do
        s \leftarrow a \text{ new NODE}
        PARENT-NODE[s] \leftarrow node; ACTION[s] \leftarrow action; STATE[s] \leftarrow result
        PATH-COST[s] \leftarrow PATH-COST[node] + STEP-COST(node, action, s)
        \text{DEPTH}[s] \leftarrow \text{DEPTH}[node] + 1
        add s to successors
   return successors
```

Search strategies

A strategy is defined by picking the order of node expansion

Strategies are evaluated along the following dimensions: completeness—does it always find a solution if one exists? time complexity—number of nodes generated/expanded space complexity—maximum number of nodes in memory optimality—does it always find a least-cost solution?

Time and space complexity are measured in terms of *b*—maximum branching factor of the search tree *d*—depth of the least-cost solution *m*—maximum depth of the state space (may be ∞)

Uninformed search strategies

Uninformed strategies use only the information available in the problem definition

Breadth-first search

Uniform-cost search

Depth-first search

Depth-limited search

Iterative deepening search

Expand shallowest unexpanded node



Expand shallowest unexpanded node



Expand shallowest unexpanded node



Expand shallowest unexpanded node



Complete??

<u>Complete</u>?? Yes (if *b* is finite)

Time??

Complete?? Yes (if *b* is finite)

<u>Time</u>?? $1 + b + b^2 + b^3 + \ldots + b^d + b(b^d - 1) = O(b^{d+1})$, i.e., exp. in d Space??

Complete?? Yes (if *b* is finite)

<u>Time</u>?? $1 + b + b^2 + b^3 + \ldots + b^d + b(b^d - 1) = O(b^{d+1})$, i.e., exp. in d

<u>Space</u>?? $O(b^{d+1})$ (keeps every node in memory)

Optimal??

Complete?? Yes (if *b* is finite)

<u>Time</u>?? $1 + b + b^2 + b^3 + \ldots + b^d + b(b^d - 1) = O(b^{d+1})$, i.e., exp. in d

<u>Space</u>?? $O(b^{d+1})$ (keeps every node in memory)

Optimal?? Yes (if cost = 1 per step); not optimal in general

Space is the big problem; can easily generate nodes at 100MB/sec so 24hrs = 8640GB.

Uniform-cost search

Expand least-cost unexpanded node

Implementation:

fringe = queue ordered by path cost, lowest first

Equivalent to breadth-first if step costs all equal

<u>Complete</u>?? Yes, if step cost $\geq \epsilon$

<u>Time</u>?? # of nodes with $g \leq \text{ cost of optimal solution}$, $O(b^{\lceil C^*/\epsilon \rceil})$ where C^* is the cost of the optimal solution

<u>Space</u>?? # of nodes with $g \leq \text{ cost of optimal solution, } O(b^{\lceil C^*/\epsilon \rceil})$

Optimal?? Yes—nodes expanded in increasing order of g(n)

Expand deepest unexpanded node



Expand deepest unexpanded node

Implementation:

Expand deepest unexpanded node


Expand deepest unexpanded node

Implementation:



Expand deepest unexpanded node

Implementation:



Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front



Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front



Expand deepest unexpanded node

Implementation:

fringe = LIFO queue, i.e., put successors at front



Complete??

<u>Complete</u>?? No: fails in infinite-depth spaces, spaces with loops Modify to avoid repeated states along path ⇒ complete in finite spaces

Time??

<u>Complete</u>?? No: fails in infinite-depth spaces, spaces with loops Modify to avoid repeated states along path ⇒ complete in finite spaces

<u>Time</u>?? $O(b^m)$: terrible if m is much larger than dbut if solutions are dense, may be much faster than breadth-first

Space??

<u>Complete</u>?? No: fails in infinite-depth spaces, spaces with loops Modify to avoid repeated states along path ⇒ complete in finite spaces

<u>Time</u>?? $O(b^m)$: terrible if m is much larger than dbut if solutions are dense, may be much faster than breadth-first

Space?? O(bm), i.e., linear space!

Optimal??

<u>Complete</u>?? No: fails in infinite-depth spaces, spaces with loops Modify to avoid repeated states along path ⇒ complete in finite spaces

<u>Time</u>?? $O(b^m)$: terrible if m is much larger than dbut if solutions are dense, may be much faster than breadth-first

Space?? O(bm), i.e., linear space!

Optimal?? No

Depth-limited search

= depth-first search with depth limit l,

i.e., nodes at depth l have no successors

Recursive implementation:

```
function DEPTH-LIMITED-SEARCH( problem, limit) returns soln/fail/cutoff

RECURSIVE-DLS(MAKE-NODE(INITIAL-STATE[problem]), problem, limit)

function RECURSIVE-DLS(node, problem, limit) returns soln/fail/cutoff

cutoff-occurred? \leftarrow false

if GOAL-TEST(problem, STATE[node]) then return node

else if DEPTH[node] = limit then return cutoff

else for each successor in EXPAND(node, problem) do

result \leftarrow RECURSIVE-DLS(successor, problem, limit)

if result = cutoff then cutoff-occurred? \leftarrow true

else if result \neq failure then return result

if cutoff-occurred? then return cutoff else return failure
```

Iterative deepening search

```
function ITERATIVE-DEEPENING-SEARCH( problem) returns a solution
inputs: problem, a problem
for depth ← 0 to ∞ do
    result ← DEPTH-LIMITED-SEARCH( problem, depth)
    if result ≠ cutoff then return result
end
```

Iterative deepening search l = 0











Complete??

Complete?? Yes

Time??

Complete?? Yes

<u>Time</u>?? $(d+1)b^0 + db^1 + (d-1)b^2 + \ldots + b^d = O(b^d)$

Space??

Complete?? Yes

<u>Time</u>?? $(d+1)b^0 + db^1 + (d-1)b^2 + \ldots + b^d = O(b^d)$

Space?? O(bd)

Optimal??

Complete?? Yes

<u>Time</u>?? $(d+1)b^0 + db^1 + (d-1)b^2 + \ldots + b^d = O(b^d)$

Space?? O(bd)

Optimal?? Yes, if step cost = 1

Can be modified to explore uniform-cost tree

Numerical comparison for b = 10 and d = 5, solution at far right leaf:

 $N(\mathsf{IDS}) = 50 + 400 + 3,000 + 20,000 + 100,000 = 123,450$ $N(\mathsf{BFS}) = 10 + 100 + 1,000 + 10,000 + 100,000 + 999,990 = 1,111,100$

IDS does better because other nodes at depth d are not expanded BFS can be modified to apply goal test when a node is **generated**

Summary of algorithms

Criterion	Breadth-	Uniform-	Depth-	Depth-	Iterative
	First	Cost	First	Limited	Deepening
Complete?	Yes^*	Yes*	No	Yes, if $l \ge d$	Yes
Time	b^{d+1}	$b^{\lceil C^*/\epsilon \rceil}$	b^m	b^l	b^d
Space	b^{d+1}	$b^{\lceil C^*/\epsilon \rceil}$	bm	bl	bd
Optimal?	Yes^*	Yes	No	No	Yes^*

Repeated states

Failure to detect repeated states can turn a linear problem into an exponential one!



Graph search

```
function GRAPH-SEARCH(problem, fringe) returns a solution, or failure

closed \leftarrow an empty set

fringe \leftarrow INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)

loop do

if fringe is empty then return failure

node \leftarrow REMOVE-FRONT(fringe)

if GOAL-TEST(problem, STATE[node]) then return node

if STATE[node] is not in closed then

add STATE[node] to closed

fringe \leftarrow INSERTALL(EXPAND(node, problem), fringe)

end
```

Summary

Problem formulation usually requires abstracting away real-world details to define a state space that can feasibly be explored

Variety of uninformed search strategies

Iterative deepening search uses only linear space and not much more time than other uninformed algorithms

Graph search can be exponentially more efficient than tree search

INFORMED SEARCH ALGORITHMS

Chapter 4, Sections 1–2

Chapter 4, Sections 1–2 1

Outline

- \diamondsuit Best-first search
- $\diamondsuit \ \ \mathsf{A}^* \ \mathsf{search}$
- \diamondsuit Heuristics

Review: Tree search

