

## Intelligent Agents

**agent**: anything that can be viewed as perceiving its **environment** through **sensors** and acting upon that environment through **actuators**.

Agent behavior is determined by the **agent function** that maps any given percept sequence to an action.

The agent function for an artificial agent will be implemented by an **agent program**.

Slide CS472 – Problem-Solving as Search 1

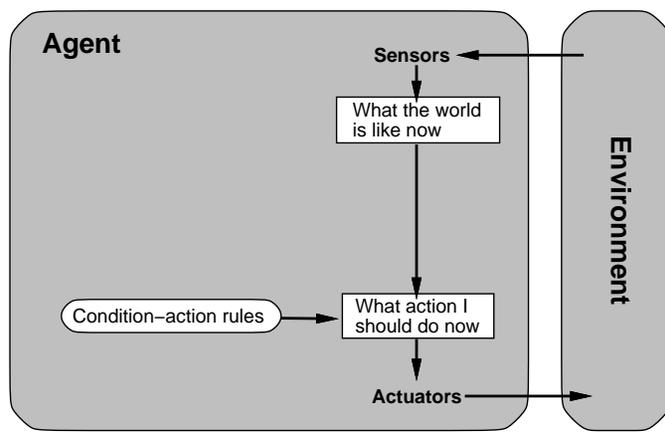
## A “Cornell AI Student” Agent

**Agent requires access to environment through sensors**: visual, aural, touch, etc.

**Available actions**: talk, walk, do arithmetic, do boolean logic, programming skills, etc.

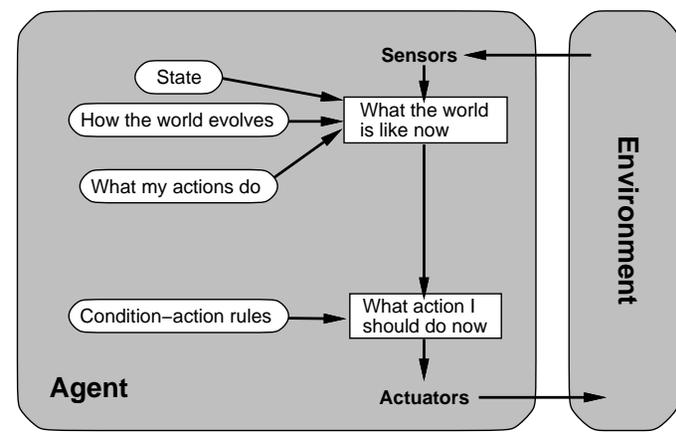
Slide CS472 – Problem-Solving as Search 2

## A Simple Reflex Agent



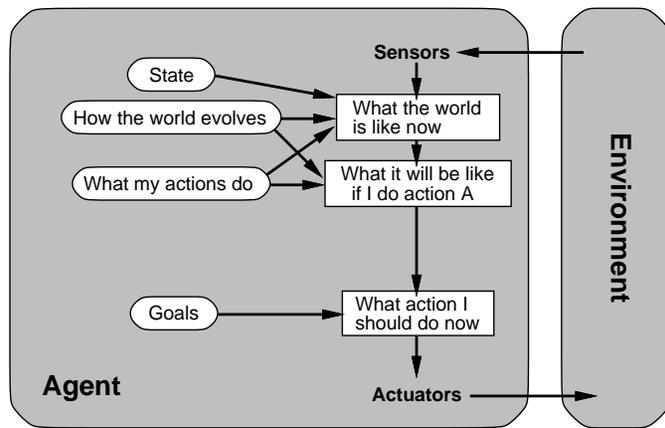
Slide CS472 – Problem-Solving as Search 3

## Agents with Internal State



Slide CS472 – Problem-Solving as Search 4

## Goal-Based Agents



Slide CS472 – Problem-Solving as Search 5

## Problem Solving as Search

### Search is a central topic in AI

- Originated with Newell and Simon's work on problem solving. Famous book: "Human Problem Solving" (1972)
- Automated reasoning is a natural search task
- More recently: Given that almost all AI formalisms (planning, learning, etc.) are NP-complete or worse, some form of search is generally **unavoidable** (no "smarter" algorithm available).

Slide CS472 – Problem-Solving as Search 6

## Defining a Search Problem

**State space** – described by

**initial state** – starting state

**actions** – possible actions available

**successor function; operators** – given a particular state  $x$ , returns a set of  $\langle \text{action}, \text{successor} \rangle$  pairs

Slide CS472 – Problem-Solving as Search 7

A **path** is any sequence of states connected by a sequence of actions.

