

Algorithm 1 Breadth First Search

```

1:  $F(\text{Frontier}) \leftarrow \text{Queue}(u)$ 
2:  $E(\text{Explored}) \leftarrow \{u\}$ 
3: while  $F$  is not empty do
4:    $u \leftarrow F.\text{pop}()$ 
5:   for all children  $v$  of  $u$  do
6:     if  $\text{GoalTest}(v)$  then
7:       return  $\text{path}(v)$ 
8:     if  $v$  not in  $E$  then
9:        $E.\text{add}(v)$ 
10:       $F.\text{push}(v)$ 
11: return Failure

```

Algorithm 6 Iterative Deepening Search

```

1:  $\text{cutoff} \leftarrow \text{initval}$ 
2: while true do
3:    $\text{ret} \leftarrow \text{FindPathToGoal}(u, \text{cutoff})$ 
4:   if  $\text{ret} == \text{Failure}$  then
5:      $\text{IncreaseCutoff}(\text{cutoff})$ 
6:   else
7:     return  $\text{ret}$ 

```

Algorithm A* Algorithm: FindPathToGoal

```

1:  $F(\text{Frontier}) \leftarrow \text{PriorityQueue}(u)$ 
2:  $E(\text{Explored}) \leftarrow \{u\}$ 
3: Initialize  $\hat{g}$   $\hat{g}[u] \leftarrow 0$ ;  $\hat{g}[v] = \infty, \forall v \neq u$ 
4: while  $F$  is not empty do
5:    $u \leftarrow F.\text{pop}()$ 
6:   if  $\text{GoalTest}(u)$  then
7:     return  $\text{path}(u)$ 
8:    $E.\text{add}(u)$ 
9:   for all successors  $v$  of  $u$  do
10:    if  $v$  not in  $E$  then
11:       $F.\text{push}(v)$ 
12:       $\hat{g}[v] = \hat{g}[u] + c(u, v)$ 
13:       $f[v] = h[v] + \hat{g}[v]$ 
14: return Failure

```

Algorithm 5 Depth First Search(DFS)

```

1:  $F(\text{Frontier}) \leftarrow \text{Stack}(u)$ 
2: while  $F$  is not empty do
3:    $u \leftarrow F.\text{pop}()$ 
4:   if  $\text{GoalTest}(u)$  then
5:     return  $\text{path}(u)$ 
6:   if  $\text{HasUnvisitedChildren}(u)$  then
7:     for all children  $v$  of  $u$  do
8:        $F.\text{push}(v)$ 
9: return Failure

```

Algorithm 5 Depth Limited Search

```

1:  $F(\text{Frontier}) \leftarrow \text{Stack}(u)$ 
2: while  $F$  is not empty do
3:    $u \leftarrow F.\text{pop}()$ 
4:   if  $\text{GoalTest}(u)$  then
5:     return  $\text{path}(u)$ 
6:   if  $\text{depth}(u) \leq \text{cutoff}$  then
7:     for all children  $v$  of  $u$  do
8:        $F.\text{push}(v)$ 
9: return Failure

```

Algorithm 4 Uniform Cost Search(UCS): Find

```

1:  $F(\text{Frontier}) \leftarrow \text{PriorityQueue}(u)$ 
2:  $E(\text{Explored}) \leftarrow \{u\}$ 
3:  $\hat{g}[u] \leftarrow 0$ 
4: while  $F$  is not empty do
5:    $u \leftarrow F.\text{pop}()$ 
6:   if  $\text{GoalTest}(u)$  then
7:     return  $\text{path}(u)$ 
8:    $E.\text{add}(u)$ 
9:   for all children  $v$  of  $u$  do
10:    if  $v$  not in  $E$  then
11:      if  $v$  in  $F$  then
12:         $\hat{g}[v] = \min(\hat{g}[v], \hat{g}[u] + c(u, v))$ 
13:      else
14:         $F.\text{push}(v)$ 
15:         $\hat{g}[v] = \hat{g}[u] + c(u, v)$ 
16: return Failure