```
\begin{aligned} & \textbf{function Tree-Search}(\textit{problem, fringe}) \textbf{ returns a solution, or failure} \\ & \textit{fringe} \leftarrow \text{INSERT}(\text{MAKE-NODE}(\text{INITIAL-STATE}[\textit{problem}]), \textit{fringe}) \\ & \textbf{loop do} \\ & \textbf{if fringe is empty then return failure} \\ & \textit{node} \leftarrow \text{REMOVE-FRONT}(\textit{fringe}) \\ & \textbf{if GOAL-TEST}[\textit{problem}] \textbf{ applied to STATE}(\textit{node}) \textbf{ succeeds return node} \\ & \textit{fringe} \leftarrow \text{INSERTALL}(\text{EXPAND}(\textit{node, problem}), \textit{fringe}) \end{aligned}
```

A strategy is defined by picking the order of node expansion

Best-first search

Idea: use an evaluation function for each node - estimate of "desirability"

 \Rightarrow Expand most desirable unexpanded node

Implementation:

fringe is a queue sorted in decreasing order of desirability

Special cases: greedy search

 A^* search

Romania with step costs in km



Greedy search

Evaluation function h(n) (heuristic)

= estimate of cost from n to the closest goal

E.g., $h_{SLD}(n) = \text{straight-line distance from } n$ to Bucharest

Greedy search expands the node that appears to be closest to goal









Properties of greedy search

Complete??
$\label{eq:complete} \underbrace{ \mbox{Complete} ?? \mbox{No-can get stuck in loops, e.g., with Oradea as goal, } \\ \mbox{Iasi} \rightarrow \mbox{Neamt} \rightarrow \mbox{Iasi} \rightarrow \mbox{Neamt} \rightarrow \\ \mbox{Complete in finite space with repeated-state checking} \end{cases}$

Time??

<u>Time</u>?? $O(b^m)$, but a good heuristic can give dramatic improvement

Space??

<u>Time</u>?? $O(b^m)$, but a good heuristic can give dramatic improvement

<u>Space</u>?? $O(b^m)$ —keeps all nodes in memory

Optimal??

<u>Time</u>?? $O(b^m)$, but a good heuristic can give dramatic improvement

<u>Space</u>?? $O(b^m)$ —keeps all nodes in memory

Optimal?? No

A^* search

Idea: avoid expanding paths that are already expensive

Evaluation function f(n) = g(n) + h(n)

g(n) = cost so far to reach nh(n) = estimated cost to goal from nf(n) = estimated total cost of path through n to goal

A* search uses an admissible heuristic i.e., $h(n) \leq h^*(n)$ where $h^*(n)$ is the **true** cost from n. (Also require $h(n) \geq 0$, so h(G) = 0 for any goal G.)

E.g., $h_{\mathrm{SLD}}(n)$ never overestimates the actual road distance

Theorem: A^* search is optimal





A^{*} search example









Optimality of A^{*} (standard proof)

Suppose some suboptimal goal G_2 has been generated and is in the queue. Let n be an unexpanded node on a shortest path to an optimal goal G_1 .



 $f(G_2) = g(G_2) \qquad \text{since } h(G_2) = 0$ > $g(G_1) \qquad \text{since } G_2 \text{ is suboptimal}$ $\geq f(n) \qquad \text{since } h \text{ is admissible}$

Since $f(G_2) > f(n)$, A^{*} will never select G_2 for expansion

Optimality of A^{*} (more useful)

Lemma: A^* expands nodes in order of increasing f value^{*}

Gradually adds "f-contours" of nodes (cf. breadth-first adds layers) Contour i has all nodes with $f = f_i$, where $f_i < f_{i+1}$



Complete??

<u>Complete</u>?? Yes, unless there are infinitely many nodes with $f \leq f(G)$ <u>Time</u>??

<u>Complete</u>?? Yes, unless there are infinitely many nodes with $f \leq f(G)$

<u>Time</u>?? Exponential in [relative error in $h \times$ length of soln.]

Space??

<u>Complete</u>?? Yes, unless there are infinitely many nodes with $f \leq f(G)$

<u>Time</u>?? Exponential in [relative error in $h \times$ length of soln.]

Space?? Keeps all nodes in memory

Optimal??

<u>Complete</u>?? Yes, unless there are infinitely many nodes with $f \leq f(G)$

<u>Time</u>?? Exponential in [relative error in $h \times$ length of soln.]

Space?? Keeps all nodes in memory

<u>Optimal</u>?? Yes—cannot expand f_{i+1} until f_i is finished

- A^* expands all nodes with $f(n) < C^*$