**Goal test** – determines whether a given state is a goal state.

**Path cost** – function that assigns a cost to a path; relevant if more than one path leads to the goal, and we want the shortest path.

Assumption: cost of a path is the sum of the costs of the individual actions along the path; sum of the **step costs**, which must be non-negative.

Slide CS472 – Problem-Solving as Search 8

## The 8-Puzzle

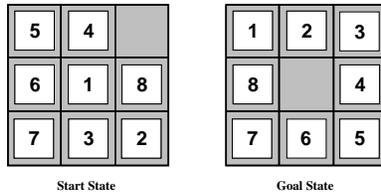
**States:**

**Initial state:**

**Goal test:**

**Successor function:**

**Path cost:**



Slide CS472 – Problem-Solving as Search 9

## Cryptarithmic

$$\begin{array}{r} \text{SEND} \\ + \text{MORE} \\ \hline \text{MONEY} \end{array}$$

Find substitution of digits for letters such that the resulting sum is arithmetically correct.

Each letter must stand for a different digit.

Slide CS472 – Problem-Solving as Search 10

## Cryptarithmic, cont.

**States:** an 8-tuple indicating a (partial) assignment of digits to letters.

**Successor function:** represents the act of assigning digits to letters.

**Goal test:** all letters have been assigned digits and sum is correct.

**Path cost:** ...all solutions are equally valid; step cost = 0.

Slide CS472 – Problem-Solving as Search 11

## Solving a Search Problem: State Space Search

**Input:**

- Initial state
- Goal test
- Successor function
- Path cost

**Output:** path from initial state to goal. Solution quality is measured by the past cost. state.

State space is **not** stored in its entirety by the computer.

Slide CS472 – Problem-Solving as Search 12

## Generic Search Algorithm

```
L = make-queue(initial-state)
loop
  node = remove-front(L) (and save in order
    to return as part of path to goal)
  if goal-test(node) = true return(path to goal)
  S = successors(node)
  insert(S,L)
until L is empty
return failure
```

Slide CS472 – Problem-Solving as Search 13

Search procedure defines a search tree

**root node** — initial state

**children of a node** — successor states

**fringe of tree** — L: states not yet expanded

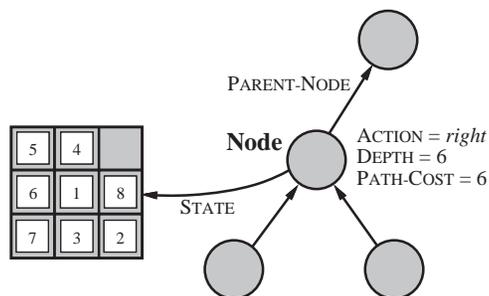
**stack:** Depth-First Search (DFS).

**queue:** Breadth-First Search (BFS).

**Search strategy** — algorithm for deciding which leaf node to expand next.

Slide CS472 – Problem-Solving as Search 14

## Node Data Structure



Slide CS472 – Problem-Solving as Search 15

## Solving the 8-Puzzle

5	4	
6	1	8
7	3	2

Start State

1	2	3
8		4
7	6	5

Goal State

What would the search tree look like after the start state was expanded?

Slide CS472 – Problem-Solving as Search 16

### Evaluating a Search Strategy

**Completeness:** is the strategy guaranteed to find a solution when there is one?

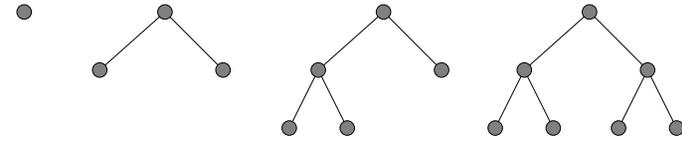
**Time Complexity:** how long does it take to find a solution?

**Space Complexity:** how much memory does it need?

**Optimality:** does the strategy find the highest-quality solution when there are several different solutions?

Slide CS472 – Problem-Solving as Search 17

### Uninformed search: BFS



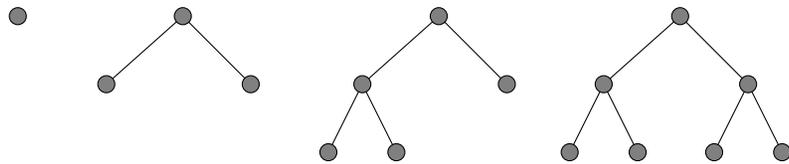
Consider paths of length 1, then of length 2, then of length 3, then of length 4,....