```

Property	BFS	UCS	DFS	IDS
Complete	Yes ¹	Yes ²	No	Yes
Optimal	No ³	Yes	No	No
Time	$\mathcal{O}(b^{d+1})$	$\mathcal{O}\left(b^{1+\left\lceil\frac{C^*}{\epsilon}\right\rceil}\right)$	$\mathcal{O}(b^{m+1})$	$\mathcal{O}(b^{d+1})$
Space	$\mathcal{O}(b^{d+1})$	$\mathcal{O}\left(b^{1+\left\lceil\frac{C^*}{\epsilon}\right\rceil}\right)$	$\mathcal{O}(bm)$	$\mathcal{O}(bd)$

1. if b is finite.
2. If b is finite and step cost $\geq \epsilon$

Problem defined by: **states** S and **initial state**, **actions** A , **state transition model** $T(s_1 \in S, a \in A) \mapsto s_2 \in S$, **goal test**, **cost function** $C(a): A \mapsto \mathbb{R}$.

Consistency: $\forall v_i, v_j, h(v_i) \leq h(v_j) + c(v_i, v_j)$ where where v_j is a successor of v_i and $c(v_i, v_j)$ is the cost of moving from v_i to v_j , i.e. for each step we move backwards, the value of the heuristic at the further node \leq the value at the closer node plus the cost of the step.

Admissibility: $\forall v_i, h(v_i) \leq C^*(v_i) = g(w) - g(v_i)$ where $C^*(v)$ is the true optimal cost to reach the goal from v , i.e. the value of the heuristic is always less than or equal to the true cost of reaching the goal.

A^* is optimal for consistent h , or admissible h with tree search. Consistency and $h(w) = 0 \implies$ admissibility. Admissible heuristics can be created by relaxing problem constraints.

For cycle checking, we keep track of all previously seen nodes (expanded or in frontier) *and* their cost. When adding a new node to the frontier, if it matches a previously seen node and has higher (or equal) cost, don't add it. If it matches but has lower cost, add it and prune the higher cost node from the frontier. Note for A^* the cost is taken as $g(n)$ and not $f(n)$.

BACKTRACKINGSEARCH(*prob, assign*)

```
1: if ALLVARSASSIGNED(prob, assign) then
2:   if ISCONSISTENT(assign) then
3:     return assign
4:   else
5:     return failure
6: var ← PICKUNASSIGNEDVAR(prob, assign)
7: for value ∈ ORDERDOMAINVALUE(var, prob, assign) do
8:   assign ← assign ∪ {var = value}
9:   result ← BACKTRACKINGSEARCH(prob, assign)
10:  if result! = failure then return result
11:  assign ← assign \ {var = value}
12: return failure
```

INFER(*prob, var, assign*)

```
1: inference ← ∅
2: varQueue ← [var]
3: while varQueue is not empty do
4:   y ← varQueue.pop()
5:   for each constraint C in prob where y ∈ Vars(C) do
6:     for all x ∈ Vars(C) \ y do
7:       S ← COMPUTEDOMAIN(x, assign, inference)
8:       for each value v in S do
9:         if no valid value exists for all var ∈ Var(C) \ x s.t. C[x ⊢ v] is satisfied then
10:          inference ← inference ∪ {x ∉ {v}}
11:       T ← COMPUTEDOMAIN(x, assign, inference)
12:       if T = ∅ then return failure
13:       if S ≠ T then
14:         varQueue.add(x)
15: return inference
```

Algorithm 3 BacktrackingSearch_with_Inference(*prob, assign*)

```
1: if ALLVARIABLESASSIGNED(prob, assign) then return assign
2: var ← PICKUNASSIGNEDVAR(prob, assign)
3: for value in ORDERDOMAINVALUE(var, prob, assign) do
4:   if VALISCONSISTENTWITHASSIGNMENT(value, assign) then
5:     assign ← assign ∪ {(var = value)}
6:     inference ← INFER(prob, var, assign)
7:     assign ← assign ∪ inference
8:     if inference! = failure then
9:       result ← BACKTRACKINGSEARCH(prob, assign)
10:      if result! = failure then return result
11:      assign ← assign \ inference
12:      assign ← assign \ {(var = value)}
13: return failure
```

Algorithm 2 BacktrackingSearch(*prob, assign*)

```
1: if ALLVARASSIGNED(prob, assign) then return assign
2: var ← PICKUNASSIGNEDVAR(prob, assign)
3: for value in ORDERDOMAINVALUE(var, prob, assign) do
4:   if VALISCONSISTENTWITHASSIGNMENT(value, assign) then
5:     assign ← assign ∪ {(var = value)}
6:     result ← BACKTRACKINGSEARCH(prob, assign)
7:     if result! = failure then return result
8:     assign ← assign \ {(var = value)}
9: return failure
```

Problem defined by: **variables** $X = \{x_1, x_2, \dots, x_n\}$, **domains** $D = \{D_{x_1}, D_{x_2}, \dots, D_{x_n}\}$, **constraints** C relating to domains. NP-Hard.

