Problem solving and search
by Stuart Russell

modified by Jacek Malec for LTH lectures
January 24th, 2014

Chapter 3 of AIMA

Outline

◇ Problem-solving agents

◇ Problem types

◇ Problem formulation

◇ Example problems

◇ Basic (uninformed) search algorithms

◇ Informed 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 ← UPDATE-STATE(state, percept)
    if seq is empty then
        goal ← FORMULATE-GOAL(state)
        problem ← FORMULATE-PROBLEM(state, goal)
        seq ← SEARCH( problem)
    action ← RECOMMENDATION(seq, state)
    seq ← REMAINDER(seq, state)
    return action
```

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

Example: Blocket

©Jorchr / Wikimedia Commons / CC-BY-SA-3.0 / GFDL

Example: Blocket

Service robot Odin, delivering drugs to divisions. Currently in the Pharmacy.
There is a drug order from Intensive Care Unit.

Formulate goal:

      be in Intensive Care Unit
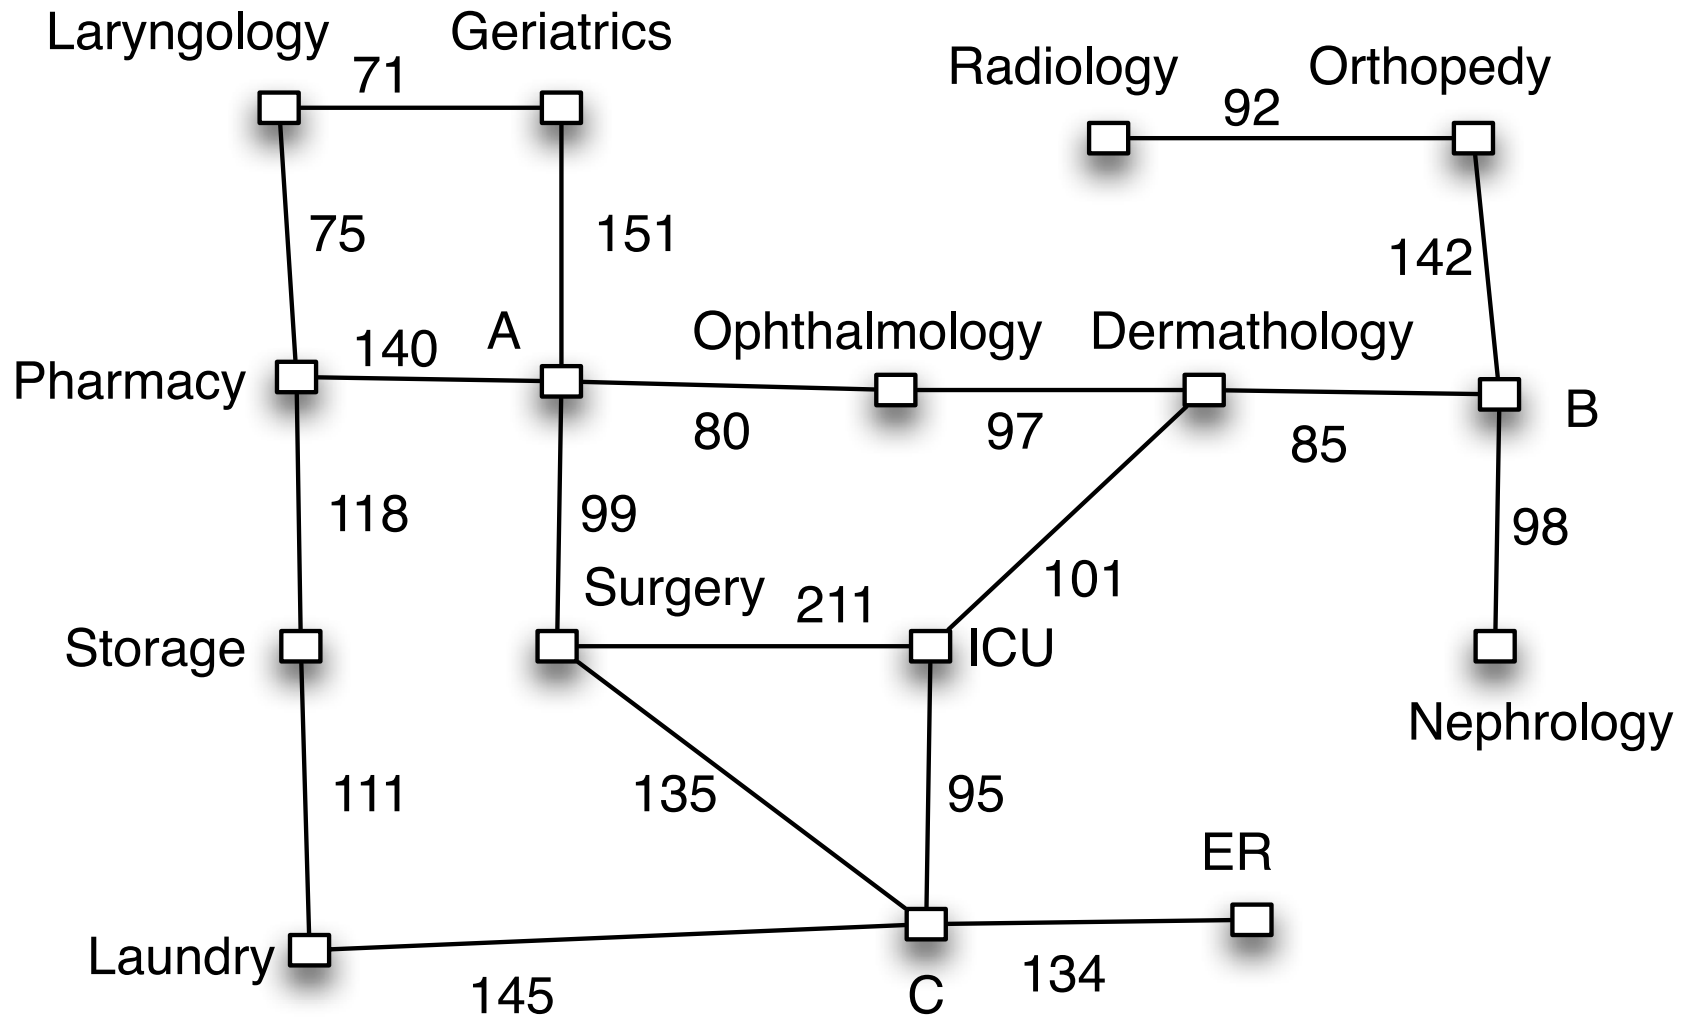
Formulate problem:

      states: various locations
      actions: drive between locations

Find solution:

      sequence of locations, e.g., Pharmacy, Elevator A, Surgery, ICU

Example: Blocket

Laryngology — 71 — Geriatrics

Radiology — 92 — Orthopedy

Laryngology–Pharmacy: 75

Geriatrics–A: 151

Orthopedy–B: 142

Pharmacy — 140 — A

Pharmacy–A: 140

A — 80 — Ophthalmology — 97 — Dermathology — 85 — B

Dermathology–Orthopedy: 142

Pharmacy–Storage: 118

A–Surgery: 99

B–Nephrology: 98

Surgery — 211 — ICU

ICU–Dermathology: 101

Storage–Laundry: 111

Surgery–C: 135

ICU–C: 95

ER

Laundry — 145 — C — 134 — ER

Problem types

Deterministic, fully observable $\implies$ single-state problem
Agent knows exactly which state it will be in; solution is a sequence

Non-observable $\implies$ conformant problem
Agent may have no idea where it is; solution (if any) is a sequence

Nondeterministic and/or partially observable $\implies$ contingency problem
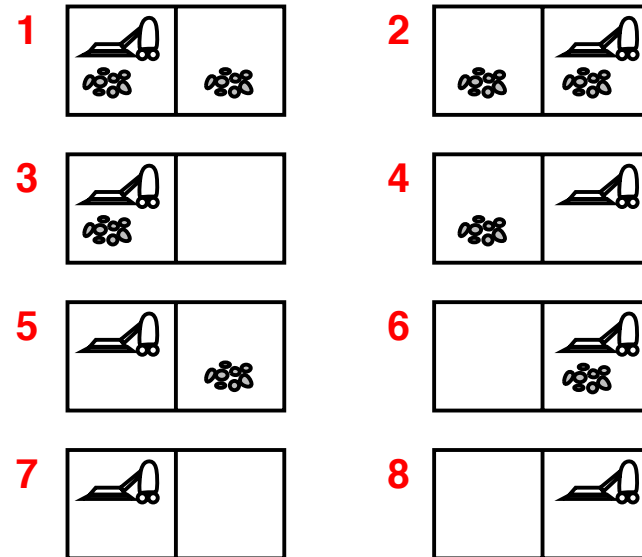percepts provide **new** information about current state
solution is a contingent plan or a policy
often **interleave** search, execution

Unknown state space $\implies$ exploration problem ("online")

Example: vacuum world
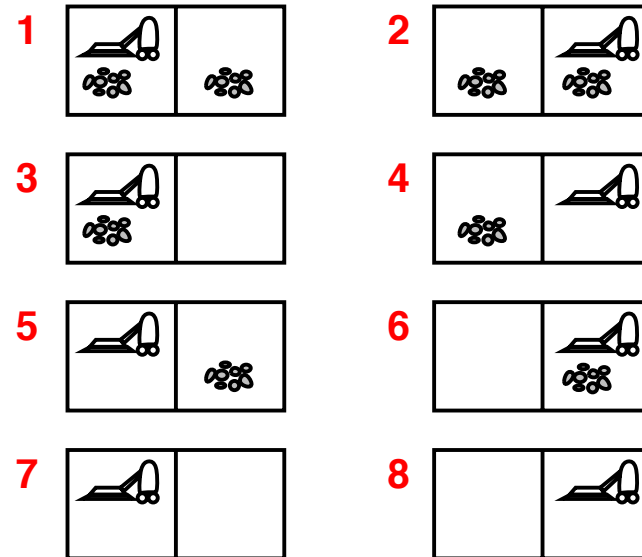
Single-state, start in #5. Solution??

Example: vacuum world

Single-state, start in #5. <u>Solution??</u>
$[Right, Suck]$

Conformant, start in $\{1, 2, 3, 4, 5, 6, 7, 8\}$
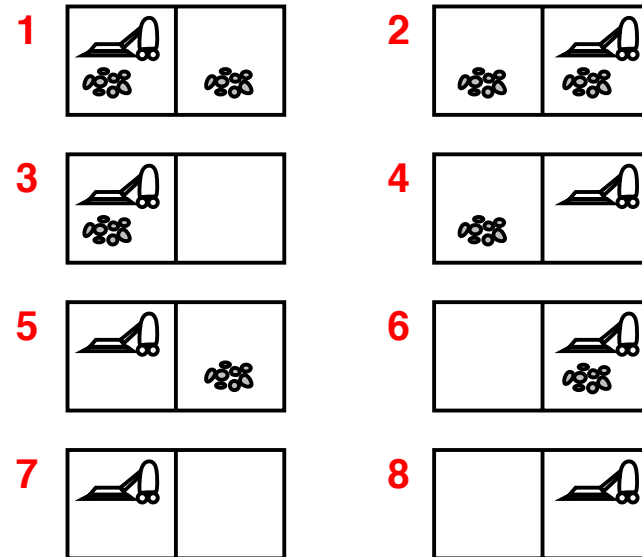e.g., $Right$ goes to $\{2, 4, 6, 8\}$. <u>Solution??</u>

Single-state, start in #5. Solution??
$[Right, Suck]$

Conformant, start in $\{1, 2, 3, 4, 5, 6, 7, 8\}$
e.g., $Right$ goes to $\{2, 4, 6, 8\}$. Solution??
$[Right, Suck, Left, Suck]$

Contingency, start in #5
Murphy's Law: $Suck$ can dirty a clean carpet
Local sensing: dirt, location only.
Solution??

Example: vacuum world

Single-state, start in #5. <u>Solution</u>??
$[Right, Suck]$

Conformant, start in $\{1, 2, 3, 4, 5, 6, 7, 8\}$
e.g., $Right$ goes to $\{2, 4, 6, 8\}$. <u>Solution</u>??
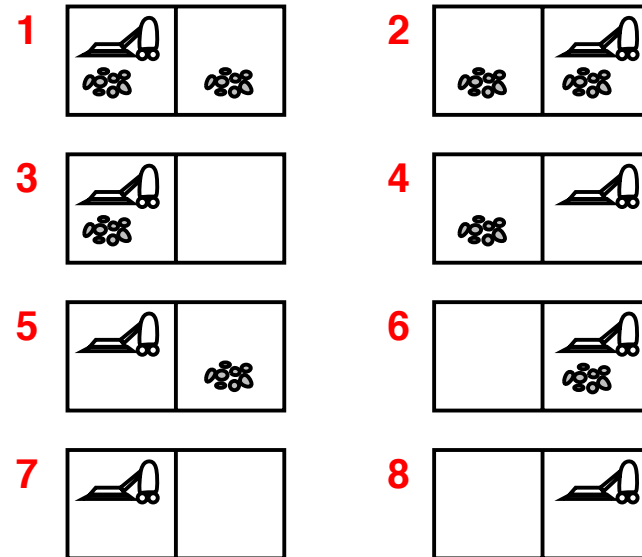$[Right, Suck, Left, Suck]$

Contingency, start in #5
Murphy's Law: $Suck$ can dirty a clean carpet
Local sensing: dirt, location only.
<u>Solution</u>??
