Solving problems by searching

In which we see how an agent can find a sequence of actions that achieves its goals when no single action will do

Piotr Fulmański

Wydział Matematyki i Informatyki,
Uniwersytet Łódzki, Polska

March 4, 2010

Spis treści

Basic definitions

Our discussion of problem solving begins with precise definitions of **problems** and their **solutions**. We give several examples to illustrate these definitions.

We then describe several general-purpose search algorithms that can be used to solve these problems.

Uninformed vs. informed search algorithms

- We will see several **uninformed** search algorithms – algorithms that are given no information about the problem other than its definition. Although some of these algorithms can solve any solvable problem, none of them can do so efficiently.
- **Informed** search algorithms, on the other hand, can do quite well given some guidance on where to look for solutions.

Goal formulation

Goal formulation, based on the current situation and the agent's performance measure, is the first step in problem solving. We will consider a goal to be a set of world states – exactly those states in which the goal is satisfied. **The agent's task is to find out how to act, now and in the future, so that it reaches a goal state.** Before it can do this, it needs to decide what sort of actions and states it should consider.

Problem formulation

Problem formulation is the process of deciding what **actions** and **states** to consider, given a goal.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.
- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.
- We will assume the **environment is known**, so the agent knows which states are reached by each action.
- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.

- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.

- We will assume the **environment is known**, so the agent knows which states are reached by each action.

- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.
- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.
- We will assume the **environment is known**, so the agent knows which states are reached by each action.
- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.
- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.
- We will assume the **environment is known**, so the agent knows which states are reached by each action.
- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.

- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.

- We will assume the **environment is known**, so the agent knows which states are reached by each action.

- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

Problem formulation

- For now, we assume that the environment is **observable**, so the agent always knows the current state.
- We also assume the environment is **discrete**, so at any given state there are only finitely many actions to choose from.
- We will assume the **environment is known**, so the agent knows which states are reached by each action.
- Finally, we assume that the **environment is deterministic**, so each action has exactly one outcome.

Under these assumptions, the solution to any problem is a fixed sequence of actions.

formulate–state–execute

- The process of looking for a sequence of actions that reaches the goal is called **search**.

- A search algorithm takes a problem as input and returns a **solution** in the form of action sequence.

- Once a solution is found, the actions it recommends can be carried out. This is called the **execution** phase.

Thus, we have a simple **formulate–state–execute** design for the agent.

formulate–state–execute

- The process of looking for a sequence of actions that reaches the goal is called **search**.
- A search algorithm takes a problem as input and returns a **solution** in the form of action sequence.
- Once a solution is found, the actions it recommends can be carried out. This is called the **execution** phase.

Thus, we have a simple **formulate–state–execute** design for the agent.

formulate–state–execute

- The process of looking for a sequence of actions that reaches the goal is called **search**.
- A search algorithm takes a problem as input and returns a **solution** in the form of action sequence.
- Once a solution is found, the actions it recommends can be carried out. This is called the **execution** phase.

Thus, we have a simple **formulate–state–execute** design for the agent.

formulate–state–execute

- The process of looking for a sequence of actions that reaches the goal is called **search**.
- A search algorithm takes a problem as input and returns a **solution** in the form of action sequence.
- Once a solution is found, the actions it recommends can be carried out. This is called the **execution** phase.

Thus, we have a simple **formulate–state–execute** design for the agent.

formulate–state–execute

- The process of looking for a sequence of actions that reaches the goal is called **search**.
- A search algorithm takes a problem as input and returns a **solution** in the form of action sequence.
- Once a solution is found, the actions it recommends can be carried out. This is called the **execution** phase.

Thus, we have a simple **formulate–state–execute** design for the agent.

Goal after goal

After formulating a goal and a problem to solve, the agent calls a search procedure to solve it. It then uses the solution to guide its actions, doing whatever the solution recommends as the next thing to do – typically, the first action of the sequence – and then removing that step from the sequence. Once the solution has been executed, the agent will formulate a new goal.