Slide CS472 – Problem-Solving as Search 18

### Time and Memory Requirements for BFS – $O(b^{d+1})$

Let  $b$  = branching factor,  $d$  = solution depth, then the maximum number of nodes *generated* is:

$$b + b^2 + \dots + b^d + (b^{d+1} - b)$$



Slide CS472 – Problem-Solving as Search 19

### Time and Memory Requirements for BFS – $O(b^{d+1})$

$b = 10$

10000 nodes/second

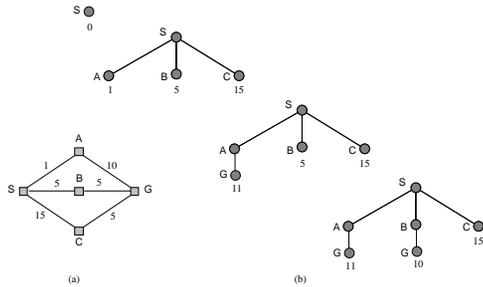
each node requires 1000 bytes of storage

depth	nodes	time	memory
2	1100	.11 sec	1 meg
4	111,100	11 sec	106 meg
6	$10^7$	19 min	10 gig
8	$10^9$	31 hrs	1 tera
10	$10^{11}$	129 days	101 tera
12	$10^{13}$	35 yrs	10 peta
14	$10^{15}$	3523 yrs	1 exa

Slide CS472 – Problem-Solving as Search 20

### Uniform-cost Search

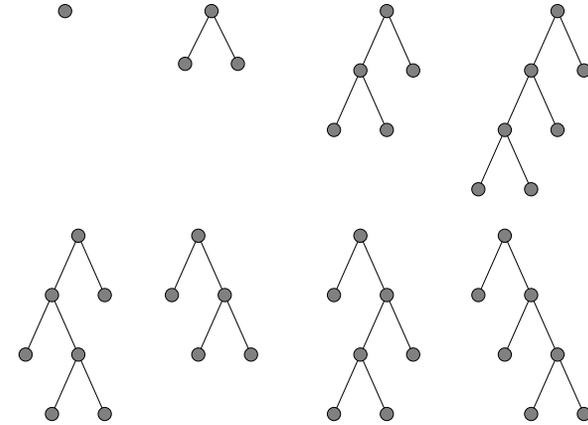
Use BFS, but always expand the lowest-cost node on the fringe as measured by path cost  $g(n)$ .



Requirement:  $g(\text{Successor}(n)) \geq g(n)$

Slide CS472 – Problem-Solving as Search 21

### Uninformed search: DFS



Slide CS472 – Problem-Solving as Search 22

### DFS vs. BFS

	Complete?	Optimal?	Time	Space
BFS	YES	YES	$O(b^{d+1})$	$O(b^{d+1})$
DFS	finite depth	NO	$O(b^m)$	$O(bm)$

$m$  is maximum depth

#### Time

$m = d$  — DFS typically wins

$m > d$  — BFS might win

$m$  is infinite — BFS probably will do better

#### Space

DFS almost always beats BFS

Slide CS472 – Problem-Solving as Search 23

### Which search should I use?

Depends on the problem.

If there may be infinite paths, then depth-first is probably bad. If goal is at a known depth, then depth-first is good.

If there is a large (possibly infinite) branching factor, then breadth-first is probably bad.

(Could try **nondeterministic** search. Expand an open node at random.)

Slide CS472 – Problem-Solving as Search 24

## Iterative Deepening [Korf 1985]

### Idea:

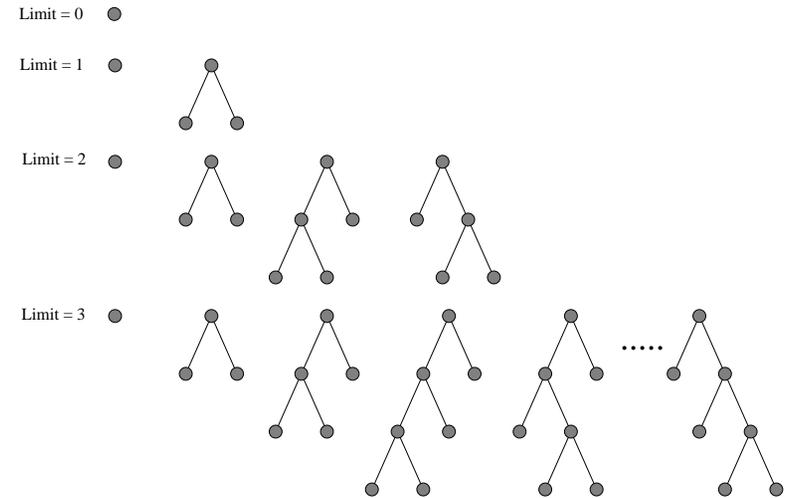
Use an *artificial* depth cutoff,  $c$ .