$[Right, \mathbf{if}\ dirt\ \mathbf{then}\ Suck]$

Single-state problem formulation

A problem is defined by four items:

initial state   e.g., "at Pharmacy"

successor function $S(x)$ = set of action–state pairs
    e.g., $S(Pharmacy) = \{\langle Pharmacy \to Storage, Storage\rangle, \ldots\}$

goal test, can be
        explicit, e.g., $x =$ "at ICU"
        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 ≥ 0

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

Selecting a state space

Real world is absurdly complex
$\Rightarrow$ state space must be **abstracted** for problem solving

(Abstract) state = set of real states

(Abstract) action = complex combination of real actions
e.g., "Pharmacy $\rightarrow$ Storage" represents a complex set
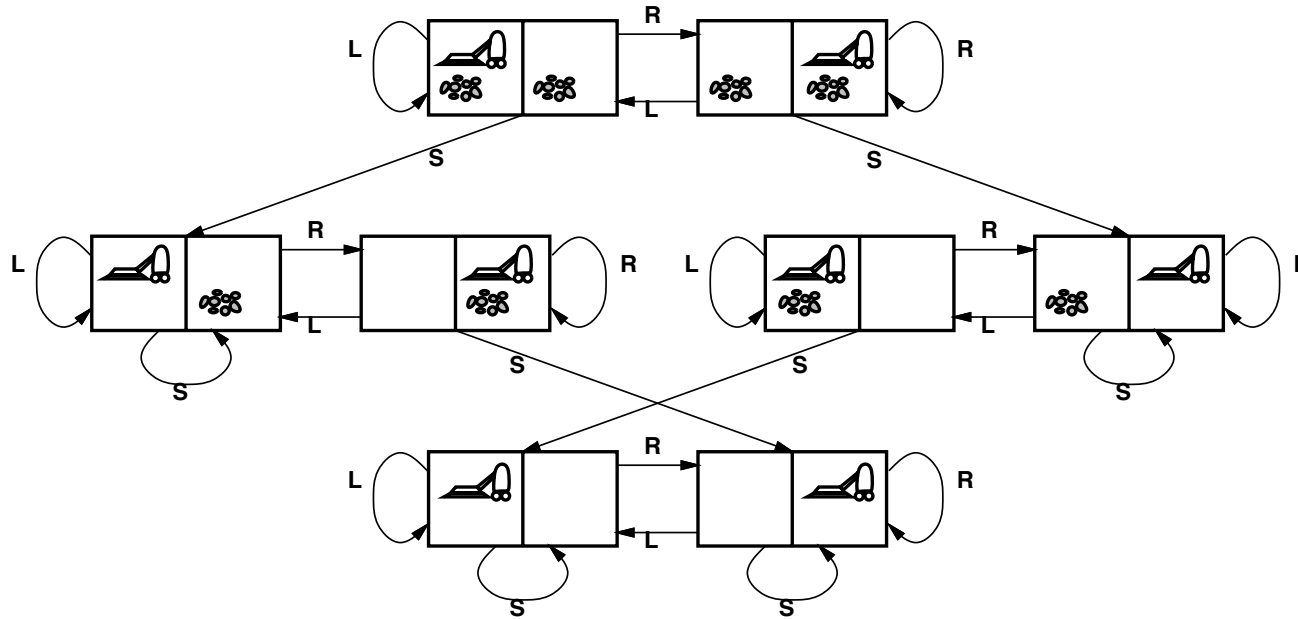of possible routes, detours, rest stops, etc.
For guaranteed realizability, **any** real state "in Pharmacy"
must get to some real state "in Storage"

(Abstract) solution =
set of real paths that are solutions in the real world

Each abstract action should be "easier" than the original problem!

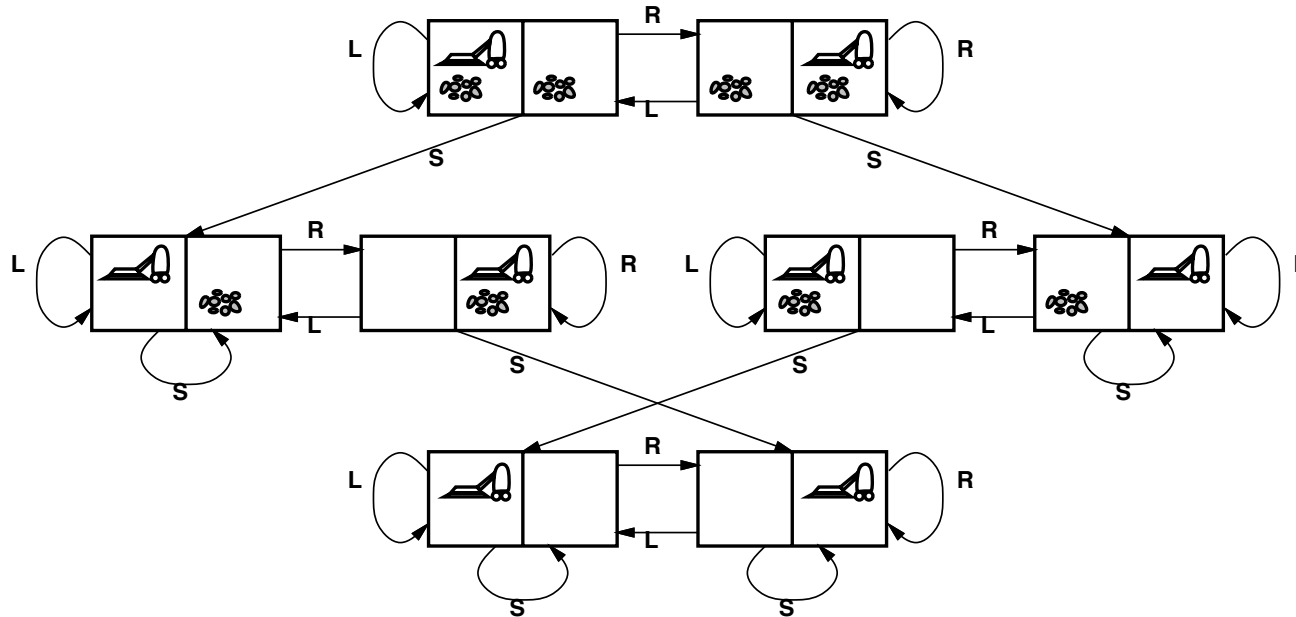# Example: vacuum world state space graph

states??
actions??
goal test??
path cost??

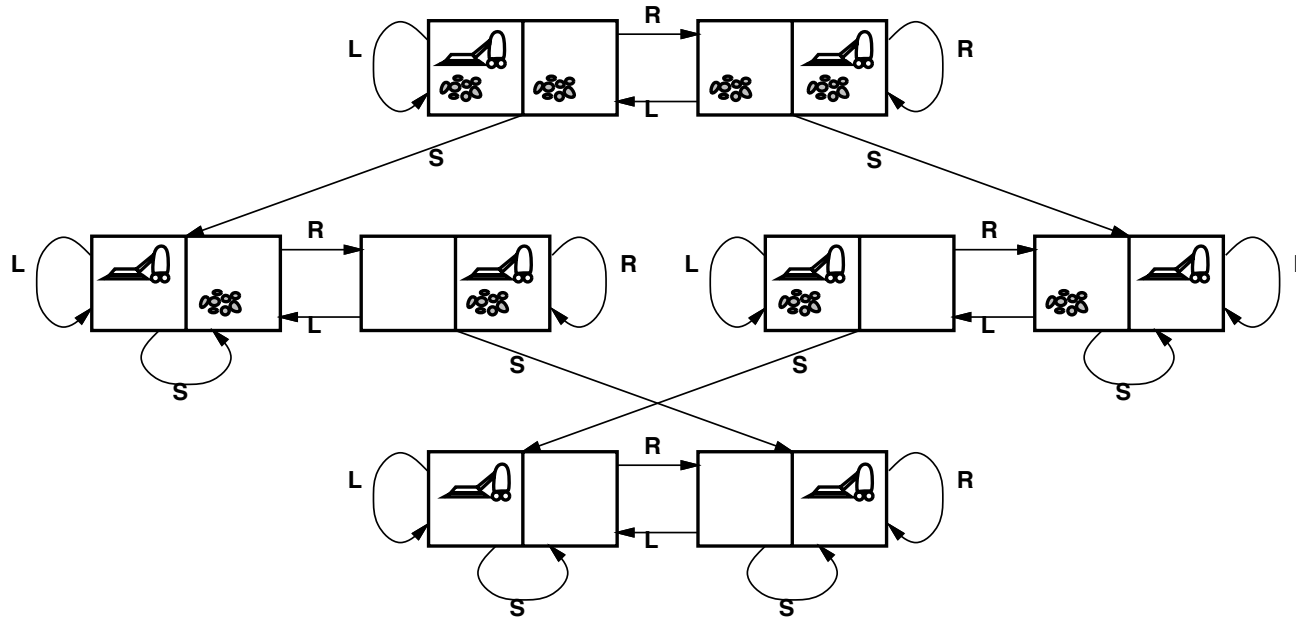# Example: vacuum world state space graph

states??: integer dirt and robot locations (ignore dirt amounts etc.)
actions??
goal test??
path cost??

Example: vacuum world state space graph

states??: integer dirt and robot locations (ignore dirt amounts etc.)
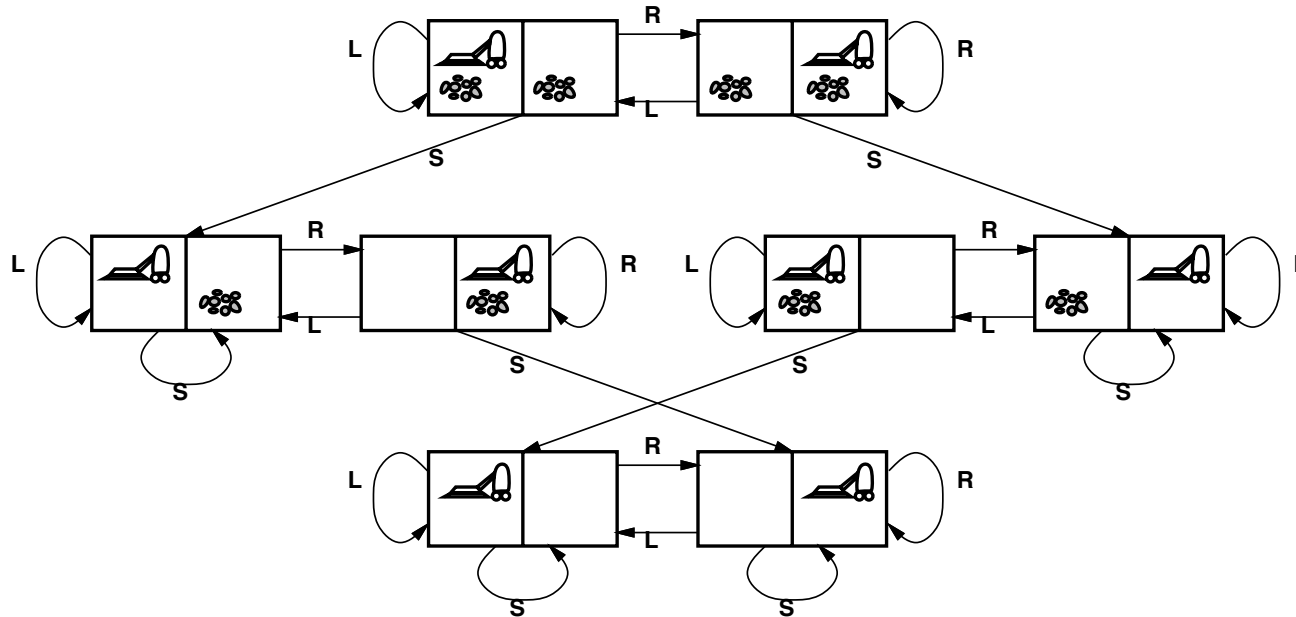
actions??: *Left*, *Right*, *Suck*, *NoOp*

goal test??

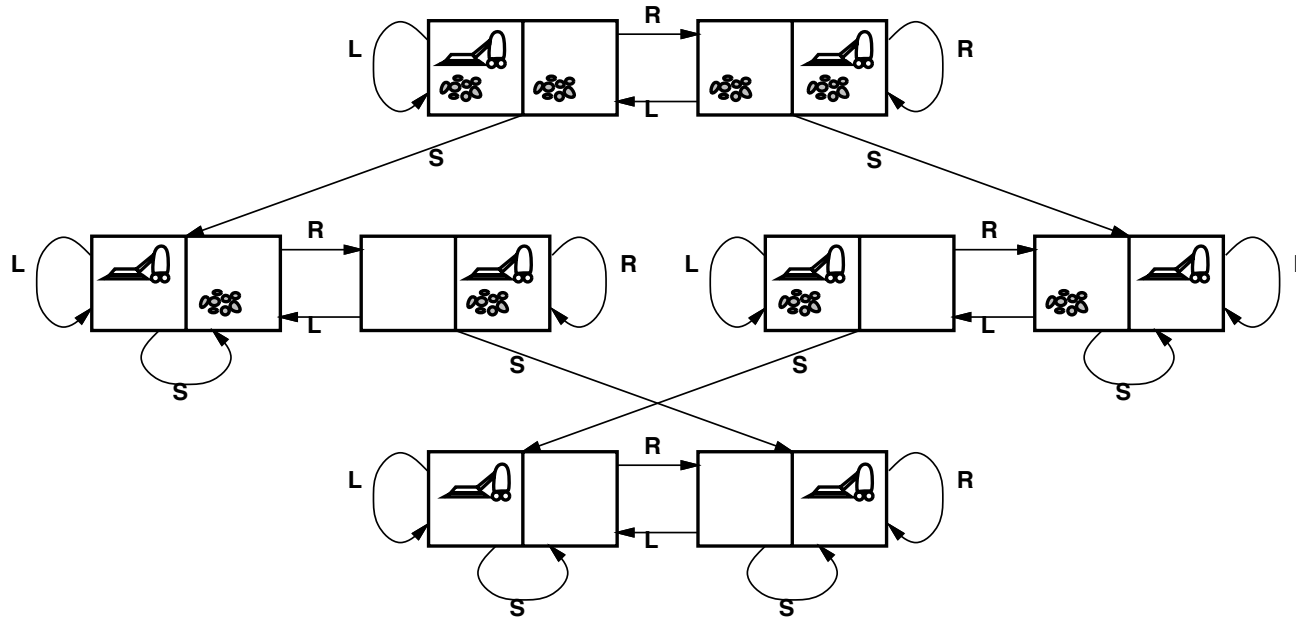path cost??

Example: vacuum world state space graph

<u>states</u>??: integer dirt and robot locations (ignore dirt amounts etc.)

<u>actions</u>??: *Left*, *Right*, *Suck*, *NoOp*

<u>goal test</u>??: no dirt

<u>path cost</u>??

Example: vacuum world state space graph

<u>states</u>??: integer dirt and robot locations (ignore dirt amounts etc.)

<u>actions</u>??: *Left*, *Right*, *Suck*, *NoOp*

<u>goal test</u>??: no dirt

<u>path cost</u>??: 1 per action (0 for *NoOp*)

Example: The 8-puzzle

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

states??
actions??
goal test??
path cost??

Example: The 8-puzzle

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

states??: integer locations of tiles (ignore intermediate positions)

actions??

goal test??

path cost??