Simple problem solving agent (SPSA)

```
function SPSA(percept) return an action
{
  state:=UPDATE_STATE(state,percept)
  if (seq is empty) then
  {
    goal := FORMULATE_GOAL(state)
    problem := FORMULATE_PROBLEM(state,goal)
    seq := SEARCH(problem)

    if (seq = failure) then
      return a null action
  }
  action := FIRST(seq)
  seq := REST(seq)
  return action
}
```

open-loop

Notice that while the agent is executing the solution sequence it ignores its percepts when choosing an action because it knows in advance what they will be. An agent that carries out its plans with its eyes closed, so to speak, must be quite certain of what is going on. Control theorists call this an **open-loop** system, because ignoring the percepts breaks the loop between agent and environment.

Problem

A **problem** can be defined formally by five components

- The **initial state** that the agent starts in.
- A description of the possible **actions** available to the agent. Given a particular state s function ACTION(s) returns the set of actions that can be executed in s. We say that each of these actions is **applicable** in s.
- A description of what each action does; the formal name for this is the **transition model**, specified by a function RESULT(s,a) that returns the state that results from doing action a in state s. We also use the term **successor** to refer to any state reachable from a given state by a single action.

Together, the initial state, actions, and transition model implicitly define the **state space** of the problem – the set of all states reachable from the initial state by any sequence of actions. The state space forms the directed network or graph in which the nodes are states and the links between nodes are actions. A **path** in the state space is a sequence of states connected by a sequence of actions.

Problem

A **problem** can be defined formally by five components

- The **initial state** that the agent starts in.
- A description of the possible **actions** available to the agent. Given a particular state s function ACTION(s) returns the set of actions that can be executed in s. We say that each of these actions is **applicable** in s.
- A description of what each action does; the formal name for this is the **transition model**, specified by a function RESULT(s,a) that returns the state that results from doing action a in state s. We also use the term **successor** to refer to any state reachable from a given state by a single action.

Together, the initial state, actions, and transition model implicitly define the **state space** of the problem – the set of all states reachable from the initial state by any sequence of actions. The state space forms the directed network or graph in which the nodes are states and the links between nodes are actions. A **path** in the state space is a sequence of states connected by a sequence of actions.

Problem

A **problem** can be defined formally by five components

- The **initial state** that the agent starts in.
- A description of the possible **actions** available to the agent. Given a particular state s function ACTION(s) returns the set of actions that can be executed in s. We say that each of these actions is **applicable** in s.
- A description of what each action does; the formal name for this is the **transition model**, specified by a function RESULT(s,a) that returns the state that results from doing action a in state s. We also use the term **successor** to refer to any state reachable from a given state by a single action.

Together, the initial state, actions, and transition model implicitly define the **state space** of the problem – the set of all states reachable from the initial state by any sequence of actions. The state space forms the directed network or graph in which the nodes are states and the links between nodes are actions. A **path** in the state space is a sequence of states connected by a sequence of actions.

Problem

A **problem** can be defined formally by five components

- The **initial state** that the agent starts in.
- A description of the possible **actions** available to the agent. Given a particular state s function ACTION(s) returns the set of actions that can be executed in s. We say that each of these actions is **applicable** in s.
- A description of what each action does; the formal name for this is the **transition model**, specified by a function RESULT(s, a) that returns the state that results from doing action a in state s. We also use the term **successor** to refer to any state reachable from a given state by a single action.

Together, the initial state, actions, and transition model implicitly define the **state space** of the problem – the set of all states reachable from the initial state by any sequence of actions. The state space forms the directed network or graph in which the nodes are states and the links between nodes are actions. A **path** in the state space is a sequence of states connected by a sequence of actions.

Problem

A **problem** can be defined formally by five components

- The **initial state** that the agent starts in.
- A description of the possible **actions** available to the agent. Given a particular state s function ACTION(s) returns the set of actions that can be executed in s. We say that each of these actions is **applicable** in s.
- A description of what each action does; the formal name for this is the **transition model**, specified by a function RESULT(s,a) that returns the state that results from doing action a in state s. We also use the term **successor** to refer to any state reachable from a given state by a single action.