- A^* expands some nodes with $f(n)=C^*$
- A^* expands no nodes with $f(n) > C^*$

Proof of lemma: Consistency

A heuristic is consistent if

$$h(n) \le c(n, a, n') + h(n')$$

If h is consistent, we have

$$f(n') = g(n') + h(n') = g(n) + c(n, a, n') + h(n') \geq g(n) + h(n) = f(n)$$

I.e., f(n) is nondecreasing along any path.



Admissible heuristics

E.g., for the 8-puzzle:

 $h_1(n) =$ number of misplaced tiles $h_2(n) =$ total Manhattan distance (i.e., no. of squares from desired location of each tile) **Start State Goal State**

 $\frac{h_1(S) = ??}{h_2(S) = ??}$

Admissible heuristics

E.g., for the 8-puzzle:

 $h_1(n) =$ number of misplaced tiles $h_2(n) =$ total Manhattan distance (i.e., no. of squares from desired location of each tile) **Start State Goal State**

 $\frac{h_1(S) = ?? \ 6}{h_2(S) = ?? \ 4 + 0 + 3 + 3 + 1 + 0 + 2 + 1 = 14}$

Dominance

If $h_2(n) \ge h_1(n)$ for all n (both admissible) then h_2 dominates h_1 and is better for search

Typical search costs:

 $\begin{array}{ll} d = 14 & \mathsf{IDS} = \texttt{3,473,941} \ \mathsf{nodes} \\ & \mathsf{A}^*(h_1) = \texttt{539} \ \mathsf{nodes} \\ & \mathsf{A}^*(h_2) = \texttt{113} \ \mathsf{nodes} \\ d = 24 & \mathsf{IDS} \approx \texttt{54,000,000,000} \ \mathsf{nodes} \\ & \mathsf{A}^*(h_1) = \texttt{39,135} \ \mathsf{nodes} \\ & \mathsf{A}^*(h_2) = \texttt{1,641} \ \mathsf{nodes} \end{array}$

Given any admissible heuristics h_a , h_b ,

 $h(n) = \max(h_a(n), h_b(n))$

is also admissible and dominates h_a , h_b

Relaxed problems

Admissible heuristics can be derived from the **exact** solution cost of a **relaxed** version of the problem

If the rules of the 8-puzzle are relaxed so that a tile can move **anywhere**, then $h_1(n)$ gives the shortest solution

If the rules are relaxed so that a tile can move to any adjacent square, then $h_2(n)$ gives the shortest solution

Key point: the optimal solution cost of a relaxed problem is no greater than the optimal solution cost of the real problem

Relaxed problems contd.

Well-known example: travelling salesperson problem (TSP) Find the shortest tour visiting all cities exactly once



Minimum spanning tree can be computed in $O(n^2)$ and is a lower bound on the shortest (open) tour

Summary

Heuristic functions estimate costs of shortest paths

Good heuristics can dramatically reduce search cost

Greedy best-first search expands lowest h

- incomplete and not always optimal
- A^* search expands lowest g+h
 - complete and optimal
 - also optimally efficient (up to tie-breaks, for forward search)

Admissible heuristics can be derived from exact solution of relaxed problems

GAME PLAYING

CHAPTER 6

Outline

- $\diamondsuit \ \mathsf{Games}$
- \diamondsuit Perfect play
 - minimax decisions
 - $\alpha \beta$ pruning
- \diamondsuit Resource limits and approximate evaluation
- \diamondsuit Games of chance
- \diamondsuit Games of imperfect information

Games vs. search problems

"Unpredictable" opponent \Rightarrow solution is a strategy specifying a move for every possible opponent reply

Time limits \Rightarrow unlikely to find goal, must approximate

Plan of attack:

- Computer considers possible lines of play (Babbage, 1846)
- Algorithm for perfect play (Zermelo, 1912; Von Neumann, 1944)
- Finite horizon, approximate evaluation (Zuse, 1945; Wiener, 1948; Shannon, 1950)
- First chess program (Turing, 1951)
- Machine learning to improve evaluation accuracy (Samuel, 1952–57)
- Pruning to allow deeper search (McCarthy, 1956)

Types of games

perfect information

imperfect information

deterministic	chance
chess, checkers,	backgammon
go, othello	monopoly
battleships,	bridge, poker, scrabble
blind tictactoe	nuclear war



Minimax

Perfect play for deterministic, perfect-information games

Idea: choose move to position with highest minimax value = best achievable payoff against best play



Minimax algorithm

