


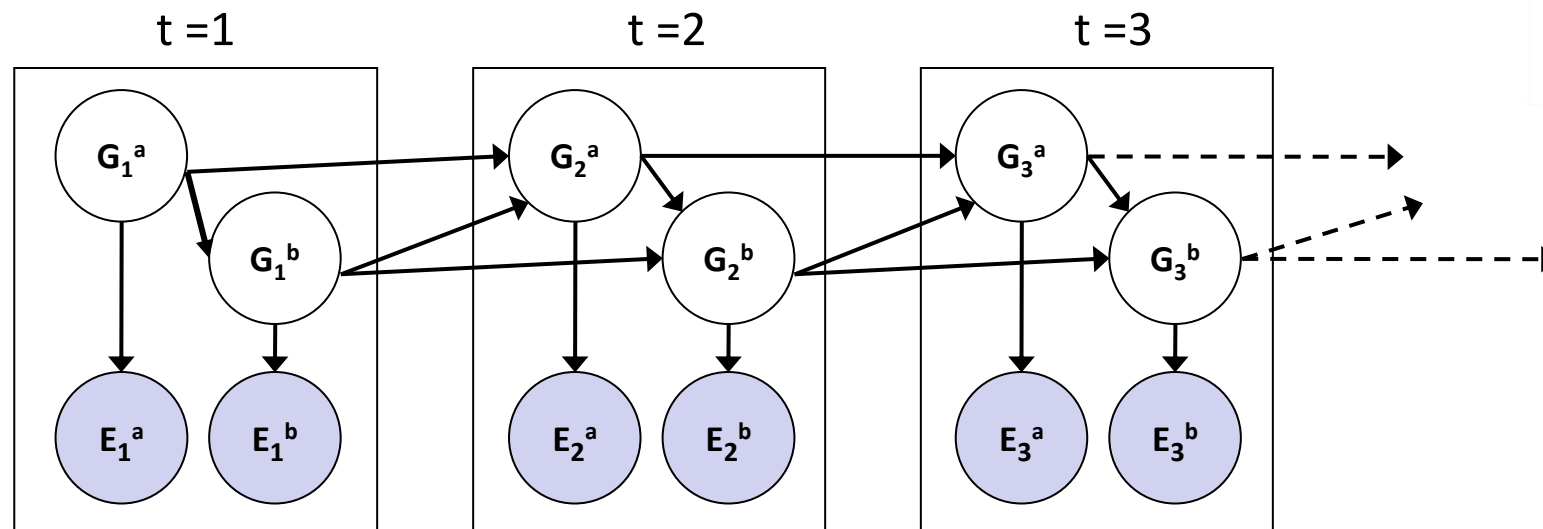

# Dynamic Bayes Nets

---



# Dynamic Bayes Nets (DBNs)


- We want to track multiple variables over time, using multiple sources of evidence
- Idea: Repeat a fixed Bayes net structure at each time
- Variables from time  $t$  can condition on those from  $t-1$



- Dynamic Bayes nets are a generalization of HMMs

# DBNs and HMMs


- Every HMM is a single-variable DBN
- Every discrete DBN is an HMM
  - HMM state is Cartesian product of DBN state variables



- Sparse dependencies  $\Rightarrow$  exponentially fewer parameters in DBN
  - E.g., 20 state variables, 3 parents each;  
DBN has  $20 \times 2^3 = 160$  parameters, HMM has  $2^{20} \times 2^{20} \approx 10^{12}$  parameters

# Exact Inference in DBNs

- Variable elimination applies to dynamic Bayes nets
- Procedure: “unroll” the network for  $T$  time steps, then eliminate variables until  $P(X_T | e_{1:T})$  is computed



- Online belief updates: Eliminate all variables from the previous time step; store factors for current time only