Example: The 8-puzzle

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

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

Example: The 8-puzzle

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

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

Example: The 8-puzzle

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

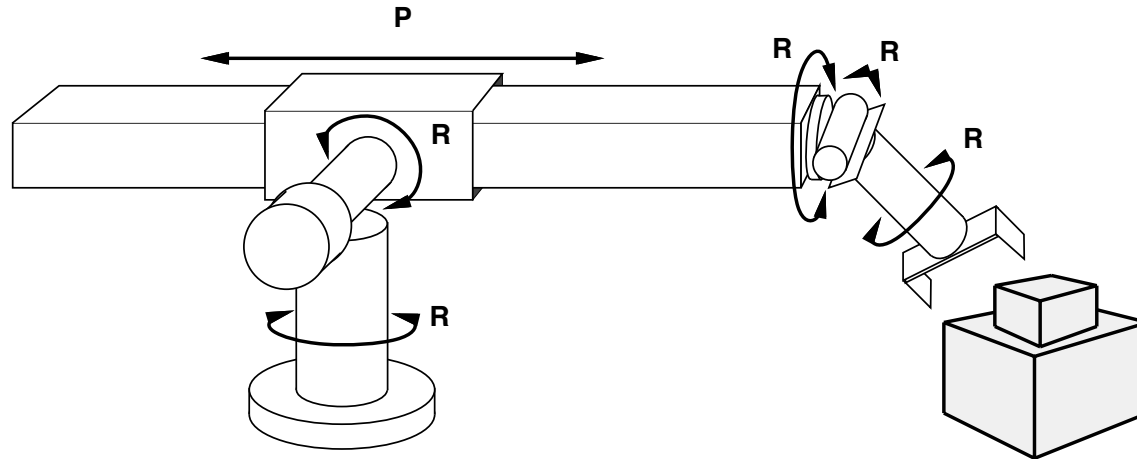states??: integer locations of tiles (ignore intermediate positions)
actions??: move blank left, right, up, down (ignore unjamming etc.)
goal test??: = goal state (given)
path cost??: 1 per move

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

Example: robotic assembly

states??: real-valued coordinates of robot joint angles
      parts of the object to be assembled

actions??: continuous motions of robot joints

goal test??: complete assembly **with no robot included!**

path cost??: time to execute

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

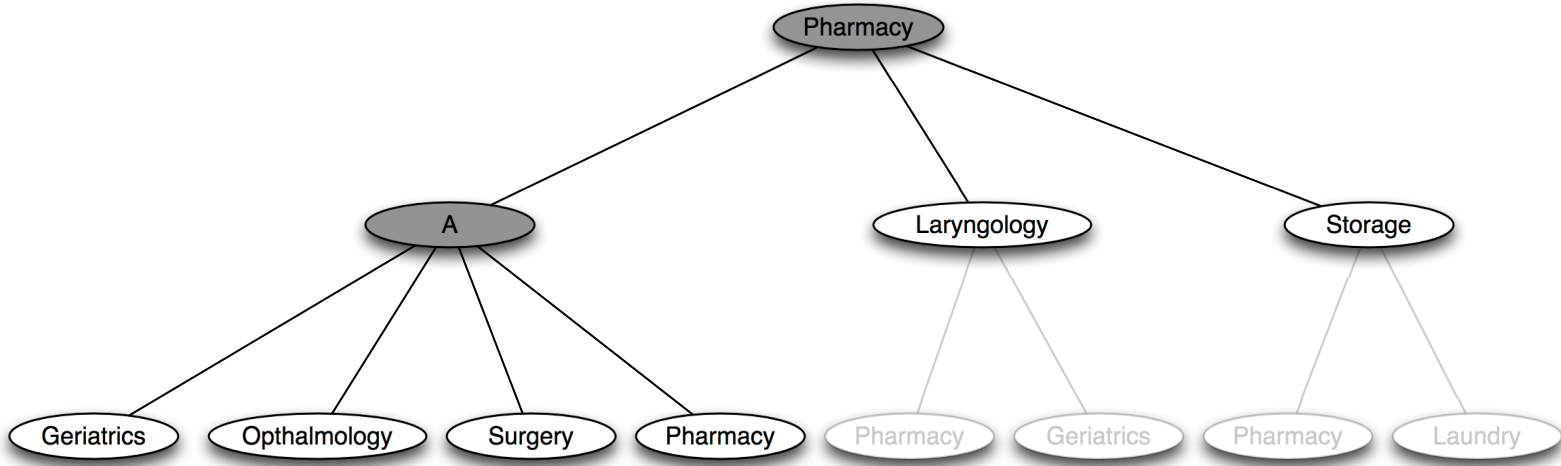
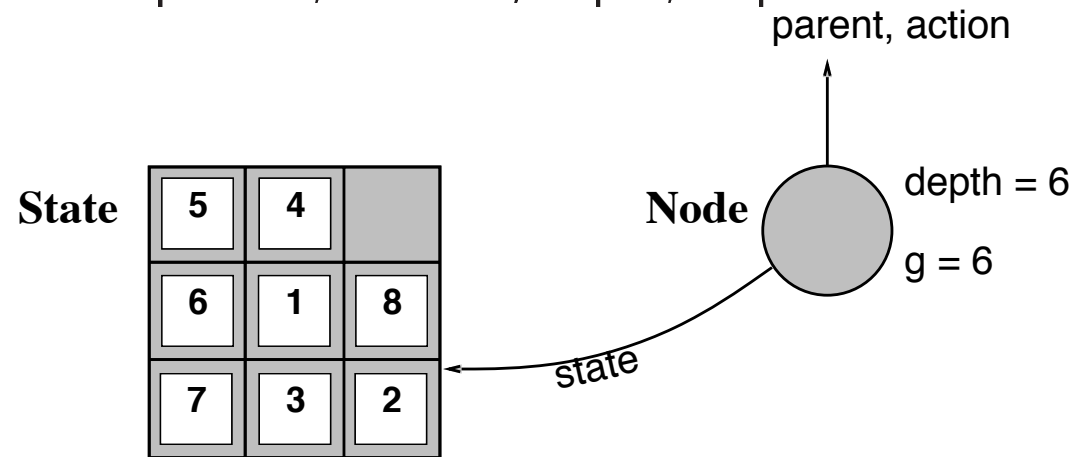# Tree search example
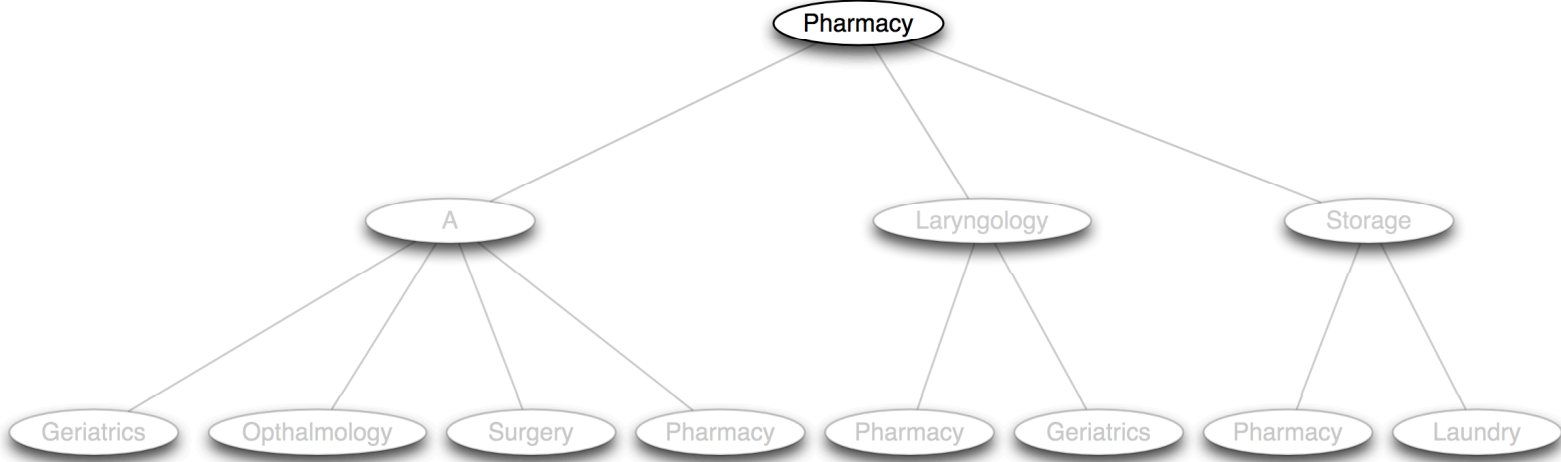
Tree search example
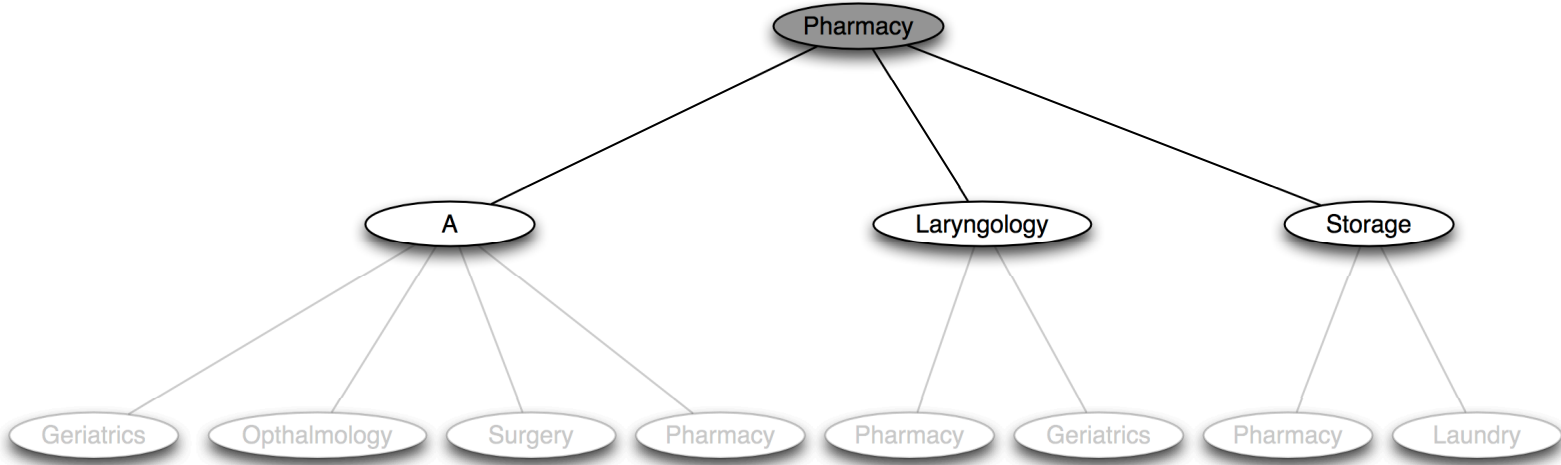
Tree search example

Implementation: states vs. nodes

A state is a (representation of) a physical configuration
A node is a data structure constituting part of a search tree
      includes parent, children, depth, path cost $g(x)$
States do not have parents, children, depth, or path cost!

State

5	4	
6	1	8
7	3	2

Node

parent, action

depth = 6

g = 6

state

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 ← INSERT(MAKE-NODE(INITIAL-STATE[*problem*]), *fringe*)
    loop do
        if *fringe* is empty **then return** failure
        node ← REMOVE-FRONT(*fringe*)
        if GOAL-TEST(*problem*, STATE(*node*)) **then return** *node*
        fringe ← INSERTALL(EXPAND(*node, problem*), *fringe*)
function EXPAND(*node, problem*) **returns** a set of nodes
    successors ← the empty set
    for each *action, result* **in** SUCCESSOR-FN(*problem*, STATE[*node*]) **do**
        s ← a new NODE
        PARENT-NODE[*s*] ← *node*; ACTION[*s*] ← *action*; STATE[*s*] ← *result*
        PATH-COST[*s*] ← PATH-COST[*node*] + STEP-COST(STATE[*node*], *action*,
result)
        DEPTH[*s*] ← 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

Sometimes called **blind** search strategies

Breadth-first search

Uniform-cost search

Depth-first search

Depth-limited search
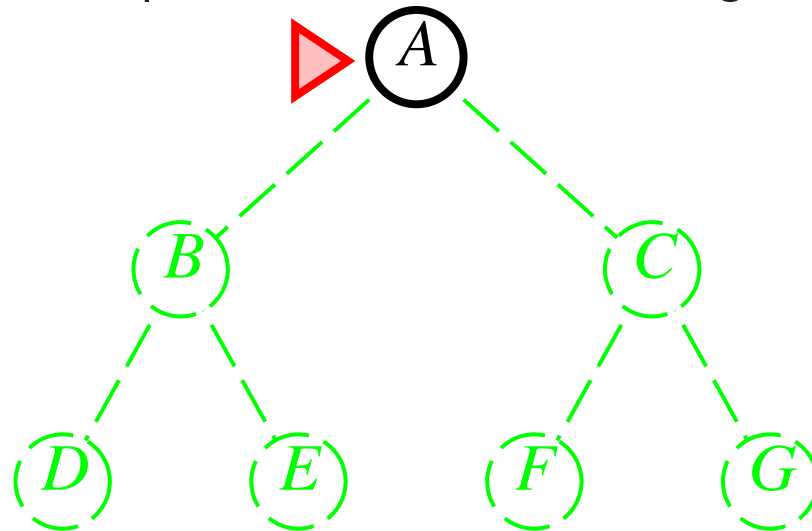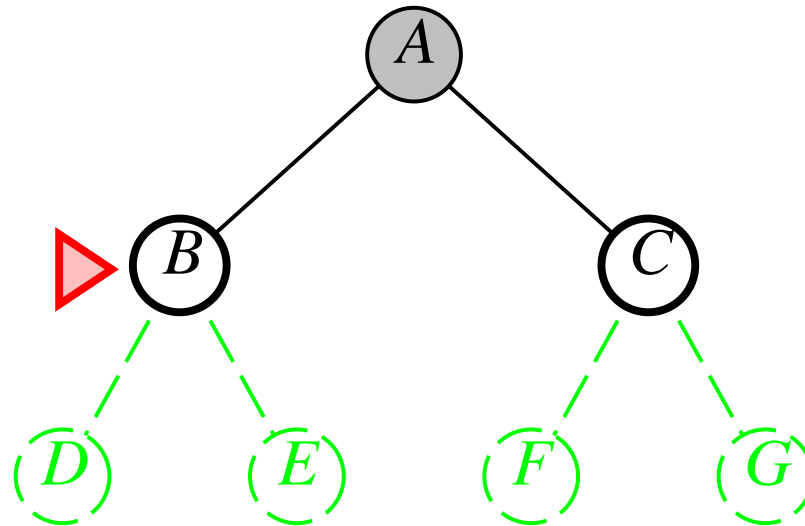
Iterative deepening search

Breadth-first search

Expand shallowest unexpanded node

Implementation:

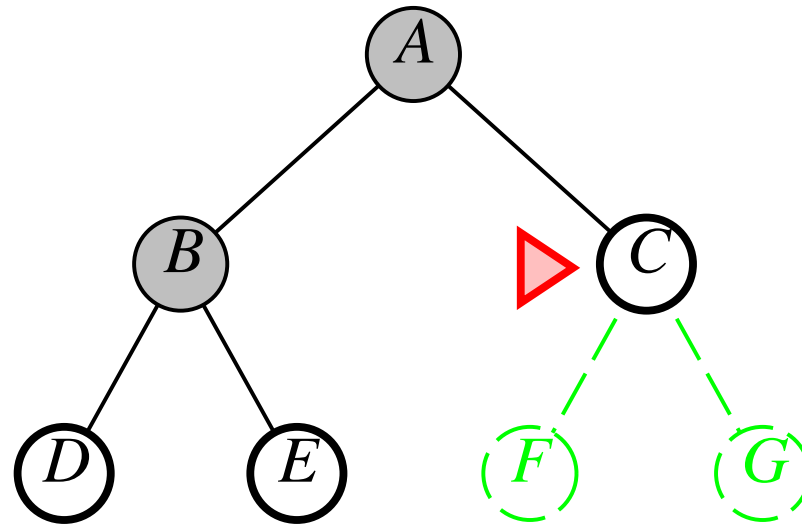fringe is a FIFO queue, i.e., new successors go at end

Breadth-first search

Expand shallowest unexpanded node

Implementation:

    fringe is a FIFO queue, i.e., new successors go at end

Breadth-first search

Expand shallowest unexpanded node

Implementation:

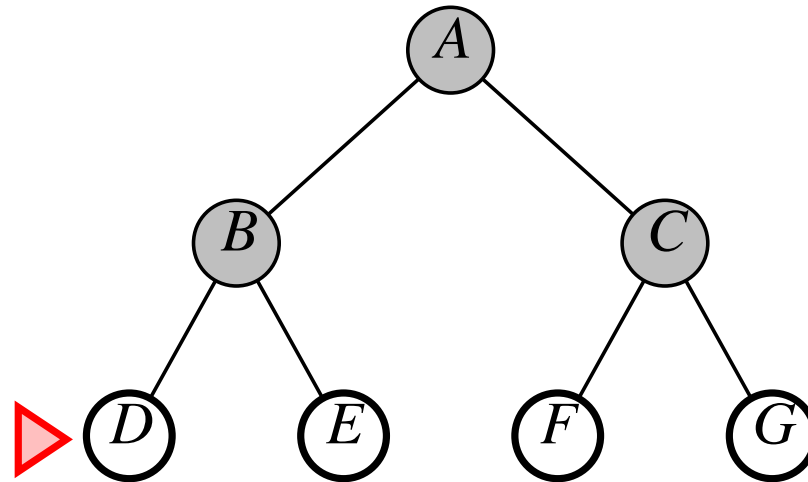   *fringe* is a FIFO queue, i.e., new successors go at end

Breadth-first search

Expand shallowest unexpanded node

Implementation:

fringe is a FIFO queue, i.e., new successors go at end

Properties of breadth-first search

Complete??