```
function MINIMAX-DECISION(state) returns an action
```

inputs: *state*, current state in game

```
return the a in ACTIONS(state) maximizing MIN-VALUE(RESULT(a, state))
```

```
function MAX-VALUE(state) returns a utility value
if TERMINAL-TEST(state) then return UTILITY(state)
v \leftarrow -\infty
for a, s in SUCCESSORS(state) do v \leftarrow MAX(v, MIN-VALUE(s))
return v
```

```
function MIN-VALUE(state) returns a utility value
if TERMINAL-TEST(state) then return UTILITY(state)
v \leftarrow \infty
for a, s in SUCCESSORS(state) do v \leftarrow MIN(v, MAX-VALUE(s))
return v
```

Complete??

<u>Complete</u>?? Only if tree is finite (chess has specific rules for this). NB a finite strategy can exist even in an infinite tree!

Optimal??

Complete?? Yes, if tree is finite (chess has specific rules for this)

Optimal?? Yes, against an optimal opponent. Otherwise??

Time complexity??

Complete?? Yes, if tree is finite (chess has specific rules for this)

Optimal?? Yes, against an optimal opponent. Otherwise??

Time complexity?? $O(b^m)$

Space complexity??
Properties of minimax

Complete?? Yes, if tree is finite (chess has specific rules for this)

Optimal?? Yes, against an optimal opponent. Otherwise??

Time complexity?? $O(b^m)$

Space complexity?? O(bm) (depth-first exploration)

For chess, $b \approx 35$, $m \approx 100$ for "reasonable" games \Rightarrow exact solution completely infeasible

But do we need to explore every path?











Why is it called $\alpha - \beta$?



 α is the best value (to MAX) found so far off the current path If V is worse than α , MAX will avoid it \Rightarrow prune that branch Define β similarly for MIN

The $\alpha - \beta$ algorithm

```
function ALPHA-BETA-DECISION(state) returns an action
   return the a in ACTIONS(state) maximizing MIN-VALUE(RESULT(a, state))
function MAX-VALUE(state, \alpha, \beta) returns a utility value
   inputs: state, current state in game
            \alpha, the value of the best alternative for MAX along the path to state
            \beta, the value of the best alternative for MIN along the path to state
   if TERMINAL-TEST(state) then return UTILITY(state)
   v \leftarrow -\infty
   for a, s in SUCCESSORS(state) do
      v \leftarrow Max(v, MIN-VALUE(s, \alpha, \beta))
      if v \geq \beta then return v
      \alpha \leftarrow MAX(\alpha, v)
   return v
```

function MIN-VALUE(*state*, α , β) **returns** *a utility value* same as MAX-VALUE but with roles of α , β reversed

Properties of $\alpha - \beta$

Pruning does not affect final result

Good move ordering improves effectiveness of pruning

With "perfect ordering," time complexity = $O(b^{m/2})$ \Rightarrow **doubles** solvable depth

A simple example of the value of reasoning about which computations are relevant (a form of metareasoning)

Unfortunately, 35^{50} is still impossible!

Resource limits

Standard approach:

• Use CUTOFF-TEST instead of TERMINAL-TEST e.g., depth limit (perhaps add quiescence search)

 \bullet Use Eval instead of $\operatorname{UTILITY}$

i.e., evaluation function that estimates desirability of position

Suppose we have 100 seconds, explore 10^4 nodes/second

 $\Rightarrow 10^6$ nodes per move $\approx 35^{8/2}$

 $\Rightarrow \alpha – \beta$ reaches depth 8 \Rightarrow pretty good chess program

Evaluation functions



Black to move

White slightly better

1 1 ¥ Ŧ £ <u>¥</u> 윺 - ₽ ₩ 윤 윺 윤 율 且 White to move

f

Black winning

f

I

Ŧ

For chess, typically linear weighted sum of features

 $Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \ldots + w_n f_n(s)$

e.g., $w_1 = 9$ with $f_1(s) = ($ number of white queens) - (number of black queens), etc.



Behaviour is preserved under any monotonic transformation of EVAL

Only the order matters:

payoff in deterministic games acts as an ordinal utility function