# DBN Particle Filters

---

- A particle is a complete sample for a time step
- **Initialize:** Generate prior samples for the  $t=1$  Bayes net
  - Example particle:  $\mathbf{G}_1^a = (3,3)$   $\mathbf{G}_1^b = (5,3)$
- **Elapse time:** Sample a successor for each particle
  - Example successor:  $\mathbf{G}_2^a = (2,3)$   $\mathbf{G}_2^b = (6,3)$
- **Observe:** Weight each entire sample by the likelihood of the evidence conditioned on the sample
  - Likelihood:  $P(\mathbf{E}_1^a | \mathbf{G}_1^a) * P(\mathbf{E}_1^b | \mathbf{G}_1^b)$
- **Resample:** Select prior samples (tuples of values) in proportion to their likelihood


# CS 188: Artificial Intelligence

## Midterm Review



Instructors: Evgeny Pobachienko – UC Berkeley

(Slides Credit: Dan Klein, Pieter Abbeel, Anca Dragan, Stuart Russell, Satish Rao, Ketrina Yim, and many others)



Kidramat 2012

# Midterm: Topics in Scope


---

- Utilities and Rationality, MEU Principle
- Search and Planning
- Constraint Satisfaction Programming
- Game Trees, Minimax, Pruning, Expectimax
- Probabilistic Inference, Bayesian Networks, Variable Elimination, D-Separation, Sampling
- Markov Models, HMMs



# Agents and environments


---



- An agent *perceives* its environment through *sensors* and *acts* upon it through *actuators* (or *effectors*, depending on whom you ask)
- The *agent function* maps percept sequences to actions
- It is generated by an *agent program* running on a *machine*

# The task environment - PEAS

- Performance measure
  - -1 per step; + 10 food; +500 win; -500 die; +200 hit scared ghost
- Environment
  - Pacman dynamics (incl ghost behavior)
- Actuators
  - Left Right Up Down or NSEW
- Sensors
  - Entire state is visible (except power pellet duration)



# Agent design


---

- **The environment type largely determines the agent design**
  - *Partially observable* => agent requires *memory* (internal state)
  - *Stochastic* => agent may have to prepare for *contingencies*
  - *Multi-agent* => agent may need to behave *randomly*
  - *Static* => agent has time to compute a rational decision
  - *Continuous time* => continuously operating *controller*
  - *Unknown physics* => need for *exploration*
  - *Unknown perf. measure* => observe/interact with *human principal*

# Utilities and Rationality

- Utility: map state of world to real value
- Rational Preferences

Orderability:  $(A > B) \vee (B > A) \vee (A \sim B)$   
Transitivity:  $(A > B) \wedge (B > C) \Rightarrow (A > C)$   
Continuity:  $(A > B > C) \Rightarrow \exists p [p, A; 1-p, C] \sim B$   
Substitutability:  $(A \sim B) \Rightarrow [p, A; 1-p, C] \sim [p, B; 1-p, C]$   
Monotonicity:  $(A > B) \Rightarrow$   
 $(p \geq q) \Leftrightarrow [p, A; 1-p, B] \geq [q, A; 1-q, B]$




Given Rational Preferences, Exists  $U(X)$  s.t.


$$U(A) \geq U(B) \Leftrightarrow A \geq B$$

$$U([p_1, S_1; \dots; p_n, S_n]) = p_1 U(S_1) + \dots + p_n U(S_n)$$

# Maximize Your Expected Utility




# Search Problems




# Search Problems

- A **search problem** consists of:

- A state space




- A successor function  
(with actions, costs)



- A start state and a goal test
- A **solution** is a sequence of actions (a plan) which transforms the start state to a goal state

# State Space Graphs vs. Search Trees


State Space Graph



*Each NODE in  
in the search  
tree is an  
entire PATH in  
the state  
space graph.*

*We construct  
only what we  
need on demand*

Search Tree





# General Tree Search

---

```
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
```

- Important ideas:
  - Fringe
  - Expansion
  - Exploration strategy
- Main question: which fringe nodes to explore?