Together, the initial state, actions, and transition model implicitly define the **state space** of the problem – the set of all states reachable from the initial state by any sequence of actions. The state space forms the directed network or graph in which the nodes are states and the links between nodes are actions. A **path** in the state space is a sequence of states connected by a sequence of actions.

Problem

- The **goal test**, which determines whether a given state is a goal state. Sometimes there is an explicit set of possible goal states, and the test simply checks whether the given state is one of them. Sometimes the goal is specified by an abstract property rather than an explicitly enumerated set of states. For example, in chess, the goal is to reach a state called „checkmate", where the opponent's king is under attack and can't escape.

- A **path cost** function that assigns a numeric cost to each path. The problem-solving agent chooses a cost function that reflects its own performance measure.

A **solution** to a problem is an action sequence that leads from the initial state to a goal state. Solution quality is measured by the path cost function, and an **optimal solution** has the lowest path cost among all solutions.

Problem

- The **goal test**, which determines whether a given state is a goal state. Sometimes there is an explicit set of possible goal states, and the test simply checks whether the given state is one of them. Sometimes the goal is specified by an abstract property rather than an explicitly enumerated set of states. For example, in chess, the goal is to reach a state called „checkmate", where the opponent's king is under attack and can't escape.

- A **path cost** function that assigns a numeric cost to each path. The problem-solving agent chooses a cost function that reflects its own performance measure.

A **solution** to a problem is an action sequence that leads from the initial state to a goal state. Solution quality is measured by the path cost function, and an **optimal solution** has the lowest path cost among all solutions.

Problem

- The **goal test**, which determines whether a given state is a goal state. Sometimes there is an explicit set of possible goal states, and the test simply checks whether the given state is one of them. Sometimes the goal is specified by an abstract property rather than an explicitly enumerated set of states. For example, in chess, the goal is to reach a state called „checkmate", where the opponent's king is under attack and can't escape.

- A **path cost** function that assigns a numeric cost to each path. The problem-solving agent chooses a cost function that reflects its own performance measure.

A **solution** to a problem is an action sequence that leads from the initial state to a goal state. Solution quality is measured by the path cost function, and an **optimal solution** has the lowest path cost among all solutions.

Problem

- The **goal test**, which determines whether a given state is a goal state. Sometimes there is an explicit set of possible goal states, and the test simply checks whether the given state is one of them. Sometimes the goal is specified by an abstract property rather than an explicitly enumerated set of states. For example, in chess, the goal is to reach a state called „checkmate", where the opponent's king is under attack and can't escape.

- A **path cost** function that assigns a numeric cost to each path. The problem-solving agent chooses a cost function that reflects its own performance measure.

A **solution** to a problem is an action sequence that leads from the initial state to a goal state. Solution quality is measured by the path cost function, and an **optimal solution** has the lowest path cost among all solutions.

Abstraction

Abstraction

Every time we propose a formulation of the problem, even it seems reasonable, it is still a **model** – an abstract mathematical description – and not the real thing. Among others we have to leave out of our state description all these considerations which are irrelevent to the problem. The process of removing detail from a representation is called **abstraction**. In addition to abstracting the state description, we must abstract the actions themselves.

Validity and usefulness

- The abstraction is **valid** if we can expand any abstract solution into a solution in the more detailed world.
- The abstraction is **useful** if carrying out each of the action in the solution is easier than the oryginal problem.

The choice of a good abstraction thus involves removing as much detail as possible while retaining validity and ensuring that the abstract actions are easy to carry out. Were it not for the ability to construct useful abstractions, intelligent agents would be completely swamped by the real world.

Validity and usefulness

- The abstraction is **valid** if we can expand any abstract solution into a solution in the more detailed world.
- The abstraction is **useful** if carrying out each of the action in the solution is easier than the oryginal problem.

The choice of a good abstraction thus involves removing as much detail as possible while retaining validity and ensuring that the abstract actions are easy to carry out. Were it not for the ability to construct useful abstractions, intelligent agents would be completely swamped by the real world.