Default inference is expensive. Methods to limit depth: **forward checking**: only check the effect of assigning the current variable (don't add x to queue); or add to queue only if the variable has fewer than N valid values left, $|T| \leq N$.

For InFER, only operate on the variables that haven't been assigned. Heuristics: **Minimum Remaining Value**: PickUnassignedVar choose the unassigned variable with the smallest effective domain size (allows quick backtracking); **Least Constraining Value**: OrderDomainValue choose the value that rules out the fewest domain values for other variables (so we maximize the chances of success).

Variants: boolean satisfiability (B-SAT), where $D_{x_i} = \{0, 1\}$; binary CSP, where every constraint $C(x_i, x_j)$ is over 2 variables. 2-SAT: combination of the two, is the only one with polynomial time solutions.

Type	Characteristics	Analogy
Simple Reflex	Action depends only on percept	Infant <ul style="list-style-type: none">• If Hungry then Cry• If Happy then Sleep
Model-based Reflex	Action depends on internal state (based on percept history), model of the world, and percept	Kid If Want candy & Didn't Receive candy then Ask for candy
Goal-based	Action depends on current state, percepts, model of the world Action plan to achieve desired goal	Teenager Goal: Want to get into UofT Plan: Study and Sports Activities based on the Goal.
Utility-based	Useful for multiple (possible) conflicting goals A weighted combination of goals	Adult $\alpha * \text{Job} + \beta * \text{Partner} + \gamma * \text{Health}$

Fully observable (vs. partially observable):

- sensors provide access to the complete state of the environment at each point in time.

Deterministic (vs. stochastic)

- The next state of the environment is completely determined by the current state and the action executed by the agent.

Static (vs. dynamic)

- The environment is unchanged while an agent is deliberating.

Discrete (vs. continuous)

- A finite number of distinct states, percepts, and actions.

Single agent (vs. multi-agent)

- An agent operating by itself in an environment.

$$\text{MiniMax}(s) = \begin{cases} \text{Utility}(s) & s \text{ is a terminal state} \\ \max_{a \in A} \text{MiniMax}(\text{Result}(s, a)) & \text{MAX player} \\ \min_{a \in A} \text{MiniMax}(\text{Result}(s, a)) & \text{MIN player} \end{cases}$$

DFS: $\mathcal{O}(b^d)$ time, $\mathcal{O}(bd)$ space (best and worst).
 Searches entire tree; needs finite search space.
 Estimate utils for non-terminal states if too deep.

Algorithm 6 α - β -MaxValue(s, α, β)

```

1: if IsTerminal(s)=True then return UTILITY(s)
2: v ← -∞
3: for all action a in Actions(s) do
4:   v ← max(v, α-β-MinValue(Result(s, a), α, β))
5:   if v ≥ β then return v
6:   α ← max(α, v)
7: return v
  
```

Algorithm 5 α - β -MinValue(s, α, β)

```

1: if IsTerminal(s)=True then return UTILITY(S)
2: v ← ∞
3: for all action a in Actions(s) do
4:   v ← min(v, α-β-MaxValue(Result(s, a), α, β))
5:   if v ≤ α then return v
6:   β ← min(β, v)
7: return v
  
```

Algorithm 4 α - β -MinimaxDecision(S)

```

1: bestAction ← null
2: max ← -∞
3: for all action a in Actions(s) do
4:   util ← α-β-MinValue(Result(s, a), max, ∞)
5:   if util > max then
6:     bestAction ← a
7:     max ← util
8: return bestAction
  
```




α : Best value for MAX along path to root (init $-\infty$). β : Best value for MIN along path to root (init ∞).
 If $\alpha \geq \beta$, prune remaining children. $\mathcal{O}(b^{\frac{d}{2}})$ time (best), $\mathcal{O}(b^{\frac{3}{2}d})$ (average), $\mathcal{O}(b^d)$ (worst).