# Depth-First Search

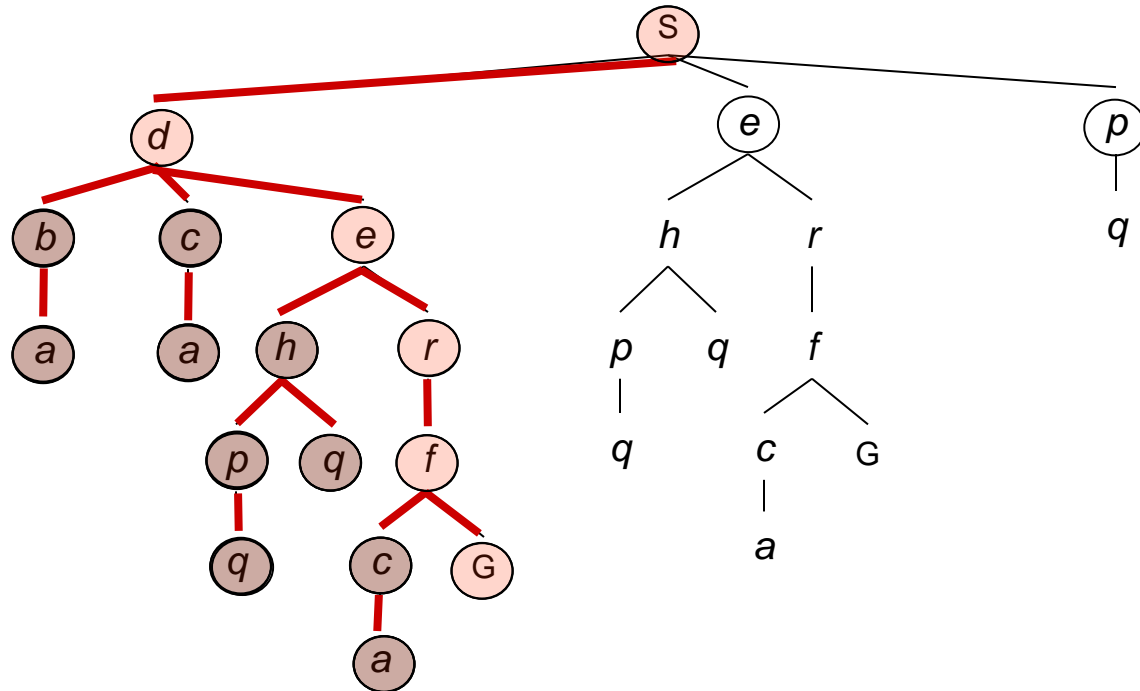

---



# Depth-First Search


Strategy: expand a deepest node first

Implementation: Fringe is a LIFO stack



# Breadth-First Search


---




# Breadth-First Search

Strategy: expand a shallowest node first


Implementation: Fringe is a FIFO queue



Search  
Tiers




# Cost-Sensitive Search



BFS finds the shortest path in terms of number of actions. It does not find the least-cost path. We will now cover a similar algorithm which does find the least-cost path.

# Uniform Cost Search



---



# Uniform Cost Search

Strategy: expand a cheapest node first:




Fringe is a priority queue (priority: cumulative cost)





# Search Heuristics

- A heuristic is:
  - A function that *estimates* how close a state is to a goal
  - Designed for a particular search problem
  - *Pathing?*
  - Examples: Manhattan distance, Euclidean distance