Properties of breadth-first search

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

<u>Time</u>??

Properties of breadth-first search

Complete?? Yes (if b is finite)

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

Space??

Properties of breadth-first search

<u>Complete</u>?? 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)

<u>Optimal</u>??

Properties of breadth-first search

<u>Complete</u>?? 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)

<u>Optimal</u>?? 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

Complete?? Yes, if step cost $\geq \epsilon$

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

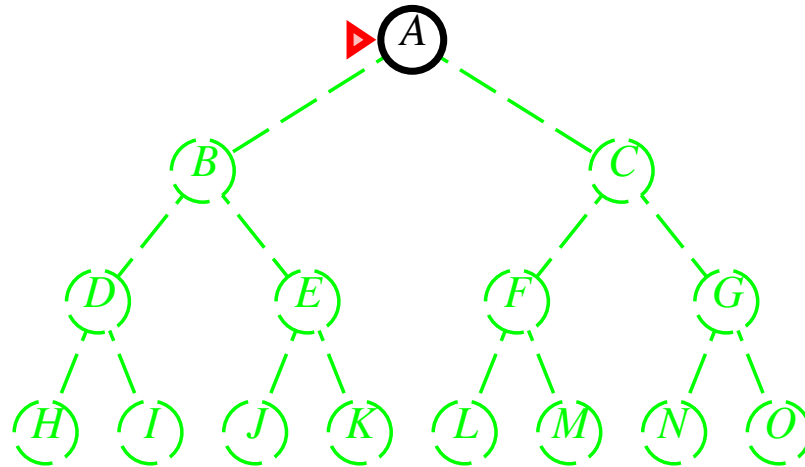Space?? # of nodes with $g \leq$ cost of optimal solution, $O(b^{\lceil C^*/\epsilon \rceil})$

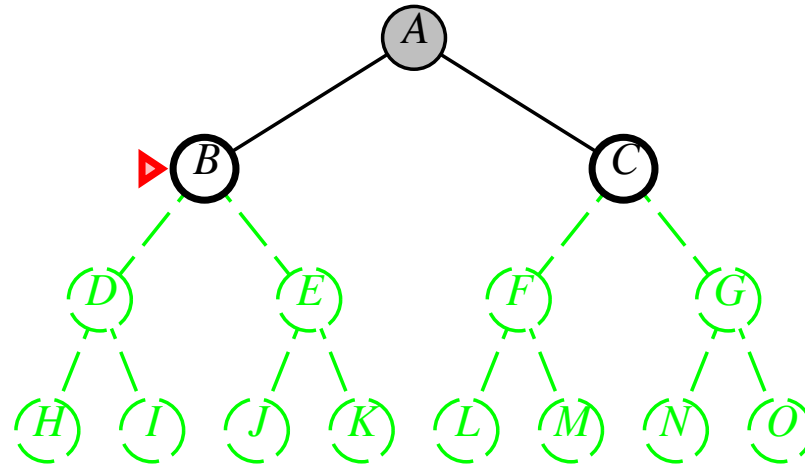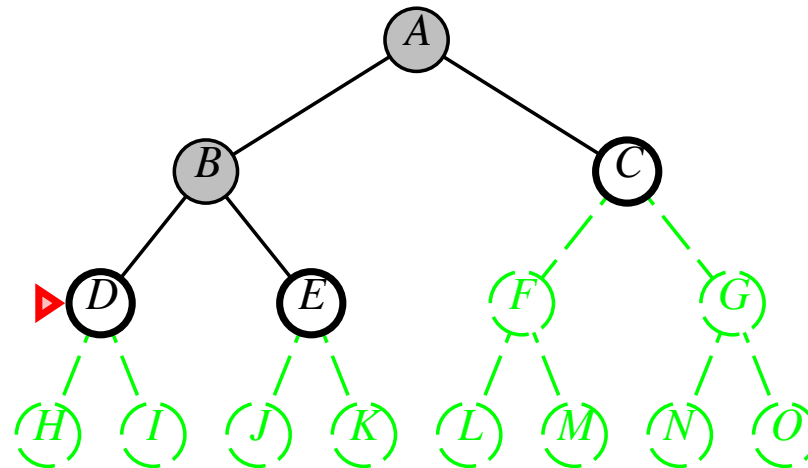Optimal?? Yes—nodes expanded in increasing order of $g(n)$

Depth-first search

Expand deepest unexpanded node

Implementation:

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

Depth-first search

Expand deepest unexpanded node

Implementation:

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

Depth-first search

Expand deepest unexpanded node

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

Depth-first search

Expand deepest unexpanded node

Implementation:

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

Depth-first search

Expand deepest unexpanded node

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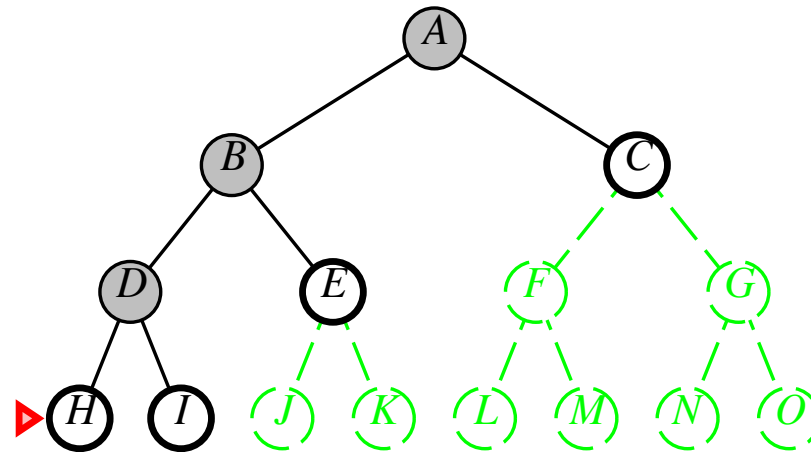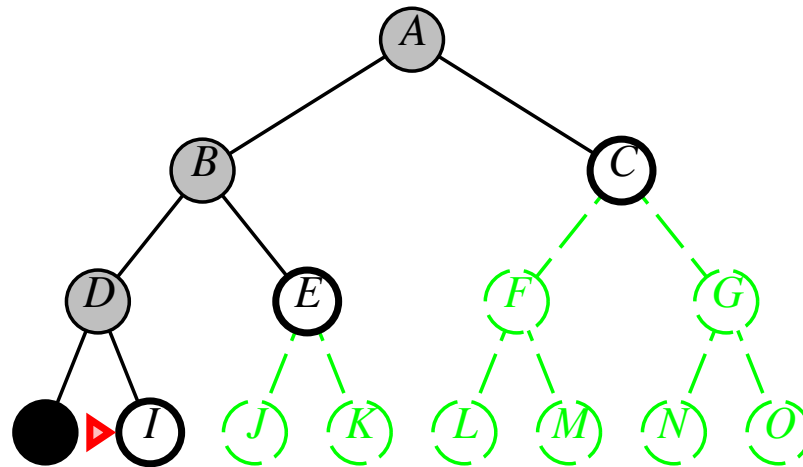$fringe$ = LIFO queue, i.e., put successors at front

Depth-first search

Expand deepest unexpanded node

Implementation:

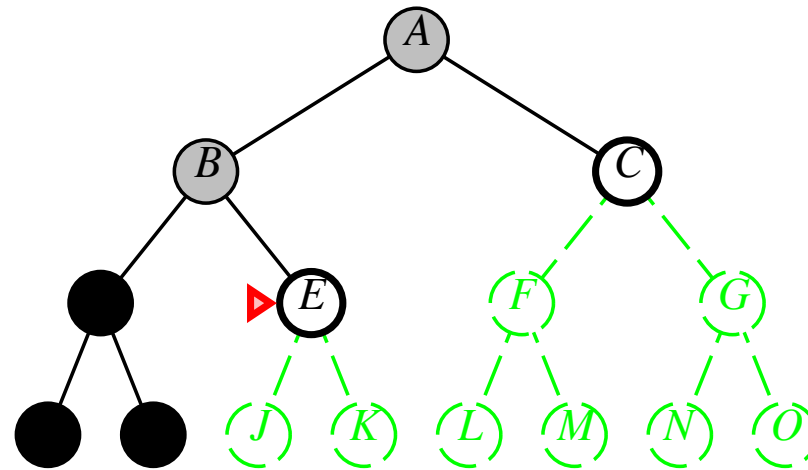$fringe =$ LIFO queue, i.e., put successors at front

Depth-first search

Expand deepest unexpanded node

Implementation:

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

Depth-first search

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

Depth-first search

Expand deepest unexpanded node

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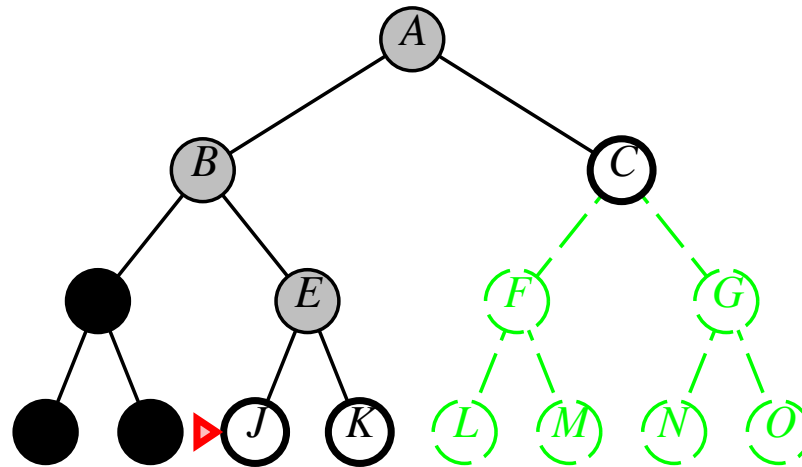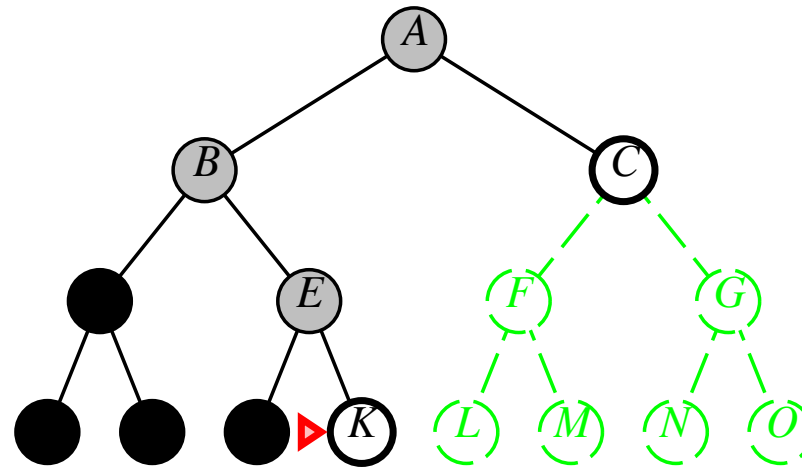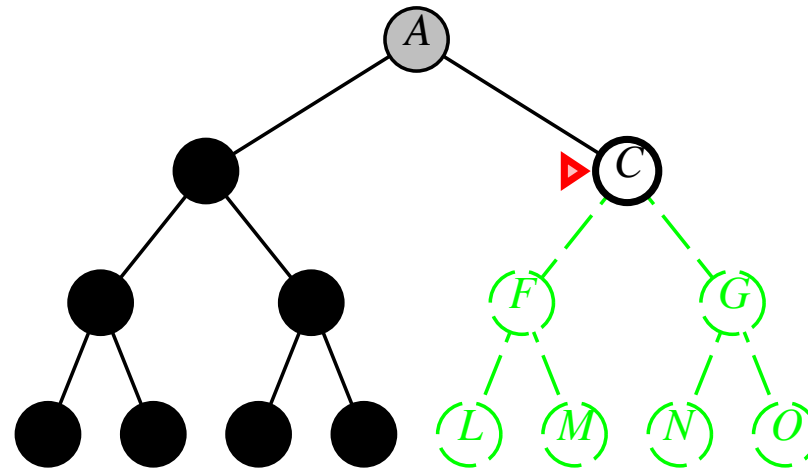 $fringe = $ LIFO queue, i.e., put successors at front

Depth-first search

Expand deepest unexpanded node

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

Depth-first search

Expand deepest unexpanded node

Implementation:
$fringe$ = LIFO queue, i.e., put successors at front
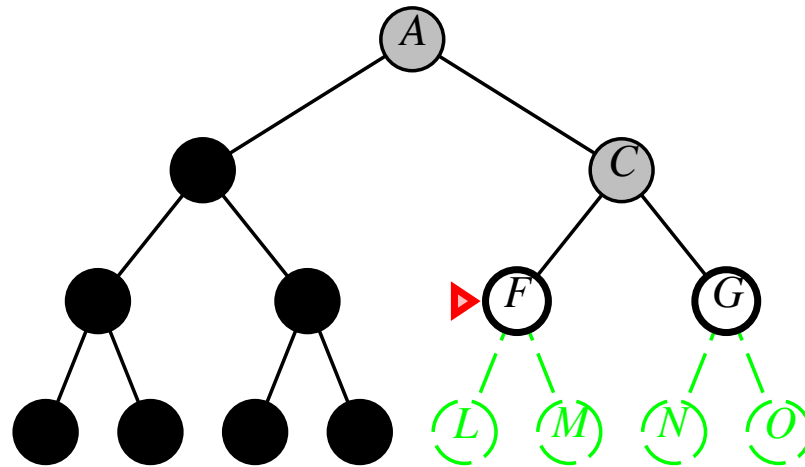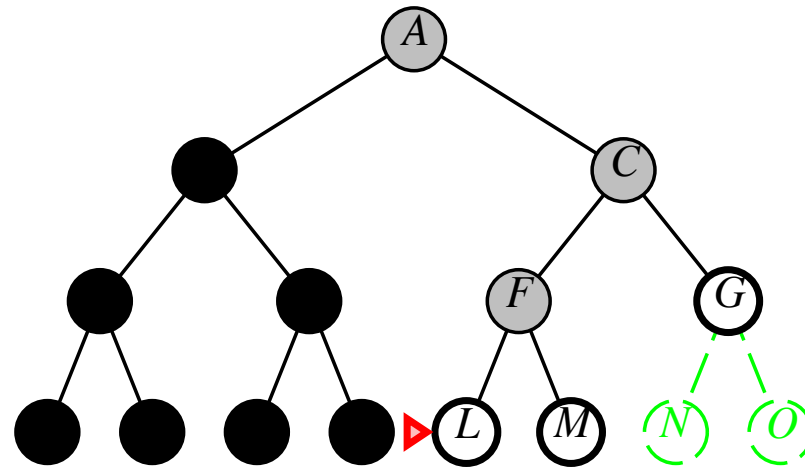
Properties of depth-first search

Complete??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops
Modify to avoid repeated states along path
$\Rightarrow$ complete in finite spaces

Time??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops
      Modify to avoid repeated states along path
          $\Rightarrow$ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d
      but if solutions are dense, may be much faster than breadth-first

Space??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops
        Modify to avoid repeated states along path
            $\Rightarrow$ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d
        but if solutions are dense, may be much faster than breadth-first

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

Optimal??