$$P(A \cup B) = P(A) + P(B) - P(A \cap B) \quad P(A|B) = \sum_{C_i} P(A|B, C_i)P(C_i|B) \quad P(X_1, \dots, X_N) = \prod_{i=1}^N P(X_i | \text{parents}(X_i))$$

Independence: $P(A|B) = P(A) \iff P(B|A) = P(B) \iff P(A, B) = P(A)P(B)$

Conditional Independence: $P(A|B, C) = P(A|C) \iff P(B|A, C) = P(B|C) \iff P(A, B|C) = P(A|C)P(B|C)$

Bayesian Network Law: Given all parents, the probability of any node is independent of its non-descendants.



- $Z \in \mathcal{E}$ and one arc on \mathcal{P} enters Z and another leaves Z
- $Z \in \mathcal{E}$ and both arcs on \mathcal{P} leave Z
- Both arcs on \mathcal{P} enter Z , and $Z, \text{descendants}(Z) \notin \mathcal{P}$

- Replace each factor $f \in F$ that mentions a variable(s) in E with its restriction $f_{E=e}$ (this might yield a factor over no variables, a constant)
- For each Z_j —in the order given—eliminate $Z_j \in Z$ as follows:
 - Compute **new factor** $g_j = \sum_{Z_j} f_1 * f_2 * \dots * f_k$, where the f_i are the factors in F that include Z_j
 - Remove** the factors f_i that mention Z_j from F and add new factor g_j to F

Restrict: $f_{E=e}$ takes only rows that have value e for E ; E is no longer a variable in the factor.

Multiply: $f_1 * f_2$ combines tables, multiplying rows with the same values for shared variables.

Sum out: $\sum_{Z_j} f$ sums rows with the same values for all other variables; Z_j is no longer a variable.

- The remaining factors refer only to the query variable Q .
Take their product and normalize to produce $\text{Pr}(Q|e)$

Complexity: Exponential in time and space to **elimination width** k .

Polytree: single path between any pair of nodes; there exists a VE order that never increases factor size.

Min-fill heuristic: always eliminate the variable that creates the smallest factor.

Guarantees linear time for polytrees in the network size (not # of vars).

Commutative law	$p \wedge q \equiv q \wedge p$	$p \vee q \equiv q \vee p$
Associative law	$(p \wedge q) \wedge r \equiv p \wedge (q \wedge r)$	$(p \vee q) \vee r \equiv p \vee (q \vee r)$
Distributive law	$p \wedge (q \vee r) \equiv (p \wedge q) \vee (p \wedge r)$	$p \vee (q \wedge r) \equiv (p \vee q) \wedge (p \vee r)$
Identity law	$p \wedge \text{true} \equiv p$	$p \vee \text{false} \equiv p$
Universal bound law	$p \vee \text{true} \equiv \text{true}$	$p \wedge \text{false} \equiv \text{false}$
Idempotent law	$p \wedge p \equiv p$	$p \vee p \equiv p$
Negation law	$p \vee \neg p \equiv \text{true}$	$p \wedge \neg p \equiv \text{false}$
Double negation law	$\neg(\neg p) \equiv p$	
de Morgan's law	$\neg(p \wedge q) \equiv \neg p \vee \neg q$	$\neg(p \vee q) \equiv \neg p \wedge \neg q$
Absorption law	$p \vee (p \wedge q) \equiv p$	$p \wedge (p \vee q) \equiv p$
Implication law	$p \rightarrow q \equiv \neg p \vee q$	

Propositional Logic

Operators: $\neg A$ (negation), $A \wedge B$ (conjunction/and), $A \vee B$ (disjunction/or), $A \rightarrow B$ (implication; $\neg A \vee B$), $A \leftrightarrow B$ (bi-implication).

τ satisfies A iff $\bar{\tau}(A)$ is true. τ satisfies Φ iff τ satisfies all formulas in Φ .

A is a **logical consequence** of Φ ($\Phi \models A$) iff for all τ , if τ satisfies Φ , then it satisfies A .

Limited by boolean variables (cannot cross reference between individuals in a statement) and the lack of quantifiers (having to list all members to specify a property).

First-Order Logic

Defined by: Variables V , functions F , predicates P . **Terms:** variables or functions of terms. 0-ary functions are constant terms (cannot be quantified). **Vocabulary:** \mathcal{L} , a set of function and predicate symbols. **Formula:** *atomic* ($P(t_1, \dots, t_n)$ where t_i are terms), or formulas combined with propositional operators, or \exists and \forall quantifiers.