Validity and usefulness

- The abstraction is **valid** if we can expand any abstract solution into a solution in the more detailed world.
- The abstraction is **useful** if carrying out each of the action in the solution is easier than the oryginal problem.

The choice of a good abstraction thus involves removing as much detail as possible while retaining validity and ensuring that the abstract actions are easy to carry out. Were it not for the ability to construct useful abstractions, intelligent agents would be completely swamped by the real world.

Validity and usefulness

- The abstraction is **valid** if we can expand any abstract solution into a solution in the more detailed world.
- The abstraction is **useful** if carrying out each of the action in the solution is easier than the oryginal problem.

The choice of a good abstraction thus involves removing as much detail as possible while retaining validity and ensuring that the abstract actions are easy to carry out. Were it not for the ability to construct useful abstractions, intelligent agents would be completely swamped by the real world.

Toy problems

- vacuum cleaner
- 8-puzzle
- 8-queens problem

state, initial state, actions, transition model, goal test, path cost

Toy problems

- vacuum cleaner
- 8-puzzle
- 8-queens problem

state, initial state, actions, transition model, goal test, path cost

Toy problems

- vacuum cleaner
- 8-puzzle
- 8-queens problem

state, initial state, actions, transition model, goal test, path cost

Toy problems

- vacuum cleaner
- 8-puzzle
- 8-queens problem

state, initial state, actions, transition model, goal test, path cost

Real-world problems

- route-finding problem
- traveling salesperson problem
- automatic assembly sequencing

state, initial state, actions, transition model, goal test, path cost

Real-world problems
- route-finding problem
- traveling salesperson problem
- automatic assembly sequencing

state, initial state, actions, transition model, goal test, path cost

Real-world problems

- route-finding problem
- traveling salesperson problem
- automatic assembly sequencing

state, initial state, actions, transition model, goal test, path cost

Real-world problems

- route-finding problem
- traveling salesperson problem
- automatic assembly sequencing

state, initial state, actions, transition model, goal test, path cost

Search tree

Having formulated some problems, we now need to solve them. A solution is an action sequence, so search algorithms work by considering various possible action sequences. The possible action sequences starting at the initial state from a **search tree** with the initial state at the root; the branches are actions and the **nodes** correspond to states in the state space of the problem.

Search tree

The root node of he tree corresponds to the initial state. Then we need to consider taking various actions. We do this by **expanding** the current state; that is applaying each legal action to the current state, thereby **generating** a new set of states. We add new **branches** from the **parent node** leading to new **child nodes**. Node with no children in the tree is a **leaf node**. The set of all leaf nodes available for expansion at any given point is called the **frontier**.

After that we must choose which of these new possibilities to consider further.

The essence of search

This is the essence of search – following up one option now and putting the others aside for leater, in case the first choice does not lead to a solution.

General tree-search algorithm

Now we present the general tree-search algorithm. Search algorithms all share this basic structure; they vary primarily according to how they choose which state to expand next – the so-called **search strategy**.

TREE-SEARCH

```
function TREE-SEARCH(problem) return a solution, or failure
{
  initialize the frontier using the initial state of problem
  loop
  {
    if (the frontier is empty) than return failure
    choose a leaf node and remove it from the frontier
    if (the node contains a gola state) then
      return the corresponding solution
    expand the chosen node, adding the resulting nodes
      to the frontier
  }
}
```

Redundant paths

Redundant paths exists whenever there is more than one way to get from one state to another. Sometimes redundant paths leads to infinite loop. The way to avoid exploring redundant paths is to remember where one has been. To do this we augment the TREE-SEARCH algorithm with a data structure called the **explored set**, which remembers every expanded node. Newly generated nodes that match previously generated nodes – ones in the explored set or the frontier – can be discarded instead of being added to the frontier. The new algorithm is called GRAPH-SEARCH.