Properties of depth-first search

Complete?? No: fails in infinite-depth spaces, spaces with loops
    Modify to avoid repeated states along path
        $\Rightarrow$ complete in finite spaces

Time?? $O(b^m)$: terrible if m is much larger than d
    but 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? ← 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 ← RECURSIVE-DLS(*successor*, *problem*, *limit*)
        if *result* = *cutoff* **then** *cutoff-occurred?* ← true
        else if *result* ≠ *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*)
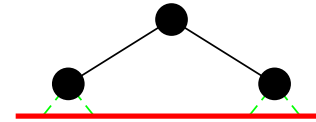   **if** *result* ≠ cutoff **then return** *result*
    end

Limit = 0

Iterative deepening search $l = 1$

Limit = 1
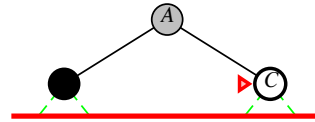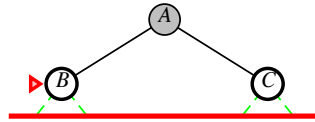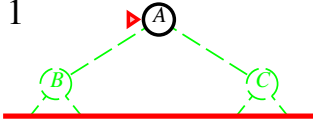
Iterative deepening search $l = 2$

Limit = 2
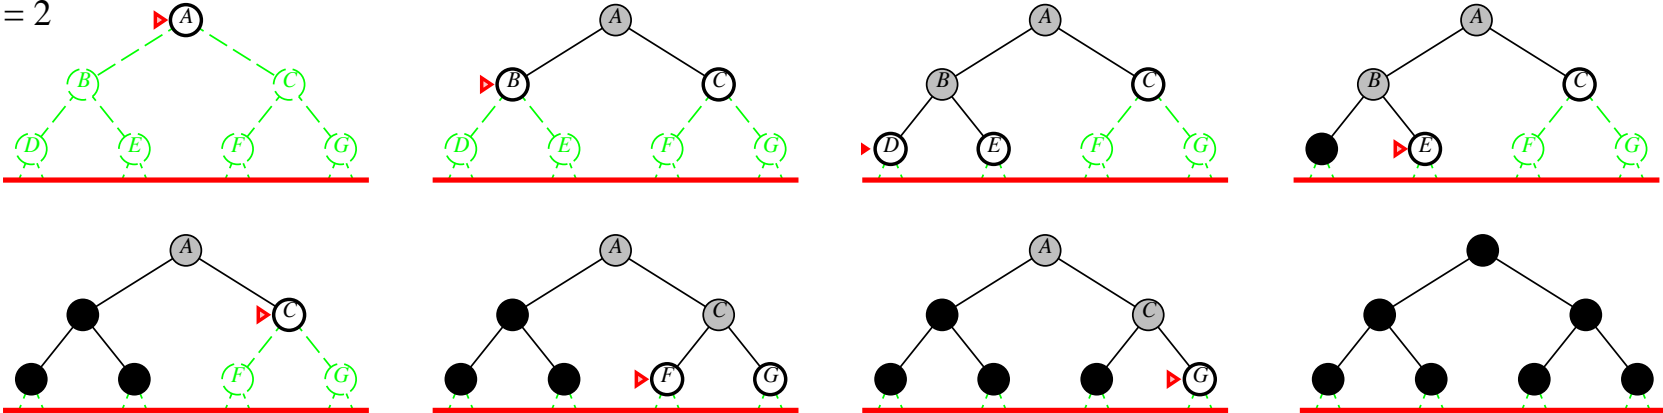
Limit = 3
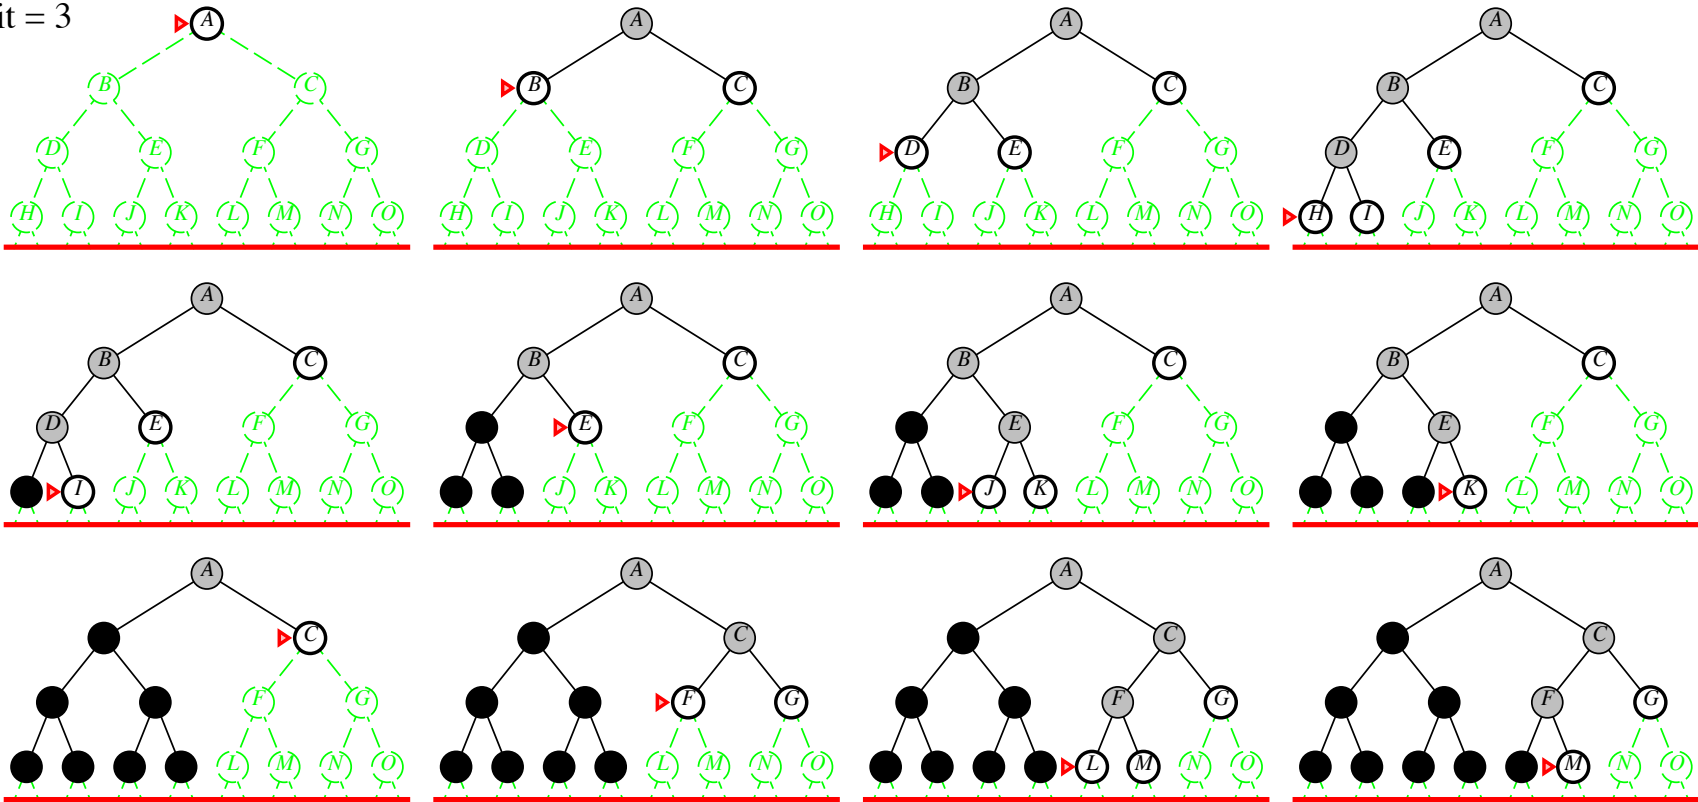
Properties of iterative deepening search

Complete??

Properties of iterative deepening search

Complete?? Yes

Time??

Properties of iterative deepening search

<u>Complete</u>?? Yes

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

<u>Space</u>??

Properties of iterative deepening search

Complete?? Yes

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

Space?? $O(bd)$

Optimal??

Properties of iterative deepening search

Complete?? Yes

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

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-First	Uniform-Cost	Depth-First	Depth-Limited	Iterative Deepening
Complete?	Yes*	Yes*	No	Yes, if $l \geq d$	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!
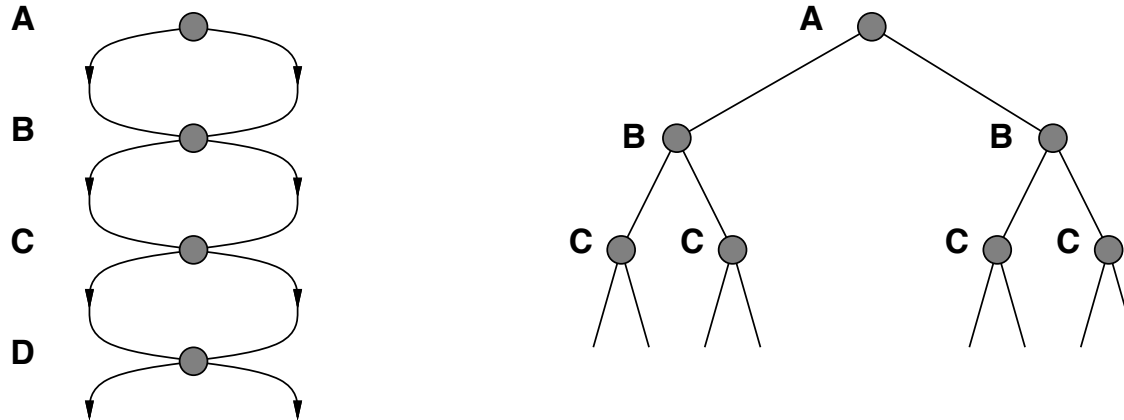
Graph search

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

    closed ← an empty set
    fringe ← INSERT(MAKE-NODE(INITIAL-STATE[*problem*]), *fringe*)
    loop do
        if *fringe* is empty **then return** failure
        node ← 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 ← INSERTALL(EXPAND(*node*, *problem*), *fringe*)
    end

Partial 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

◇ Best-first search

◇ A* search

◇ Heuristics

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

Blocket with distances in seconds

Laryngology   Geriatrics

71

75   151

Pharmacy   140   A   Ophthalmology   Dermathology

80   97   85   B

118   99

Surgery   211   101   98

Storage   ICU

111   Nephrology

135   95

ER

