


CS3243 Introduction to Artificial Intelligence

AY2021/22 Semester 2

1. Introduction

- Agent Function: $f: P \rightarrow a_t$, where P is the sequence of percepts captured by sensors and $a_t \in A$ is the selected action by activators.
- Rational agent optimises performance measure.
 - An agent that senses only **partial** information can also be perfectly rational.
- Environment Properties:
 - Fully Observable vs Partially Observable
 - Deterministic vs Stochastic: Whether immediate state can be determined based on action.
 - Episodic vs Sequential: Whether actions only impact current state or all future states.
 - Discrete vs Continuous
 - Single-agent vs Multi-agent: Opponents might be competitive or cooperative.
 - Known vs Unknown: Refers to the agent/designer.
 - Static vs Dynamic
- Taxonomy of Agents:
 - Reflex Agents: Uses if-statements.
 - Model-based Reflex Agents: Makes decisions based on an internal model.
 - Goal-based/Utility-based Agents
 - Learning Agents



2. Uninformed Search

- Formulation of search problem:
 - State Representation (s_i): ADT containing data describing an instance of the environment.
 - Initial State (s_0): Initial values of the data above.
 - Action
 - Transition Model: How each data change corresponding to the action given.
 - Step Cost
 - Goal Test
- Uninformed Search: No domain knowledge beyond search problem formulation.
- General Search Algorithm:

```

frontier = {initial state}
while frontier not empty:
    current = frontier.pop()
    if current is goal:
        return path found
    for a in actions(current):
        frontier.push(T(current, a))
return failure
    
```

BFS – queue
 DFS – stack
 UCS – pq

- State vs Node:
 - State: A representation of the environment at some timestamp.
 - Node: Includes state, parent node, action, path cost (for UCS), depth.
- Algorithm Criteria:
 - Time/Space Complexity
 - Completeness: Complete if can find a solution when one exists and report failure if it does not.
 - Optimality: Optimal if it finds a solution with the lowest path cost among all solutions.
- Tree Search vs Graph Search: In graph search, we only add nodes to frontier and reached if (1) state represented by node not previously reached and (2) path to state already reached is cheaper than the one stored.
- Performance Summary:
 - If b is finite and state space is finite or contains a goal.
 - Same as 1.
 - If action costs are all equal.
 - b is branching factor, d is depth of shallowest goal. Can be improved by early goal test (when pushing): Assuming the worst case, we can save the time and space associated with $(b^d - b)$ nodes.
 - May get caught in a cycle.
 - If search space is finite.
 - Where m is the maximum depth. Can be improved to $O(m)$ by backtracking.
 - If BFS is complete and all action cost $> \epsilon > 0$. Must perform late goal test (when popping).


9. $e = 1 + \lceil C^*/\epsilon \rceil$, where C^* is the optimal path cost and ϵ is some small positive constant.

10. Where l is the limited depth.
11. Number of nodes explored: $(d + 1)O(b^l) + dO(b^{l-1}) + \dots + O(b^d)$.

Criterion	BFS	DFS	UCS	DLS	IDS ¹¹
Complete (Tree/Graph)	\checkmark^1	\times^5	\checkmark^8	\times	\checkmark
Optimal	\times^3	\times	\checkmark	\times	\checkmark
Time (Tree/Graph)	$O(b^d)^4$	$O(b^m)$	$O(b^e)^9$	$O(b^l)^{10}$	$O(b^d)$
Space (Tree/Graph)	$O(b^d)$	$O(bm)^7$	$O(b^e)$	$O(bl)$	$O(bl)$
	$O(V + E)$				

3. Informed Search

- Heuristic Function (h): Approximates the path cost from n . state to its nearest goal G .



- $h^*(n)$: True path cost from n to G .
- Evaluation Function (f): Priority for a node n .
- Best-First Search Algorithm:


```

frontier = {initial state}
reached = {}
while frontier not empty:
    current = frontier.pop()
    if current is goal:
        return path found
    for a in actions(current):
        next = T(current, a)
        s = next.state
        if s not reached or
            next has a smaller path cost:
            reached[s] = next
            frontier.push(next)
return failure
    
```
- Greedy Best-First Search:
 - $f(n) = h(n)$
 - Tree search version is incomplete (may get stuck in a loop between nodes where h values are lowest). Graph search version is complete if search space is finite.
 - Optimal: No.
- A* Search:
 - $f(n) = g(n) + h(n)$
 - Limited Graph Search Version 1: No exceptions on lower path costs.
 - Limited Graph Search Version 2: Adds to reached when popping.
 - Complete if UCS is complete.
 - Tree search and graph search version is optimal if h is admissible. Limited graph search version 2 is optimal if h is consistent.
- Heuristics
 - Admissible: $h(n)$ is admissible if $\forall n, h(n) \leq h^*(n)$. If $h(n)$ is admissible, then A* Search using tree/graph search is optimal.
 - Consistent: $h(n)$ is consistent if $\forall n, n', h(n) \leq \text{cost}(n, a, n') + h(n')$, where n' denotes all successors of n . If $h(n)$ is consistent, then A* Search using limited graph search version 2 is optimal.
 - If h is consistent, then it is admissible.
 - If $\forall n, h_1(n) \geq h_2(n)$, then h_1 dominates h_2 . If h_1 is admissible, then it is more efficient than or as efficient as h_2 .