Deterministic games in practice

Checkers: Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994. Used an endgame database defining perfect play for all positions involving 8 or fewer pieces on the board, a total of 443,748,401,247 positions.

Chess: Deep Blue defeated human world champion Gary Kasparov in a sixgame match in 1997. Deep Blue searches 200 million positions per second, uses very sophisticated evaluation, and undisclosed methods for extending some lines of search up to 40 ply.

Othello: human champions refuse to compete against computers, who are too good.

Go: human champions refuse to compete against computers, who are too bad. In go, b > 300, so most programs use pattern knowledge bases to suggest plausible moves.

Nondeterministic games: backgammon



Nondeterministic games in general

In nondeterministic games, chance introduced by dice, card-shuffling Simplified example with coin-flipping:



Algorithm for nondeterministic games

EXPECTIMINIMAX gives perfect play

. . .

Just like $\operatorname{MINIMAX}$, except we must also handle chance nodes:

if state is a MAX node then
 return the highest EXPECTIMINIMAX-VALUE of SUCCESSORS(state)
if state is a MIN node then
 return the lowest EXPECTIMINIMAX-VALUE of SUCCESSORS(state)
if state is a chance node then
 return average of EXPECTIMINIMAX-VALUE of SUCCESSORS(state)

Nondeterministic games in practice

Dice rolls increase b: 21 possible rolls with 2 dice Backgammon \approx 20 legal moves (can be 6,000 with 1-1 roll)

depth $4 = 20 \times (21 \times 20)^3 \approx 1.2 \times 10^9$

As depth increases, probability of reaching a given node shrinks \Rightarrow value of lookahead is diminished

 $\alpha \text{-}\beta$ pruning is much less effective

$$\label{eq:total_total} \begin{split} TDGAMMON \text{ uses depth-2 search} + \text{very good } Eval \\ \approx \text{world-champion level} \end{split}$$



Behaviour is preserved only by positive linear transformation of EVAL

Hence Eval should be proportional to the expected payoff

Games of imperfect information

E.g., card games, where opponent's initial cards are unknown

Typically we can calculate a probability for each possible deal

Seems just like having one big dice roll at the beginning of the game *

Idea: compute the minimax value of each action in each deal, then choose the action with highest expected value over all deals*

Special case: if an action is optimal for all deals, it's optimal.*

GIB, current best bridge program, approximates this idea by1) generating 100 deals consistent with bidding information2) picking the action that wins most tricks on average

LOCAL SEARCH ALGORITHMS

CHAPTER 4, SECTIONS 3–4

Chapter 4, Sections 3–4 1

Outline

- \diamond Hill-climbing
- \diamondsuit Simulated annealing
- ♦ Genetic algorithms (briefly)
- \diamond Local search in continuous spaces (very briefly)

Iterative improvement algorithms

In many optimization problems, **path** is irrelevant; the goal state itself is the solution

Then state space = set of "complete" configurations; find **optimal** configuration, e.g., TSP or, find configuration satisfying constraints, e.g., timetable

In such cases, can use iterative improvement algorithms; keep a single "current" state, try to improve it

Constant space, suitable for online as well as offline search

Example: Travelling Salesperson Problem

Start with any complete tour, perform pairwise exchanges



Variants of this approach get within 1% of optimal very quickly with thousands of cities

Example: *n*-queens

Put n queens on an $n \times n$ board with no two queens on the same row, column, or diagonal

Move a queen to reduce number of conflicts



Almost always solves *n*-queens problems almost instantaneously for very large *n*, e.g., n = 1million

Hill-climbing (or gradient ascent/descent)