GRAPH-SEARCH

```
function GRAPH-SEARCH(problem) return a solution, or failure
{
  initialize the frontier using the initial state of problem
  initialize the explored set to be empty
  loop
  {
    if (the frontier is empty) than return failure
    choose a leaf node and remove it from the frontier
    if (the node contains a gola state) then
      return the corresponding solution
    add the node to the explored set
    expand the chosen node, adding the resulting nodes
      to the frontier only if not in the frontier
      or explored set
  }
}
```

Infrastructure for searching algorithms

Infrastructure for searching algorithms

Search algorithms required a data structure to keep track of the search tree that is being constructed. For each node n of the tree, we have a structure that contains four components:

- `n.STATE`: the state in the state space to which the node corresponds;
- `n.PARENT`: the node in the search tree that generated this node;
- `n.ACTION`: the action that was applied to the parent to generate the node;
- `n.PATH-COST`: the cost, traditionally denoted by $g(n)$, of the path from the initial state to the node, as indicated by the parent pointers.

We also use `CHILD-NODE` function which takes as arguments a parent node and an action and returns the resulting child node.

Measuring problem solving performance

We can evaluate an algorithm's performance in four ways

- **completeness**: is the algorithm guaranteed to find a solution when there is one?
- **optimality**: does the strategy find the optimal solution (**optimal solution** is such a solution which has the lowest path cost among all solutions)?
- **time complexity**: how long does it take to find a soltion?
- **space complexity**: how much memory is needed to perform the search?

Measuring problem solving performance

Measuring problem solving performance

We can evaluate an algorithm's performance in four ways

- **completeness**: is the algorithm guaranteed to find a solution when there is one?
- optimality: does the strategy find the optimal solution (**optimal solution** is such a solution which has the lowest path cost among all solutions)?
- time complexity: how long does it take to find a soltion?
- space complexity: how much memory is needed to perform the search?

Measuring problem solving performance

Measuring problem solving performance

We can evaluate an algorithm's performance in four ways

- **completeness**: is the algorithm guaranteed to find a solution when there is one?
- **optimality**: does the strategy find the optimal solution (**optimal solution** is such a solution which has the lowest path cost among all solutions)?
- **time complexity**: how long does it take to find a soltion?
- **space complexity**: how much memory is needed to perform the search?

Measuring problem solving performance

Measuring problem solving performance

We can evaluate an algorithm's performance in four ways

- **completeness**: is the algorithm guaranteed to find a solution when there is one?
- **optimality**: does the strategy find the optimal solution (**optimal solution** is such a solution which has the lowest path cost among all solutions)?
- **time complexity**: how long does it take to find a soltion?
- **space complexity**: how much memory is needed to perform the search?

Measuring problem solving performance

Measuring problem solving performance

We can evaluate an algorithm's performance in four ways

- **completeness**: is the algorithm guaranteed to find a solution when there is one?
- **optimality**: does the strategy find the optimal solution (**optimal solution** is such a solution which has the lowest path cost among all solutions)?
- **time complexity**: how long does it take to find a soltion?
- **space complexity**: how much memory is needed to perform the search?

Uninformed and informed search strategies

Uninformed and informed search strategies

no we cover several strategies that comes under the heading of **uninformed search** (also called **blind search**). The term means that the strategies have no additional information about states beyond that provided in the problem definition. All they can do is generate successors and distinguish a goal state from a non-gola state. All search strategies are distinguished by the order in which nodes are expanded. In contrary, strategies that know whether one non-goal state is „more promising" than other are called **informed search** or **heuristic search** strategies.

Breadth-first search

Breadth-first search is a simple strategy in which the root node is expanded first, than all the successors of the root node are expanded next, than their successors, and so on. In general, all the nodes are expanded at a given depth in the search tree before any nodes at the next level are expanded.

Breadth-first search is an instance of the general graph-search algorithm in which the shallowest unexpanded node is chosen for expansion. This is achieved very simply by using FIFO queue for the frontier. Thus, new nodes (which are always deeper than their parents) go to the back of the queue, and old nodes, which are shallower than the new nodes, get expanded first. There is one slight tweak on the general graph-search algorithm, which is that the goal test is applied to each node when it is generated rather than it is selected for expansion.