4. Local Search

- Hill-Climbing Algorithm:


```


current = initial state
while true:
    neighbour = highest-valued successor
    if neighbour.value <= current.value:
        return current
    current = neighbour
    
```

 - $f(n) = -h(n)$
 - May get stuck at (1) local maxima, (2) shoulder or plateau and (3) ridge (sequence of local maxima).
 - Sideways Move: Replaces \leq with $<$. This allows the algorithm to traverse shoulders.
 - Stochastic Hill Climbing: Chooses randomly among states with values better than $current$. May take longer to find a solution but sometimes leads to better solutions.
 - First-choice Hill Climbing: Randomly generating successors until one better than $current$ is found.
 - Random-restart Hill Climbing: Adds an outer loop which randomly picks a new starting state. Keeps attempting random restarts until a solution is found.
- Local Beam Search:
 - Always stores k states instead of 1.
 - Begins with k random starts and chooses best k among all successors until a solution is found.

- o May be improved by stochastic.

5. Constraint Satisfaction Problems

- Formulation of CSPs:
 - o State Representation (s_i): Variables ($X = \{x_1, x_2, \dots, x_n\}$) and their domains ($D = \{d_1, d_2, \dots, d_n\}$).
 - o Initial State (s_0): All variables unassigned.
 - o Action
 - o Transition Model
 - o Goal Test: Whether all constraints $C = \{c_1, c_2, \dots, c_m\}$ is satisfied. Each constraint corresponds to a subset of X . Each constraint contains a scope and a relation (e.g. scope = (x_1, x_2) ; relation = $x_1 < x_2$).
- Constraint Graph:
 - o One binary constraint is two arcs.



- Backtracking Algorithm:


```

assignment = {}
function backtrack(csp, assignment):
    if assignment is complete:
        return assignment
    var = select unassigned variable
        from assignment
    for each value in domain of var:
        if value consistent with csp:
            assignment.add {var=value}
            inference = infer(assignment)
            if inference not failure:
                csp.add(inference)
                result = backtrack(csp,
                    assignment)
                if result not failure:
                    return result
            csp.remove(inference)
            assignment.remove {var=value}
    return failure
            
```

 - o Total number of leaves: d^n , where d is number of values and n is number of variables.
 - o Solution is found at depth n .
 - o Variable order: Minimum-remaining-value heuristic + degree heuristic
 - o Value order: Least-constraining-value heuristic
 - o Maximum backtrack times: $O(2^n)$ (whole tree)
- AC-3 Algorithm

```

function AC3(csp):
    queue = a queue of all arcs
    while queue is not empty do:
        (Xi, Xj) = queue.pop()
        if REVISE(csp, Xi, Xj):
            if size of Di == 0:
                return false
            for Xk in Xi.neighbours -
                {Xj}:
                add (Xk, Xi) to queue
    return true

function REVISE(csp, Xi, Xj):
    revised = false
    for each x in Di do:
        if no y in Dj satisfies
            constraint:
                delete x from Di
                revised = true
    return revised
            
```

- o Time complexity: $O(n^2d^3)$

6. Adversarial Search

- Formulation of games:
 - o State representation
 - o TO – MOVE(s): The player to move in state s .
 - o ACTIONS(s): The legal moves in state s .
 - o RESULT(s, a): The transition model.
 - o IS – TERMINAL(s): Whether game is over.
 - o UTILITY(s, p): Defines the final numerical value to player p when the game ends in state s .
- Winning Strategy: A winning strategy for player 1, s_1^* , implies that for any strategy s_2 for player 2, the game ends in a win for player 1.
- Minimax Algorithm:

$$Minimax(s) = \begin{cases} UTILITY(s, MAX) \\ \max_{a \in ACTIONS(s)} Minimax(RESULT(s, a)) \\ \min_{a \in ACTIONS(s)} Minimax(RESULT(s, a)) \end{cases}$$