```
"Like climbing Everest in thick fog with amnesia"
```

Hill-climbing contd.

Useful to consider state space landscape



Random-restart hill climbing overcomes local maxima—trivially complete Random sideways moves ©escape from shoulders ©loop on flat maxima

Simulated annealing

Idea: escape local maxima by allowing some "bad" moves but gradually decrease their size and frequency

```
function SIMULATED-ANNEALING (problem, schedule) returns a solution state
   inputs: problem, a problem
              schedule, a mapping from time to "temperature"
   local variables: current, a node
                        next, a node
                        T, a "temperature" controlling prob. of downward steps
   current \leftarrow Make-Node(INITIAL-STATE[problem])
   for t \leftarrow 1 to \infty do
        T \leftarrow schedule[t]
        if T = 0 then return current
        next \leftarrow a randomly selected successor of current
        \Delta E \leftarrow \text{VALUE}[next] - \text{VALUE}[current]
        if \Delta E > 0 then current \leftarrow next
        else current \leftarrow next only with probability e^{\Delta E/T}
```

Local beam search

Idea: keep k states instead of 1; choose top k of all their successors

Not the same as k searches run in parallel! Searches that find good states recruit other searches to join them

Problem: quite often, all k states end up on same local hill

ldea: choose k successors randomly, biased towards good ones

Observe the close analogy to natural selection!

Genetic algorithms

= stochastic local beam search + generate successors from **pairs** of states



Fitness Selection

Pairs

Mutation

Genetic algorithms contd.

GAs require states encoded as strings (GPs use programs)

Crossover helps iff substrings are meaningful components



 $GAs \neq evolution: e.g., real genes encode replication machinery!$

CONSTRAINT SATISFACTION PROBLEMS

Chapter 5

Outline

- \diamond CSP examples
- \diamond Backtracking search for CSPs
- \diamondsuit Problem structure and problem decomposition
- \diamondsuit Local search for CSPs

Constraint satisfaction problems (CSPs)

Standard search problem:

state is a "black box"—any old data structure that supports goal test, eval, successor

CSP:

state is defined by variables X_i with values from domain D_i

goal test is a set of constraints specifying allowable combinations of values for subsets of variables

Simple example of a formal representation language

Allows useful **general-purpose** algorithms with more power than standard search algorithms



Example: Map-Coloring contd.



Solutions are assignments satisfying all constraints, e.g., $\{WA = red, NT = green, Q = red, NSW = green, V = red, SA = blue, T = green\}$

Constraint graph

Binary CSP: each constraint relates at most two variables

Constraint graph: nodes are variables, arcs show constraints



General-purpose CSP algorithms use the graph structure to speed up search. E.g., Tasmania is an independent subproblem!
Varieties of CSPs

Discrete variables

finite domains; size $d \Rightarrow O(d^n)$ complete assignments

 \diamond e.g., Boolean CSPs, incl. Boolean satisfiability (NP-complete) infinite domains (integers, strings, etc.)

- \diamondsuit e.g., job scheduling, variables are start/end days for each job
- \diamond need a constraint language, e.g., $StartJob_1 + 5 \leq StartJob_3$
- \diamondsuit linear constraints solvable, nonlinear undecidable

Continuous variables

- $\diamondsuit\,$ e.g., start/end times for Hubble Telescope observations
- \diamondsuit linear constraints solvable in poly time by LP methods

Varieties of constraints

Unary constraints involve a single variable, e.g., $SA \neq green$

Binary constraints involve pairs of variables, e.g., $SA \neq WA$

Higher-order constraints involve 3 or more variables, e.g., cryptarithmetic column constraints

Preferences (soft constraints), e.g., red is better than green often representable by a cost for each variable assignment \rightarrow constrained optimization problems

Example: Cryptarithmetic

T W O + T W O F O U R



Variables: $F \ T \ U \ W \ R \ O \ X_1 \ X_2 \ X_3$ Domains: $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$ Constraints

alldiff(F, T, U, W, R, O) $O + O = R + 10 \cdot X_1, \text{ etc.}$

Real-world CSPs

Assignment problems

e.g., who teaches what class

Timetabling problems

e.g., which class is offered when and where?

Hardware configuration

Spreadsheets

Transportation scheduling

Factory scheduling

Floorplanning

Notice that many real-world problems involve real-valued variables

Standard search formulation (incremental)

Let's start with the straightforward, dumb approach, then fix it

States are defined by the values assigned so far

 \diamond Initial state: the empty assignment, $\{\}$

 ♦ Successor function: assign a value to an unassigned variable that does not conflict with current assignment.
 ⇒ fail if no legal assignments (not fixable!)

 \diamondsuit Goal test: the current assignment is complete

1) This is the same for all CSPs! \bigcirc

2) Every solution appears at depth n with n variables

 \Rightarrow use depth-first search

3) Path is irrelevant, so can also use complete-state formulation

4) $b = (n - \ell)d$ at depth ℓ , hence $n!d^n$ leaves!!!! \bigotimes

Backtracking search

Variable assignments are commutative, i.e.,

[WA = red then NT = green] same as [NT = green then WA = red]

Only need to consider assignments to a single variable at each node $\Rightarrow b = d$ and there are d^n leaves

Depth-first search for CSPs with single-variable assignments is called backtracking search

Backtracking search is the basic uninformed algorithm for CSPs

Can solve *n*-queens for $n \approx 25$

Backtracking search

```
function BACKTRACKING-SEARCH(csp) returns solution/failure
return RECURSIVE-BACKTRACKING(\{ \}, csp)
function RECURSIVE-BACKTRACKING(assignment, csp) returns soln/failure
if assignment is complete then return assignment
var \leftarrow SELECT-UNASSIGNED-VARIABLE(VARIABLES[csp], assignment, csp)
for each value in ORDER-DOMAIN-VALUES(var, assignment, csp) do
if value is consistent with assignment given CONSTRAINTS[csp] then
add {var = value} to assignment
result \leftarrow RECURSIVE-BACKTRACKING(assignment, csp)
if result \neq failure then return result
remove {var = value} from assignment
return failure
```









Improving backtracking efficiency

General-purpose methods can give huge gains in speed:

- 1. Which variable should be assigned next?
- 2. In what order should its values be tried?
- 3. Can we detect inevitable failure early?
- 4. Can we take advantage of problem structure?

Minimum remaining values

Minimum remaining values (MRV): choose the variable with the fewest legal values



Degree heuristic

Tie-breaker among MRV variables

Degree heuristic:

choose the variable with the most constraints on remaining variables



Least constraining value

Given a variable, choose the least constraining value: the one that rules out the fewest values in the remaining variables



Combining these heuristics makes 1000 queens feasible









Constraint propagation

Forward checking propagates information from assigned to unassigned variables, but doesn't provide early detection for all failures:



NT and SA cannot both be blue!

Constraint propagation repeatedly enforces constraints locally

Simplest form of propagation makes each arc consistent

 $X \to Y$ is consistent iff

for **every** value x of X there is **some** allowed y



Simplest form of propagation makes each arc consistent

 $X \to Y$ is consistent iff

for **every** value x of X there is **some** allowed y



Simplest form of propagation makes each arc consistent

X o Y is consistent iff

for **every** value x of X there is **some** allowed y



If X loses a value, neighbors of X need to be rechecked

Simplest form of propagation makes each arc consistent

X o Y is consistent iff

for **every** value x of X there is **some** allowed y



If X loses a value, neighbors of X need to be rechecked Arc consistency detects failure earlier than forward checking Can be run as a preprocessor or after each assignment

Arc consistency algorithm