Converting from English: things become constants, types/properties become unary predicates, relationships become binary (or more) predicates, associations become functions.

Structure: an \mathcal{L} -structure \mathcal{M} contains the **universe** $M \neq \emptyset$, function extensions $f^{\mathcal{M}}: M^n \mapsto M$ for each $f \in \mathcal{L}$ (specified as individual mappings), predicate extensions $P^{\mathcal{M}} \subseteq M^n$ for each $P \in \mathcal{L}$ (specified as sets of n -tuples $\langle A, B, \dots \rangle$ for which the predicate is true).

Object assignment: an object assignment σ for \mathcal{M} is a mapping from a set of variables to the universe of M . Note $\sigma(m/x)$ is an assignment mapping x to $m \in M$ (for quantifiers).

Satisfaction: \mathcal{M} is a model of C under σ (denoted $\mathcal{M} \models C[\sigma]$) if C is true, under the definitions and variable mappings of \mathcal{M} and σ .

x is **bounded** in A if it only exists in A under a quantifier; otherwise it is **free**. If σ and σ' have the same assignment for all free variables of A , then $\mathcal{M} \models A[\sigma] \iff \mathcal{M} \models A[\sigma']$. A is a **sentence** if it is **closed** (no free variables). For sentences, σ is irrelevant and can be dropped. \mathcal{M} is a **model** of Φ ($\mathcal{M} \models \Phi$) if it satisfies all sentences in Φ . Φ is **satisfiable** if there exists \mathcal{M} that models Φ .

A is a **logical consequence** of Φ ($\Phi \models A$) iff for every \mathcal{M} , $\mathcal{M} \models \Phi \implies \mathcal{M} \models A$. If $\Phi \models A$, then $\nexists \mathcal{M}$ s.t. $\mathcal{M} \models \Phi \cup \{\neg A\}$.

Knowledge base: a collection of sentences representing the agent's beliefs, can be used for inference about implicit knowledge through proof procedures. Procedures are **sound** if it only produces logical consequences of the KB, and **complete** if it can produce all logical consequences of the KB.

Resolution by refutation: a sound and complete proof procedure; first assume $\neg A$ (refutation), then convert $\neg A$ and KB to a clausal theory C , and resolve clauses in C until the empty clause is reached, where we conclude A is true.

Clausal theory: a set (conjunction) of **clauses** that must all be true; each clause is a disjunction of **literals** (at least one is true); each literal is either an atomic formula or a negated atomic formula.

Resolution:
$$\frac{a_1 \vee \dots \vee a_n \vee c \quad b_1 \vee \dots \vee b_m \vee \neg c}{a_1 \vee \dots \vee a_n \vee b_1 \vee \dots \vee b_m}$$

Conversion to clausal form:

1. Convert implications: $A \rightarrow B$ to $\neg A \vee B$.
2. Move negations inward and simplify double negations: $\neg(A \wedge B) \iff \neg A \vee \neg B$, $\neg(A \vee B) \iff \neg A \wedge \neg B$, $\neg\neg xA \iff \exists x\neg A$, $\neg\neg\neg xA \iff \forall x\neg A$.
3. Variable standardization: rename so that each quantified variable is unique.
4. Skolemization: remove \exists by replacing the variable quantified with a new unique constant $\exists xP(x) \iff P(\mathbf{a})$, or unique function which mentions every variable that scopes the existential $\forall x\exists yP(x, y) \iff \forall xP(x, g(x))$.
5. Prenex form: take all quantifiers (only \forall at this point) to the front: $\forall xP \wedge Q \iff \forall x(P \wedge Q)$, $\forall xP \vee Q \iff \forall x(P \vee Q)$.
6. Distribute \vee over \wedge : $A \vee (B \wedge C) \iff (A \vee B) \wedge (A \vee C)$.
7. Convert to clauses: remove all \forall quantifiers (implicit) and break apart \wedge .

Resolution is **refutation complete**: it can eventually prove that a clausal theory is unsatisfiable, but may not terminate when it is satisfiable. In general first-order unsatisfiability is semi-decidable (there exists an algorithm that correctly gives positive answers), but not decidable.