If search to depth  $c$  succeeds, we're done. If not, increase  $c$  by 1 and start over.

Each iteration searches using DFS.

Slide CS472 – Problem-Solving as Search 25

## Iterative Deepening



Slide CS472 – Problem-Solving as Search 26

Space requirements? Same as DFS. Each search is just a DFS.

Time requirements. Would seem very expensive!! **BUT** not much different from single BFS or DFS to depth  $d$ .

**Reason:** Almost all work is in the final couple of layers.  
E.g., binary tree: 1/2 of the nodes are in the bottom layer.  
With  $b=10$ , 9/10th of the nodes in final layer!

So, repeated runs are on much smaller trees (i.e., exponentially smaller).

Slide CS472 – Problem-Solving as Search 27

### Example:

$b=10, d=5$ , the number of nodes generated in a BFS:  
 $b + b^2 + \dots + b^d + b^{d+1} - b =$   
 $10 + 100 + 1000 + 10,000 + 100,000 + 999,990 = 1,111,100$

For IDS:

$(d)b + (d-1)b^2 + \dots + (1)b^d =$   
 $50 + 400 + 3,000 + 20,000 + 100,000 = 123,450$

Cost of repeating the work at shallow depths is not prohibitive.

Slide CS472 – Problem-Solving as Search 28

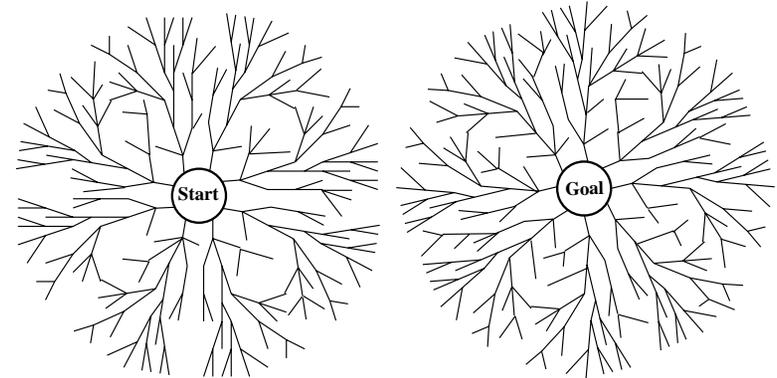
### Cost of Iterative Deepening

space:  $O(bd)$  as in DFS, time:  $O(b^d)$

b	ratio of IDS to DFS
2	3
3	2
5	1.5
10	1.2
25	1.08
100	1.02

Slide CS472 – Problem-Solving as Search 29

### Bidirectional Search



Slide CS472 – Problem-Solving as Search 30

- Search forward from the start state and backward from the goal state simultaneously and stop when the two searches meet the middle.
- If branching factor =  $b$  from both directions, and solution exists at depth  $d$ , then need only  $O(2b^{d/2}) = O(b^{d/2})$  steps.
- Example  $b = 10$ ,  $d = 6$  then BFS needs 1,111,111 nodes and bidirectional search needs only 2,222.
  - What does it mean to search backwards from a goal?
  - What if there is more than one goal state? (chess).

Slide CS472 – Problem-Solving as Search 31

### Comparing Search Strategies

Criterion	Breadth-First	Uniform-Cost	Depth-First	Iterative Deepening	Bidirectional (if applicable)
Time	$b^{d+1}$	$b^d$	$b^m$	$b^d$	$b^{d/2}$
Space	$b^{d+1}$	$b^d$	$bm$	$bd$	$b^{d/2}$
Optimal?	yes	yes	no	yes	yes
Complete?	yes	yes	no	yes	yes

\*\*\*Note that many of the “yes’s” above have caveats, which we discussed when covering each of the algorithms.

Slide CS472 – Problem-Solving as Search 32