BREADTH-FIRST-SEARCH (on a graph)

```
function BREADTH-FIRST-SEARCH(problem) return a solution, or failure
{
  node.STATE := problem.INITIAL_STATE
  node.PATH_COST := 0
  if (problem.GOAL_TEST(node.STATE)=TRUE) than
    return SOLUTION(node)
  frontier := a FIFO queue with node as the only element
  explored := an empty set
  loop{
    if (the frontier is empty) than return failure
    choose a leaf node and remove it from the frontier
    /* chooses the shallowest node in frontier */
    explored.ADD(node.STATE)

    for each action in problem.ACTIONS(node.STATE){
      child := CHILD_NODE(problem,node,action)
      if (child.STATE is not in explored or frontier) then
      {
        if (problem.GOAL_TEST(child.STATE)) then
          return SOLUTION(child)
        frontier.ADD(child)
      }
    }
  }
}
```

Uniform-cost search

When all step cost are equal, breadth-first search is optimal because it always expands the shallowest unexpanded node. By a simple extension, we can find an algorithm that is ooptimal with any step-cost function. Instead of expanding the shallowest node, **uniform-cost search** expands the node n with the lowest path cost $g(n)$. This is done by storing the frontier as a priority queue ordered by g.

UNIFORM-COST-SEARCH (on a graph)

```
function UNIFORM-COST-SEARCH(problem) return a solution, or failure
{
  node.STATE := problem.INITIAL_STATE
  node.PATH_COST := 0
  if (problem.GOAL_TEST(node.STATE)=TRUE) than
    return SOLUTION(node)
  frontier := a priority queue ordered by PATH-COST,
              with node as the only element
  explored := an empty set
  loop{
    if (the frontier is empty) than return failure
    node := POP(frontier)
    /* chooses the lowest-cost node in frontier
       and remove it from the frontier */
    if (problem.GOAL_TEST(node.STATE)) then
      return SOLUTION(node)
    explored.ADD(node.STATE)

    for each action in problem.ACTIONS(node.STATE){
      child := CHILD_NODE(problem,node,action)
      if (child.STATE is not in explored or frontier) then
      {
        frontier.ADD(child)
      }
      else if child.STATE is in frontier with higher PATH-COST then
      {
        replace the frontier node with child
      }
    }
  }
}
```

Uniform-cost search – differences

In addition on the ordering of the queue by path cost there are two other significant differences from breadth-first search. The first is that the goal test is applied to a node when it is selected for expansion (as in the generic graph-search algorithm) rather than when it is first generated. The second difference is that a test is added in case a better path is found to a node currently on the frontier.

Depth-first search

Depth-first search always expands the deepest node in the current frontier of the search tree. The search proceeds immediately to the deepest level of the search tree, where the nodes have no successors. As those nodes are expanded, they are dropped from the frontier, so then the search „backs up" to the next deepest node that still has unexplored successor. The depth-first search algorithm is an instance of the graph-search algorithm which uses LIFO queue. A LIFO queue means that the most recently generated node is chosen for expansion. This must be the deepest unexpanded node because it is one deeper than its parent – which, in turn, was the deepest unexpanded node when it was selected.

Depth-limited search

The embarrassing failure of depth-first search in infinite state spaces can be alleviated by supplying depth-first search with a predetermined dept limit l. That is, nodes at depth l are treated as if they have no successors. This approach is called **depth-limited search**.

Recursive DEPTH-LIMITED-SEARCH

```
function DEPTH-LIMITED-SEARCH(problem) return a solution, or failure/cutoff
  return Recursuve-DLS(MAKE_NODE(problem.INITIAL-STATE),problem,limit)

function Recursive-DLS(node,problem,limit) return a solution, or failure/cutoff
{
  if (problem.GOAL_TEST(node.STATE)=TRUE) than
    return SOLUTION(node)
  else if limit=0 then
    return cutoff
  else
    cutoff_occurred := FALSE
    for each action in problem.ACTIONS(node.STATE)
    {
      child := CHILD_NODE(problem,node,action)
      result := Recursive-DLS(child,problem,limit-1)
      if (result = cutoff)
        then cutoff_occurred := TRUE
      else if (result != failure) then
        return result
    }
    if (cutoff_occurred = TRUE) then
      return cutoff
    else
      return failure
}
```

Iterative deepening depth-first search

Iterative deepening depth-first search

Iterative deepening depth-first search (or iterative deepening search) is a general strategy, often use in combination with depth-first tree search, that finds the best depth limit. It does this by gradually increasing the limit – first 0, then 1, then 2, and so on – until a goal is found.
The iterative deepening combines the benefits of depth-first and breadth-first search.