Laundry   145   C   134

Radiology   92   Orthopedy

142
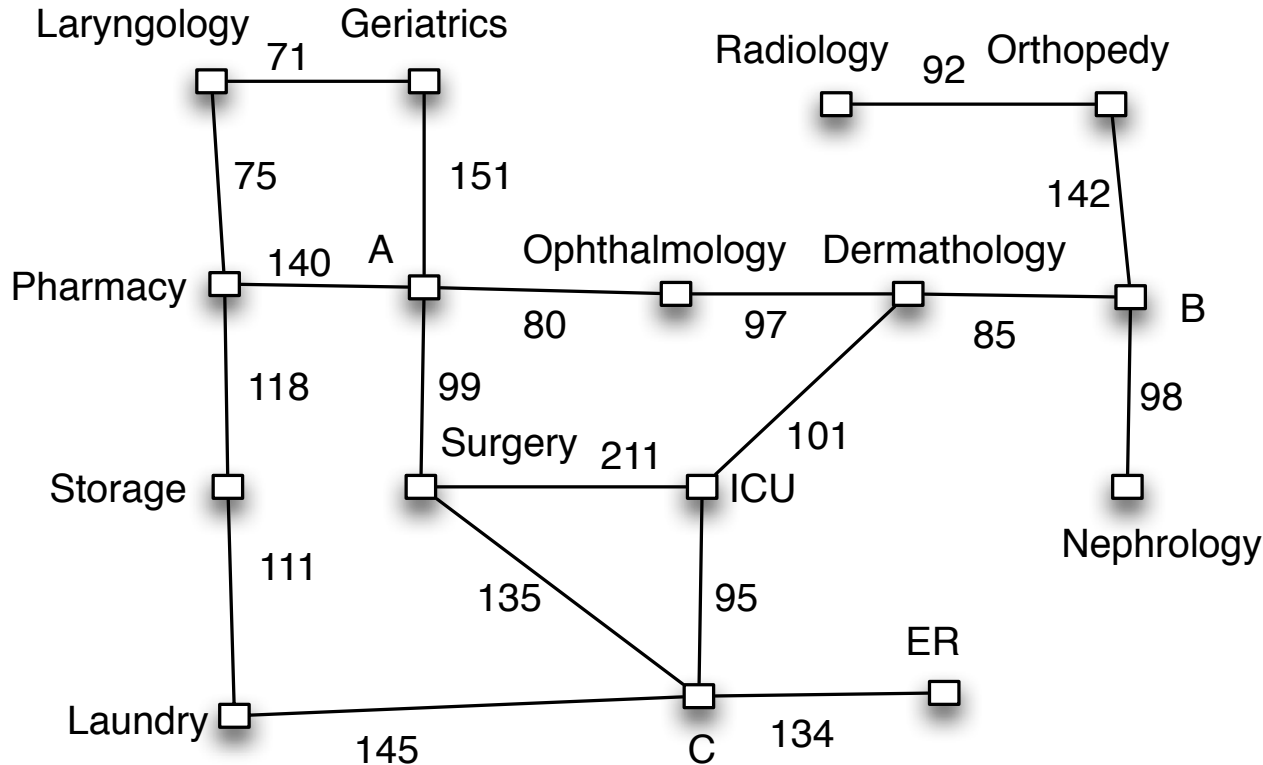
Direct 3D distances to ICU:

Radiology	260
Orthopedy	285
B	130
Nephrology	250
Dermathology	90
Ophthalmology	110
A	230
Surgery	100
C	77
ER	162
Geriatrics	280
Laryngology	312
Pharmacy	280
Storage	235
Laundry	220
ICU	0

Greedy search

Evaluation function $h(n)$ (**h**euristic)
$= $ estimate of cost from n to the closest goal

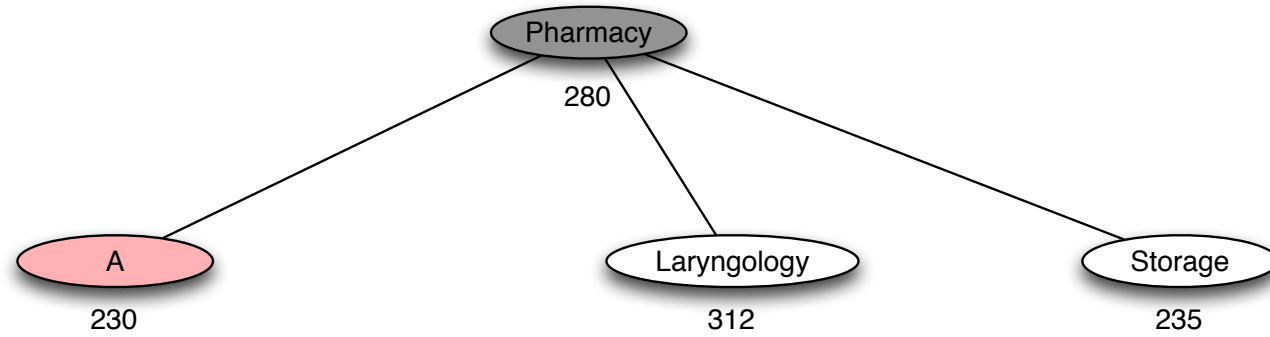E.g., $h_{\mathrm{SLD}}(n) = $ straight-line distance from n to ICU

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

Greedy search example

Pharmacy

280

Greedy search example

Greedy search example

Greedy search example

Pharmacy
280

A
230

Laryngology
312

Storage
235

Geriatrics
280

Opthalmology
110

Surgery
100

Pharmacy
280

A
230

C
77

ICU
0
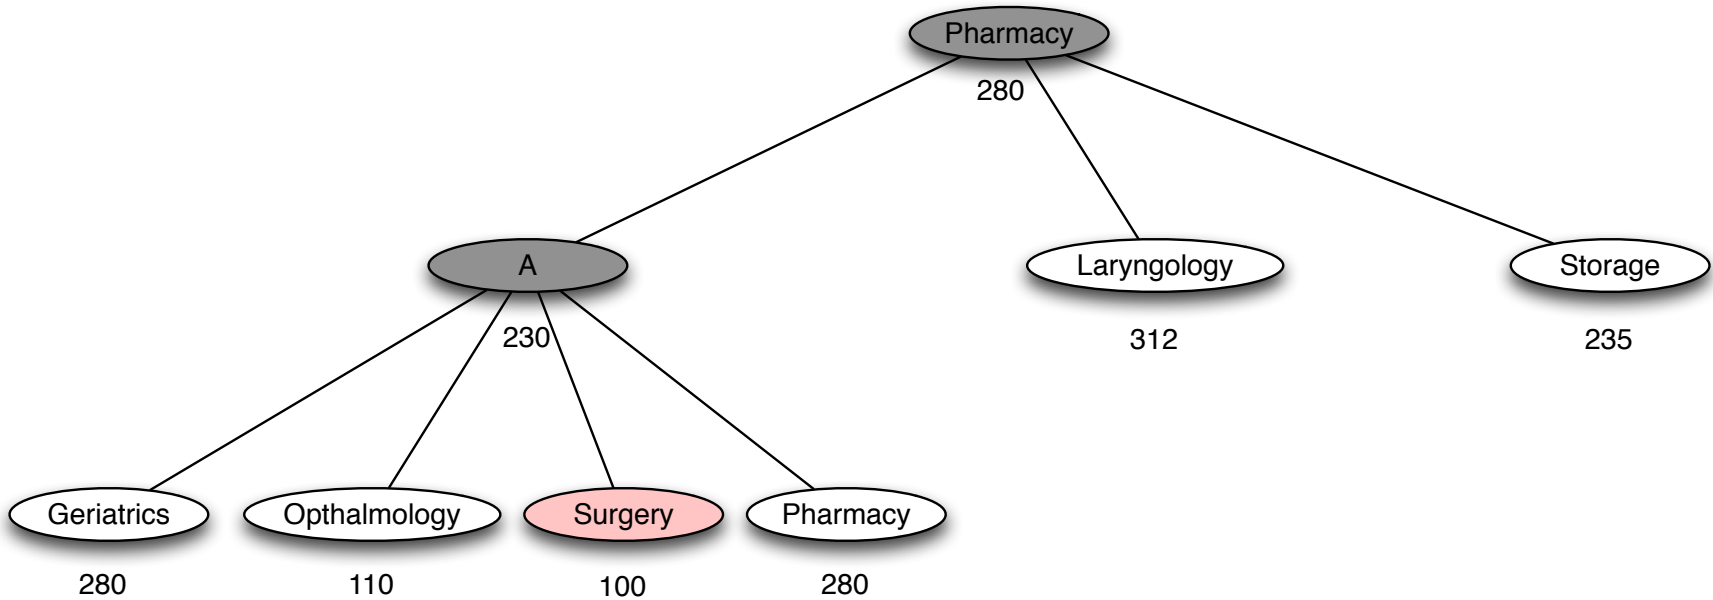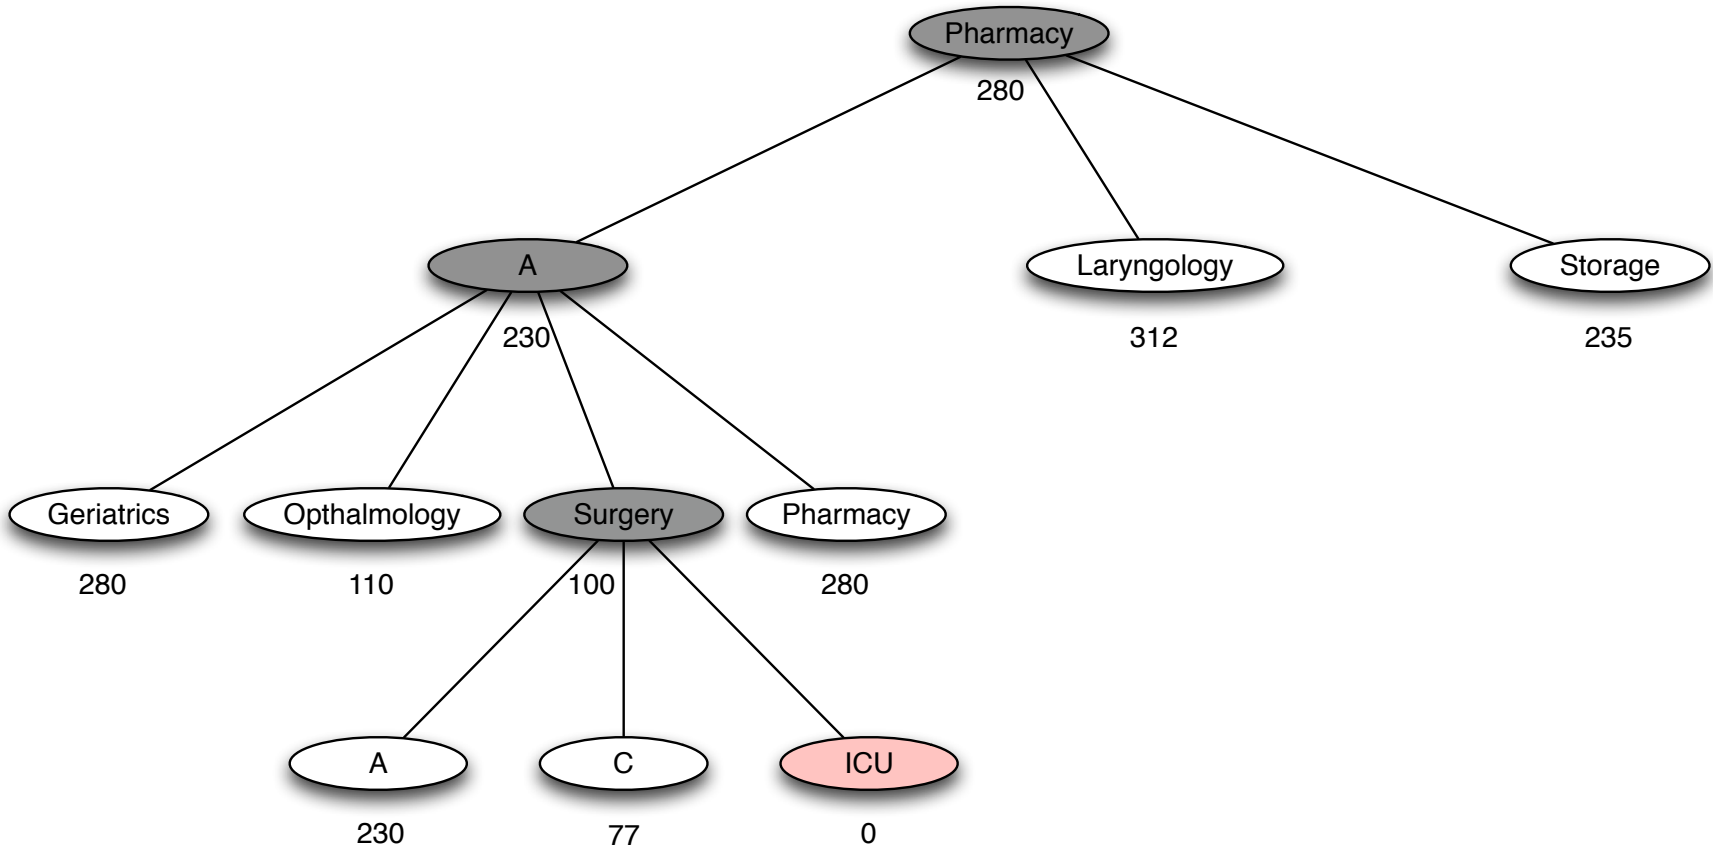
Properties of greedy search

Complete??

Properties of greedy search

<u>Complete</u>?? No–can get stuck in loops, e.g., with Geriatrics as goal,

Radiology $\rightarrow$ Orthopedy $\rightarrow$ Radiology $\rightarrow$ Orthopedy $\rightarrow$

Complete in finite space with repeated-state checking

<u>Time</u>??

Properties of greedy search

Complete?? No–can get stuck in loops, e.g.,
$$\text{Radiology} \rightarrow \text{Orthopedy} \rightarrow \text{Radiology} \rightarrow \text{Orthopedy} \rightarrow$$
Complete in finite space with repeated-state checking

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

Space??

Properties of greedy search

Complete?? No--can get stuck in loops, e.g.,

$\qquad$ Radiology $\rightarrow$ Orthopedy $\rightarrow$ Radiology $\rightarrow$ Orthopedy $\rightarrow$

Complete in finite space with repeated-state checking

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

Space?? $O(b^m)$—keeps all nodes in memory

Optimal??

Properties of greedy search

<u>Complete</u>?? No–can get stuck in loops, e.g.,

　　　　Radiology $\rightarrow$ Orthopedy $\rightarrow$ Radiology $\rightarrow$ Orthopedy $\rightarrow$

Complete in finite space with repeated-state checking

<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

<u>Optimal</u>?? 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 n
$h(n)$ = estimated cost to goal from n
$f(n)$ = estimated total cost of path through n to goal

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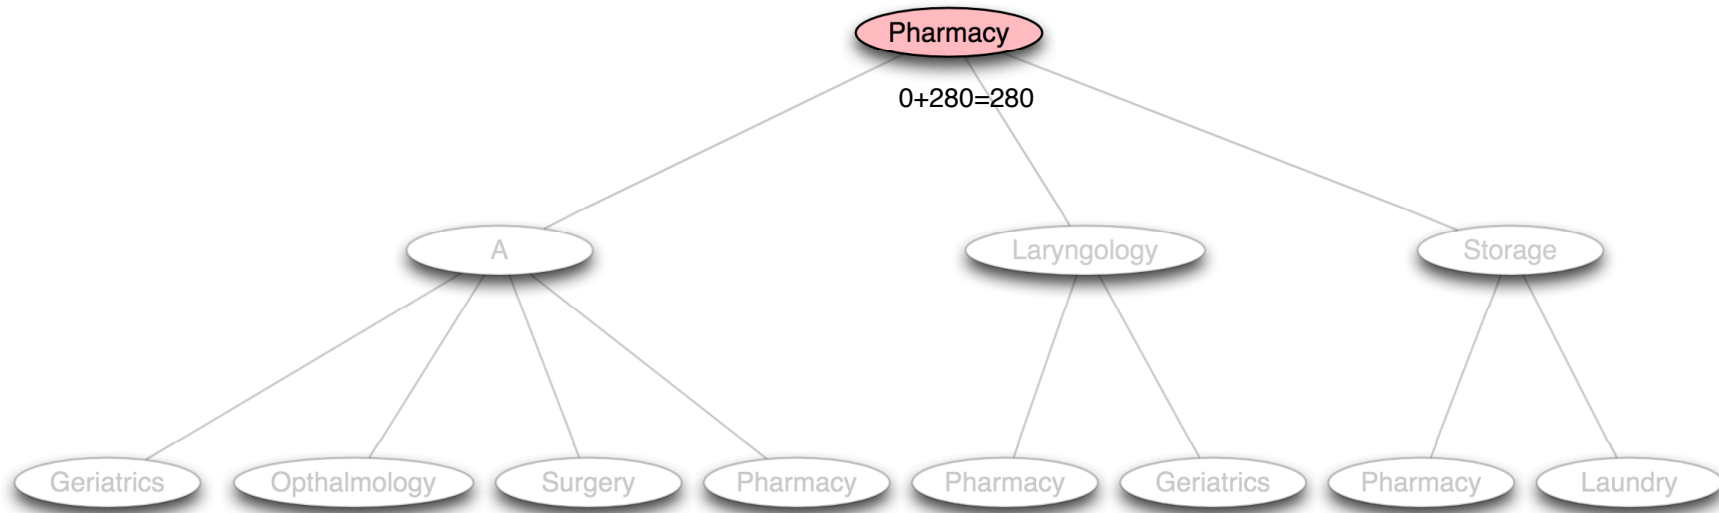i.e., $h(n) \leq h^*(n)$ where $h^*(n)$ is the **true** cost from n.