```
function AC-3(csp) returns the CSP, possibly with reduced domains

inputs: csp, a binary CSP with variables \{X_1, X_2, \ldots, X_n\}

local variables: queue, a queue of arcs, initially all the arcs in csp

while queue is not empty do

(X_i, X_j) \leftarrow \text{REMOVE-FIRST}(queue)

if REMOVE-INCONSISTENT-VALUES(X_i, X_j) then

for each X_k in NEIGHBORS[X_i] do

add (X_k, X_i) to queue
```

```
function REMOVE-INCONSISTENT-VALUES(X_i, X_j) returns true iff succeeds

removed \leftarrow false

for each x in DOMAIN[X_i] do

if no value y in DOMAIN[X_j] allows (x,y) to satisfy the constraint X_i \leftrightarrow X_j

then delete x from DOMAIN[X_i]; removed \leftarrow true

return removed
```

 $O(n^2d^3)$, can be reduced to $O(n^2d^2)$ (but detecting all is NP-hard)

Problem structure



Tasmania and mainland are independent subproblems

Identifiable as connected components of constraint graph

Problem structure contd.

Suppose each subproblem has c variables out of n total

Worst-case solution cost is $n/c \cdot d^c$, **linear** in n

E.g., n = 80, d = 2, c = 20 $2^{80} = 4$ billion years at 10 million nodes/sec $4 \cdot 2^{20} = 0.4$ seconds at 10 million nodes/sec

Tree-structured CSPs



Theorem: if the constraint graph has no loops, the CSP can be solved in ${\cal O}(n\,d^2)$ time

Compare to general CSPs, where worst-case time is $O(d^n)$

This property also applies to logical and probabilistic reasoning: an important example of the relation between syntactic restrictions and the complexity of reasoning.

Algorithm for tree-structured CSPs

1. Choose a variable as root, order variables from root to leaves such that every node's parent precedes it in the ordering



- 2. For *j* from *n* down to 2, apply REMOVEINCONSISTENT($Parent(X_j), X_j$)
- 3. For j from 1 to n, assign X_j consistently with $Parent(X_j)$

Nearly tree-structured CSPs

Conditioning: instantiate a variable, prune its neighbors' domains



Cutset conditioning: instantiate (in all ways) a set of variables such that the remaining constraint graph is a tree

Cutset size $c \ \Rightarrow \ {\rm runtime} \ O(d^c \cdot (n-c)d^2),$ very fast for small c

Iterative algorithms for CSPs

Hill-climbing, simulated annealing typically work with "complete" states, i.e., all variables assigned

To apply to CSPs: allow states with unsatisfied constraints operators **reassign** variable values

Variable selection: randomly select any conflicted variable

Value selection by min-conflicts heuristic: choose value that violates the fewest constraints i.e., hillclimb with h(n) = total number of violated constraints

Example: 4-Queens

States: 4 queens in 4 columns ($4^4 = 256$ states)

Operators: move queen in column

Goal test: no attacks

Evaluation: h(n) = number of attacks



Performance of min-conflicts

Given random initial state, can solve n-queens in almost constant time for arbitrary n with high probability (e.g., n = 10,000,000)

The same appears to be true for any randomly-generated CSP **except** in a narrow range of the ratio



Summary

CSPs are a special kind of problem: states defined by values of a fixed set of variables goal test defined by constraints on variable values

Backtracking = depth-first search with one variable assigned per node

Variable ordering and value selection heuristics help significantly

Forward checking prevents assignments that guarantee later failure

Constraint propagation (e.g., arc consistency) does additional work to constrain values and detect inconsistencies

The CSP representation allows analysis of problem structure

Tree-structured CSPs can be solved in linear time

Iterative min-conflicts is usually effective in practice