ITERATIVE-DEEPENING-SEARCH

```
function ITERATIVE-DEEPENING-SEARCH(problem) return a solution, or failure
{
  for depth := 0 to infinity do
  {
    result := DEPTH-LIMITED-SEARCH(problem,depth)
    if (result != cutoff) then return result
  }
}
```

Bidirectional search

The idea behind bidirectional search is to run two simultaneous searches – one forward from the initial state and the other backward from the goal – hoping that the two searches meet in the middle.

Thers is one very important problem. Consider the question of what we mean by „the goal" in searching „backward from the goal". For the 8-puzzle there is just one goal so the backward search is very much like the forward search. If there are several explicitly listed goal states than we can construct a new dummy goal state whose immediate predecessors are all the actual goal states. But if the goal is an abstract description, such as the goal that „no queen attacks another queen" in the n-queens problem, than bidirectional search is difficult to use.

Comparing uninformed search strategies

	Breadth-First	Uniform-Cost	Depth-First	Depth-Limited	Iterative-Deepening	Bidirectional
Complete	Yes	Yes	No	No	Yes	Yes
Time	$O(b^d)$	$O(b^{1+[C^*/\epsilon]})$	$O(b^m)$	$O(b^l)$	$O(b^d)$	$O(b^{d/2})$
Space	$O(b^d)$	$O(b^{1+[C^*/\epsilon]})$	$O(bm)$	$O(bl)$	$O(bd)$	$O(b^{d/2})$
Optimal	Yes	Yes	No	No	Yes	Yes

Informed (heuristic) search strategies

Informed (heuristic) search strategies

Now we show how an informed search strategy –one that uses problem-specific knowledge beyond the definition of the problem itself – can find solution more efficiently than can an uninformed strategy.

A general approach we consider is called **best-first search**. Best-first search is an instance of the general TREE-SEARCH or GRAPH-SEARCH algorithm in which a node is selected for expansion based on an **evaluation function**, $f(n)$. The evaluation function is construed as a cost estimate, so the node with the lowest evaluation is expanded first. The implementation of the best-first graph search is identical to that for uniform-cost search, except for the use of f instead of g to order the priority queue.

Heuristic function

The choice of f determines the search strategy. Most best-first algorithms include as a component of f a **heuristic function**, denoted $h(n)$; $h(n)$ = estimated cost of the cheapest path from the state at node n to a goal state.

Heuristic (from the Greek: „find" or „discover") is an adjective for experience-based techniques that help in problem solving, learning and discovery. A heuristic method is used to rapidly come to a solution that is hoped to be close to the best possible answer, or „optimal solution".

Greedy best-first search

Greedy best-first search tries to expand the node that is closest to the goal, on the grounds that this is likely to lead to a solution quickly. At each step algorithm tries to get as close to the goal as it can. Thus, it evaluates nodes by using just the heuristic function; that is, $f(n) = h(n)$.

A* search: minimizing the total estimated solution cost

A*

The most widely known form of best-first search is called **A* search**. It evaluates nodes y combining $g(n)$, the cost to rach the node, and $h(n)$, he cost to get from the node to the goal:

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

Since $g(n)$ gives the path cost from the start node to node n, and $h(n)$ is the estimated cost of the cheapest path from n to the goal, we have $f(n) =$ estimated cost of cheapest solution through n.

Thus if we are trying to find the cheapest solution, a reasonable thing to try first is the node with the lowest value of $g(n) + h(n)$. It turns out that this strategy is more than just reasonable: provided that the heuristic function $h(n)$ satisfies certain conditions, A* search is both complete and optimal. The algorith ise identical to uniform-cost-search except that A* uses $g + h$ instead of g.