(Also require $h(n) \geq 0$, so $h(G) = 0$ for any goal G.)

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

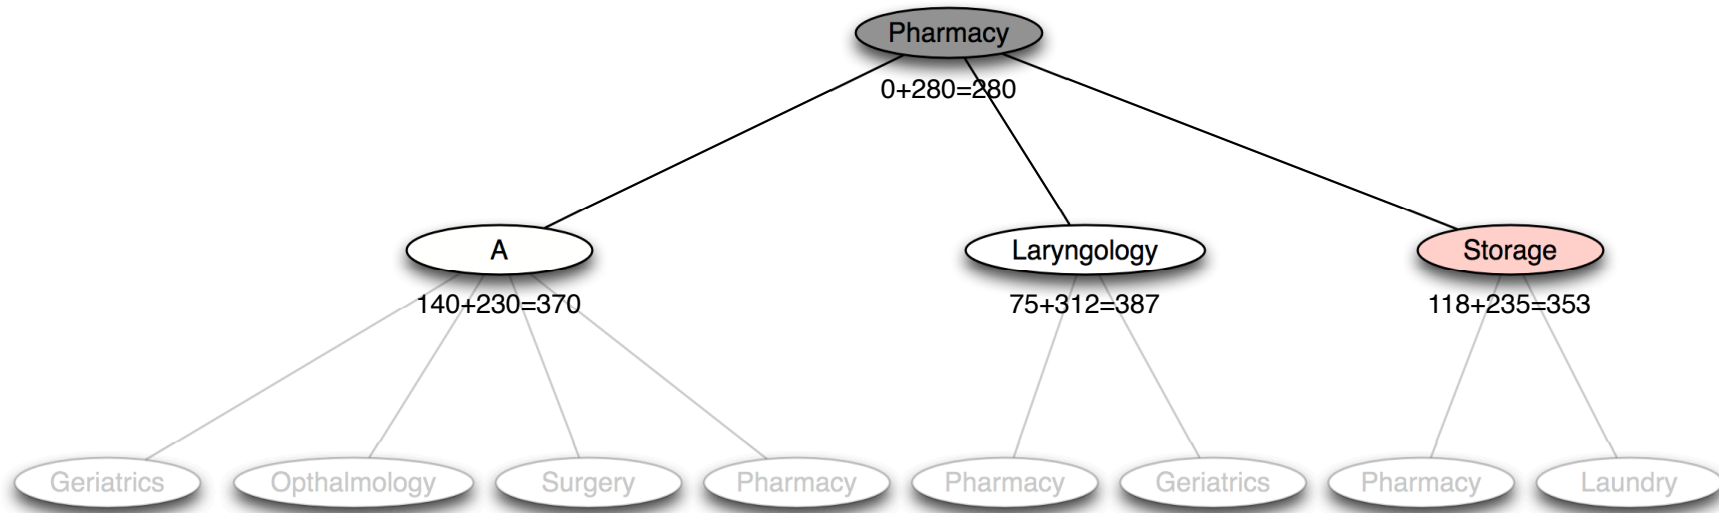Theorem: A* search is optimal

A* search example

Pharmacy
0+280=280

A   Laryngology   Storage

Geriatrics   Opthalmology   Surgery   Pharmacy   Pharmacy   Geriatrics   Pharmacy   Laundry

A* search example

A* search example

A* search example

Pharmacy
0+280=280

A
140+230=370

Laryngology
75+312=387

Storage
118+235=353

Geriatrics
291+280=571

Opthalmology
220+110=330

Surgery
239+100=339

Pharmacy
280+280=560

Pharmacy

Geriatrics

Pharmacy
230+286=516

Laundry
229+220=449

A* search example

Pharmacy

0+280=280

A

140+230=370

Laryngology

75+312=387

Storage

118+235=353

Geriatrics

291+280=571

Opthalmology

220+110=330

Surgery

239+100=339

Pharmacy

280+280=560

Pharmacy

Geriatrics

Pharmacy
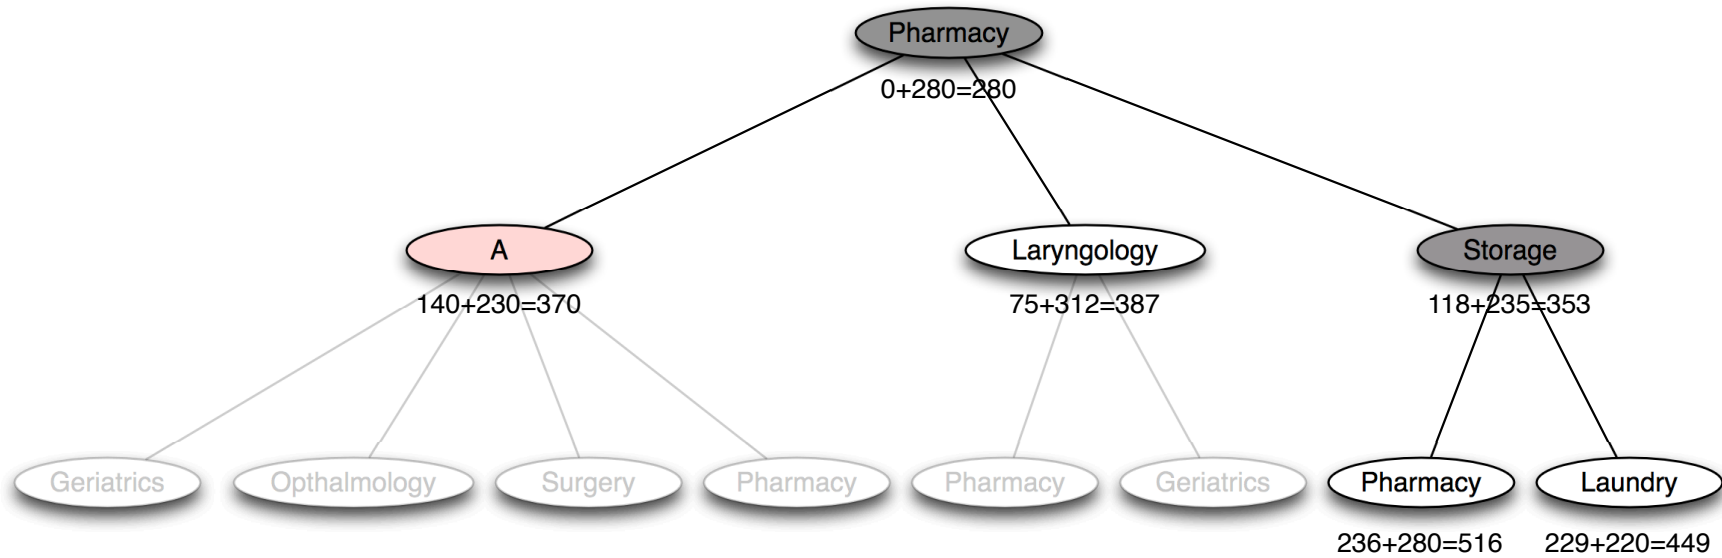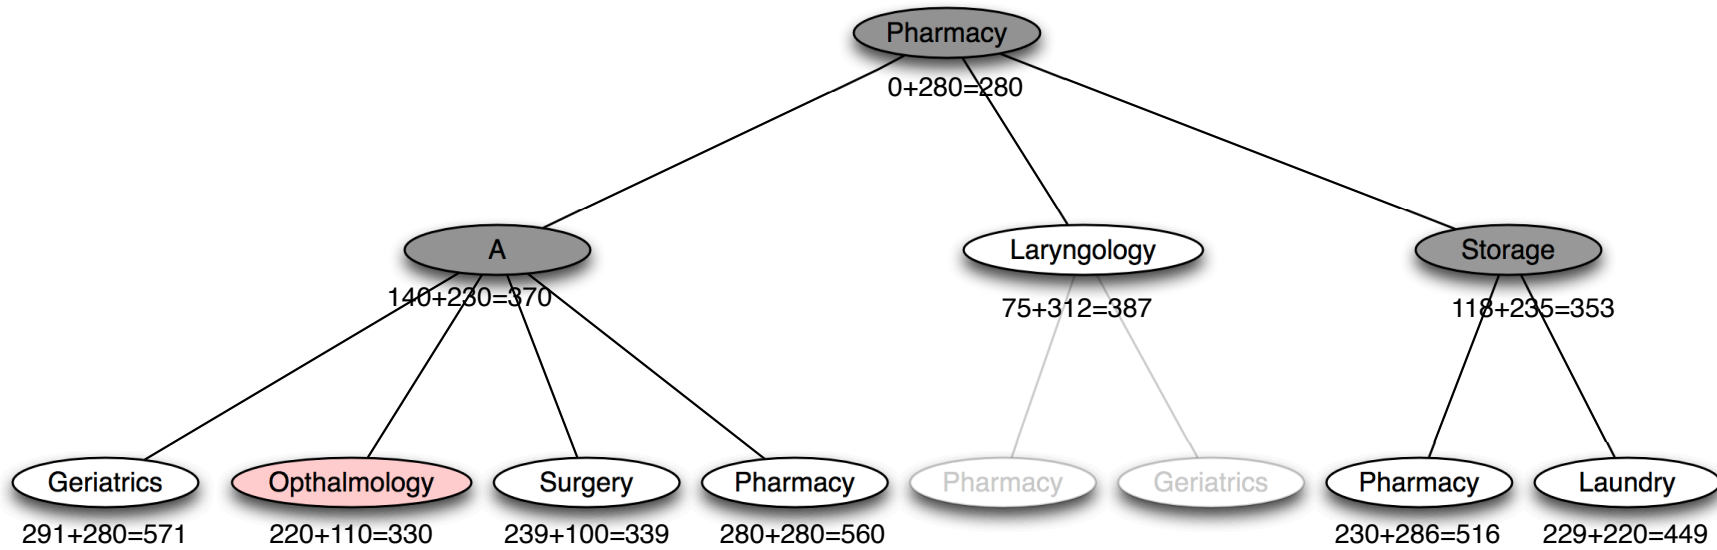
230+286=516

Laundry

229+220=449

Dermathology

317+90=407

A

300+230=530
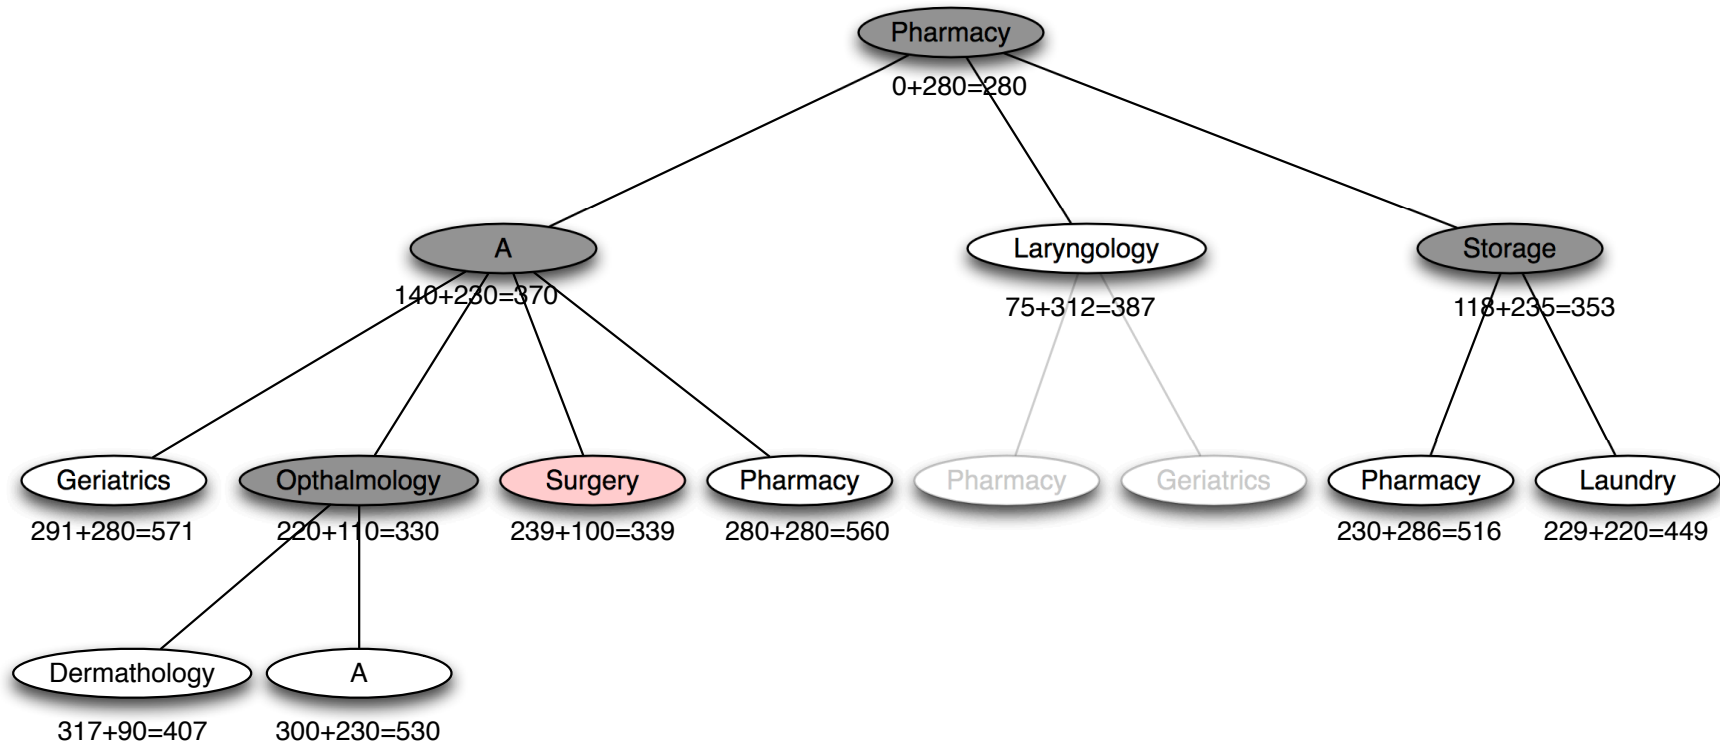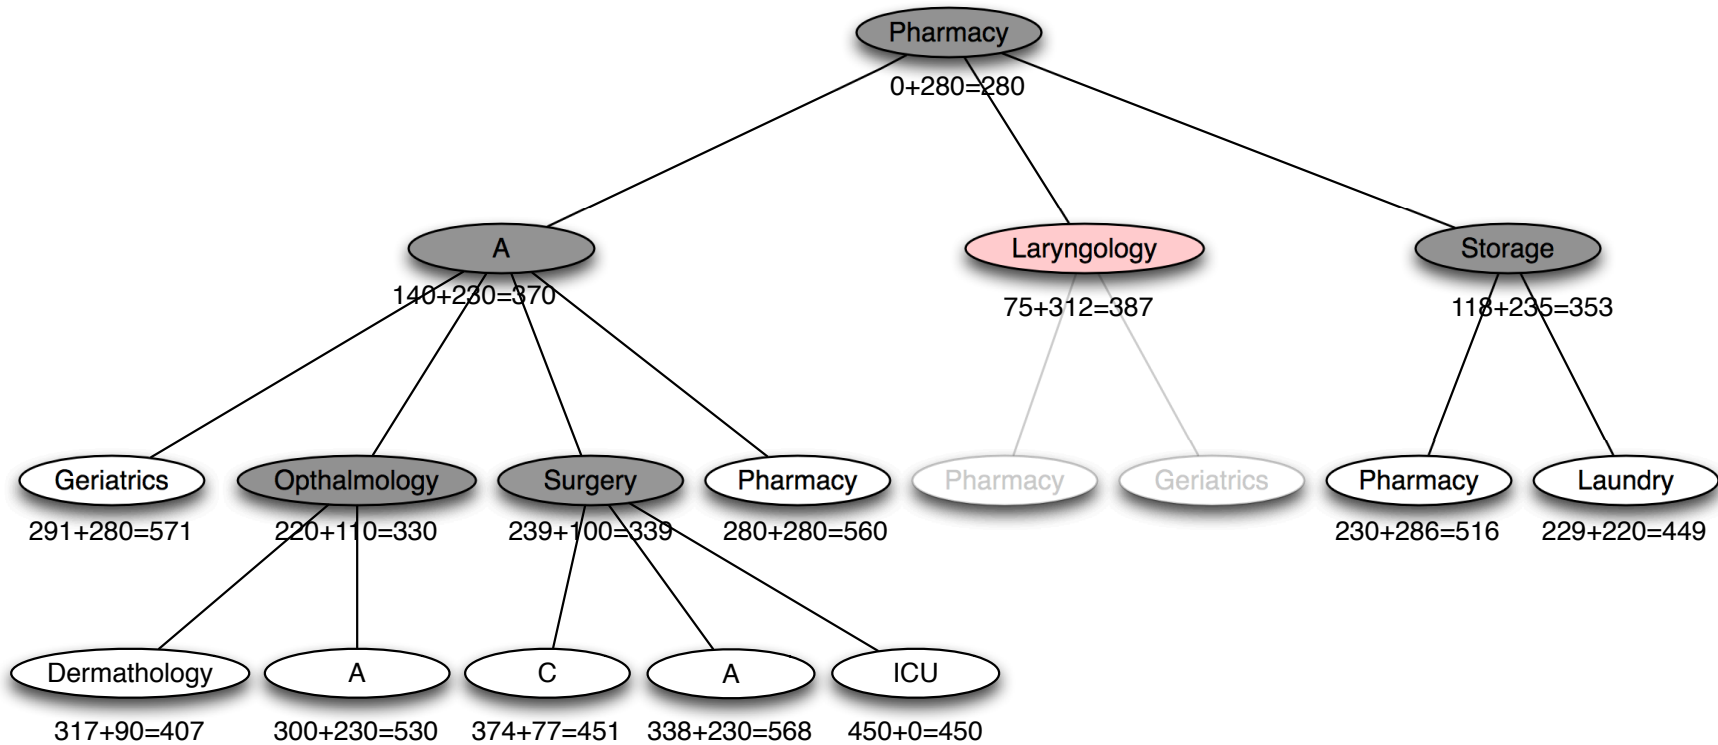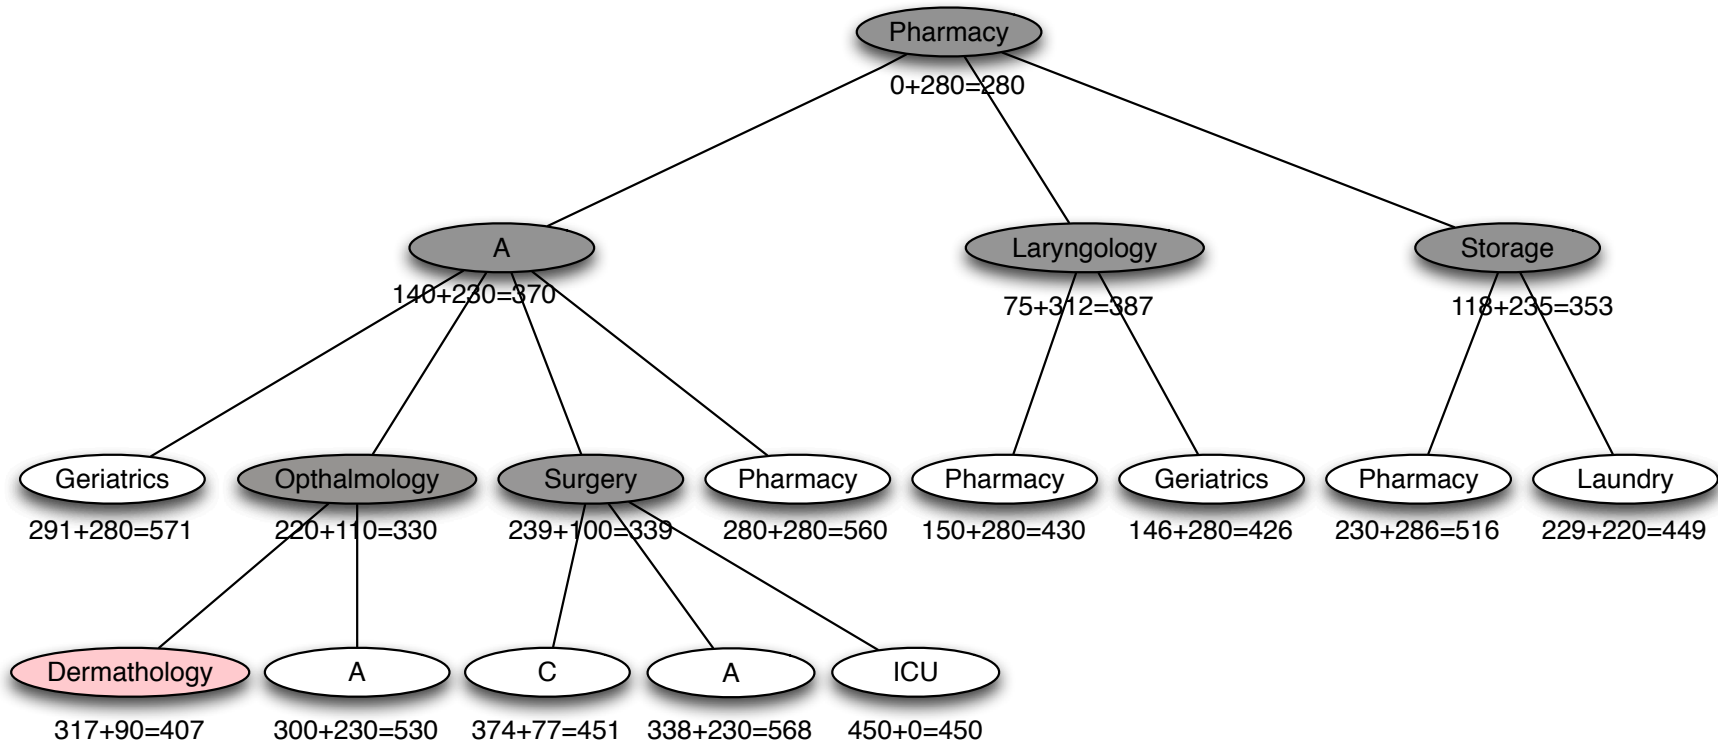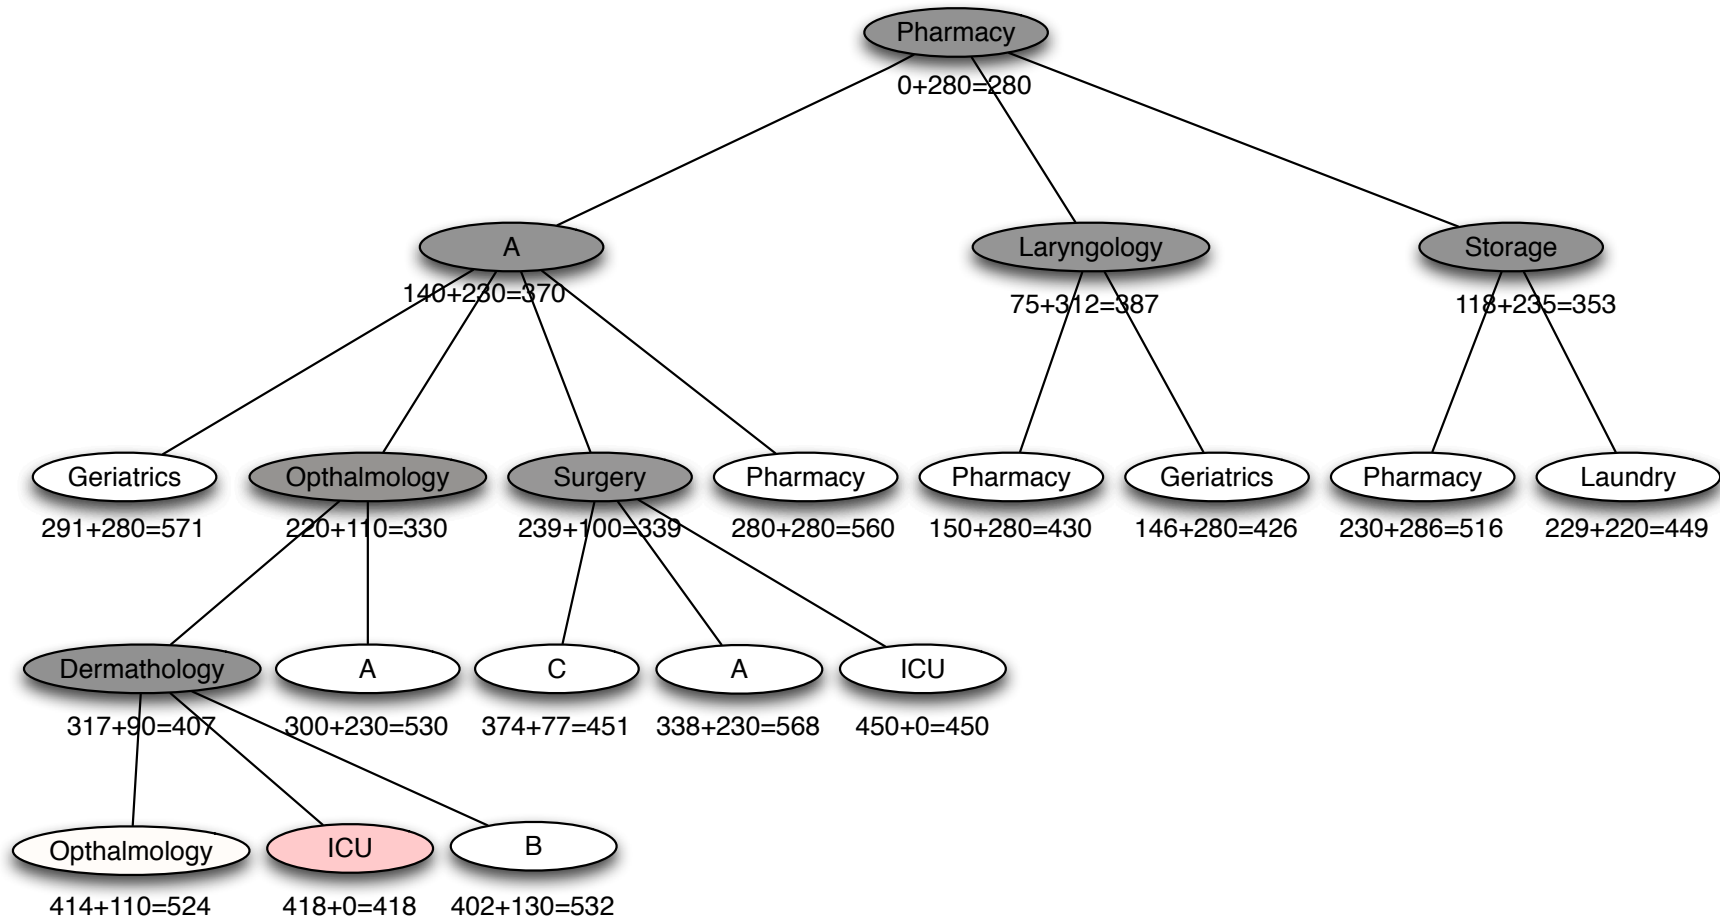
A* search example

A* search example

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.

Start

n

G G_2
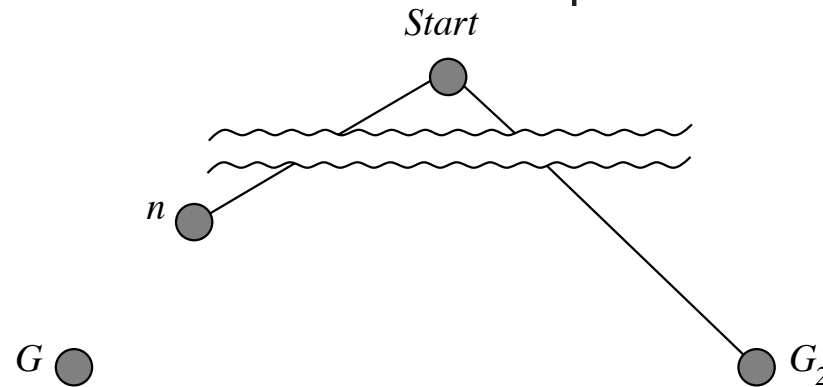
$$
\begin{aligned}
f(G_2) &= g(G_2) && \text{since } h(G_2) = 0 \\
&> g(G_1) && \text{since } G_2 \text{ is suboptimal} \\
&\geq f(n) && \text{since } h \text{ is admissible}
\end{aligned}
$$

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

Properties of A*

Complete??

Properties of A*

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

Time??

Properties of A*

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

Time?? Exponential in [relative error in $h \times$ length of solution]

Space??

Properties of A*

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

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

Space?? Keeps all nodes in memory

Optimal??

Properties of A*

<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.]

<u>Space</u>?? 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 $f(n) > C^*$

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)

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

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

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)

7	2	4
5		6
8	3	1

Start State

1	2	3
4	5	6
7	8	

Goal State

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

Dominance

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

Typical search costs:

$d = 14$ IDS $= 3{,}473{,}941$ nodes
 A*$(h_1) = 539$ nodes
 A*$(h_2) = 113$ nodes
$d = 24$ IDS $\approx 54{,}000{,}000{,}000$ nodes
 A*$(h_1) = 39{,}135$ nodes
 A*$(h_2) = 1{,}641$ nodes

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

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