- o Minimax Algorithm is complete if game tree is finite, is optimal.
- o Time: $O(b^m)$; Space: $O(bm)$.
- $\alpha - \beta$ Pruning:
 - o At MAX node n , $\alpha(n)$ is highest observed value found on path from n . Initially $\alpha(n) = -\infty$.
 - o At MIN node n , $\beta(n)$ is lowest observed value found on path from n . Initially $\beta(n) = \infty$.
 - o Given a MIN node n , stop searching below n if some MAX ancestor i with $\alpha(i) \geq \beta(n)$.
 - o Given a MAX node n , stop searching below n if some MIN ancestor i with $\alpha(n) \geq \beta(i)$.

```

function AB-PRUNING(node, alpha, beta):
    if node is leaf:
        return node.val
    if isMaxTurn:
        bestVal = -INFINITY
        for each child:
            value = AB-PRUNING(child)
            bestVal = max(bestVal, value)
            alpha = max(bestVal, alpha)
            if beta <= alpha:
                break
        return bestVal
    if isMinTurn:
        bestVal = +INFINITY
        for each child:
            value = AB-PRUNING(child, alpha, beta)
            bestVal = min(bestVal, value)
            beta = min(bestVal, beta)
            if beta <= alpha:
                break
        return bestVal
            
```

7. Knowledge Representation

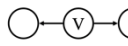
- Knowledge:
 - o Knowledge Base: Set of sentences in a formal language prepopulated with domain knowledge. It is domain-specific content, while inference engine is domain-independent algorithm.
 - o Entailment: $\alpha \models \beta \rightarrow M(\alpha) \subseteq M(\beta)$
- Algorithm Criteria:
 - o KB $\vdash_{\mathcal{A}} \alpha$: Sentence α is derived from KB by inference algorithm \mathcal{A} .
 - o Soundness: \mathcal{A} is sound if KB $\vdash_{\mathcal{A}} \alpha$ implies KB $\models \alpha$. \mathcal{A} does not infer nonsense.
 - o Completeness: \mathcal{A} is complete if KB $\models \alpha$ implies KB $\vdash_{\mathcal{A}} \alpha$. \mathcal{A} can infer any sentence KB entails.
- Truth Table Enumeration: Checks all 2^n truth assignments to verify KB entails α (DFS).
 - o Time: $O(2^n)$; Space: $O(n)$.
 - o Sound and complete.
- Resolution:
 - o Conjunctive Normal Form (CNF): Conjunction of disjunctive sentences, $R_1 \wedge R_2 \wedge \dots \wedge R_n$.
 - o Sound and complete.
 - o Rules:
 1. Convert $\alpha \leftrightarrow \beta$ to $(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)$.
 2. Convert $\alpha \Rightarrow \beta$ to $\neg \alpha \vee \beta$.
 3. Convert $\neg(\alpha \vee \beta)$ to $\neg \alpha \wedge \neg \beta$.
 4. Convert $\neg(\alpha \wedge \beta)$ to $\neg \alpha \vee \neg \beta$.
 5. Convert $\neg(\neg \alpha)$ to α .
 6. Convert $(\alpha \vee (\beta \wedge \gamma))$ to $(\alpha \vee \beta) \wedge (\alpha \vee \gamma)$.
 - o Show that $KB \wedge \neg \alpha$ is unsatisfiable.
 - o $R_1: \alpha \vee \gamma$ and $R_2: \beta \vee \neg \gamma$ are resolved to $\alpha \vee \beta$.

8. Uncertainty

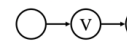
- Bayes Rule:

$$Pr[A|B] = \frac{Pr[B|A] \cdot Pr[A]}{Pr[B]}$$
- Chain Rule:

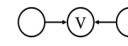
$$Pr[R_1, R_2, \dots, R_k] = \prod_{j=1,2,\dots,k} Pr[R_j|R_1, R_2, \dots, R_{j-1}]$$
- Path Blocking Scenario:



v is given



v is given



v is not given
- Suppose given S , all T_1, T_2, \dots, T_{n-1} are independent, then:

$$Pr[T_1, T_2, \dots, T_{n-1}, S] = Pr[T_1|S] \cdot \dots \cdot Pr[T_{n-1}|S] \cdot Pr[S]$$
 - o A joint distribution for n Boolean random variables results in at least $2^n - 1$ entries.
 - o A joint distribution for n Boolean random variables with conditional independence results in $2n - 1$ entries.