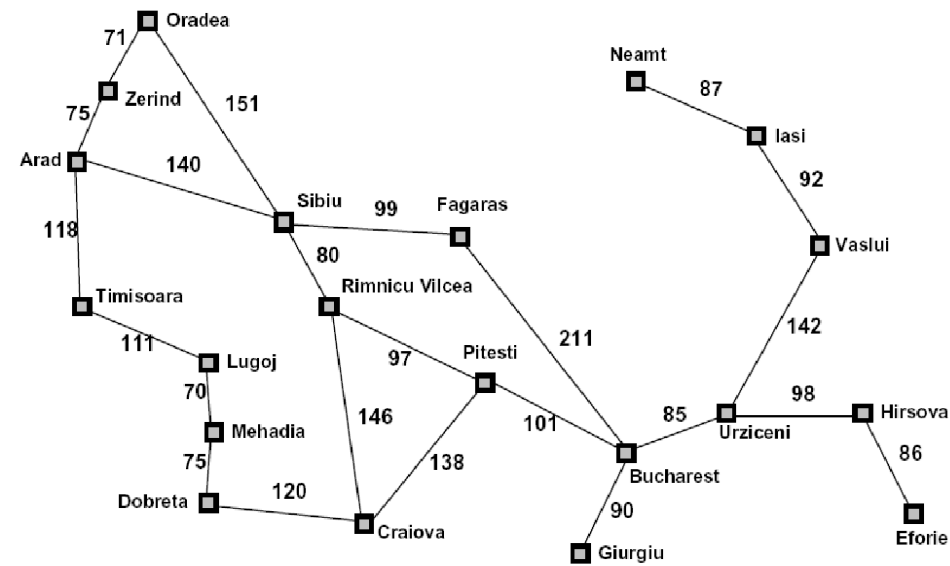
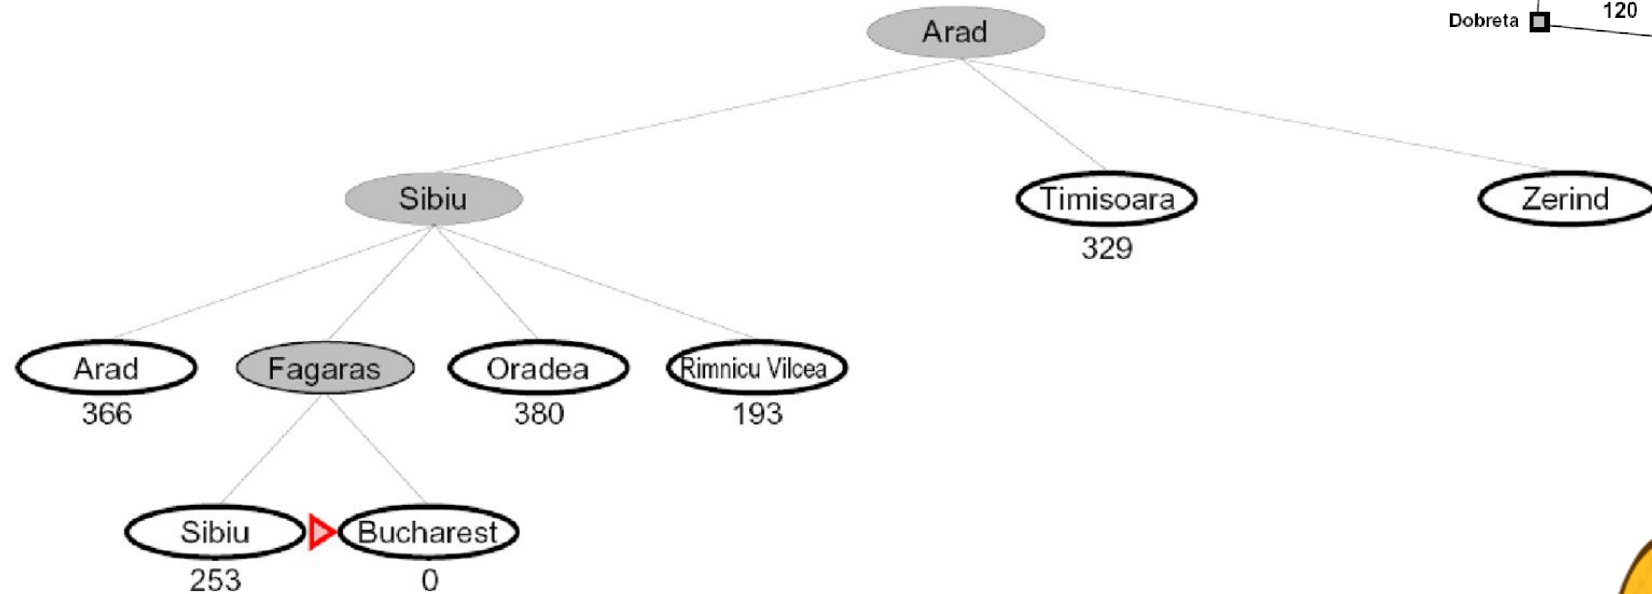
# Greedy Search

---



# Greedy Search

- Expand the node that seems closest...
  - Move to smallest heuristic value



- Is it optimal?
  - No. Resulting path to Bucharest is not the shortest!

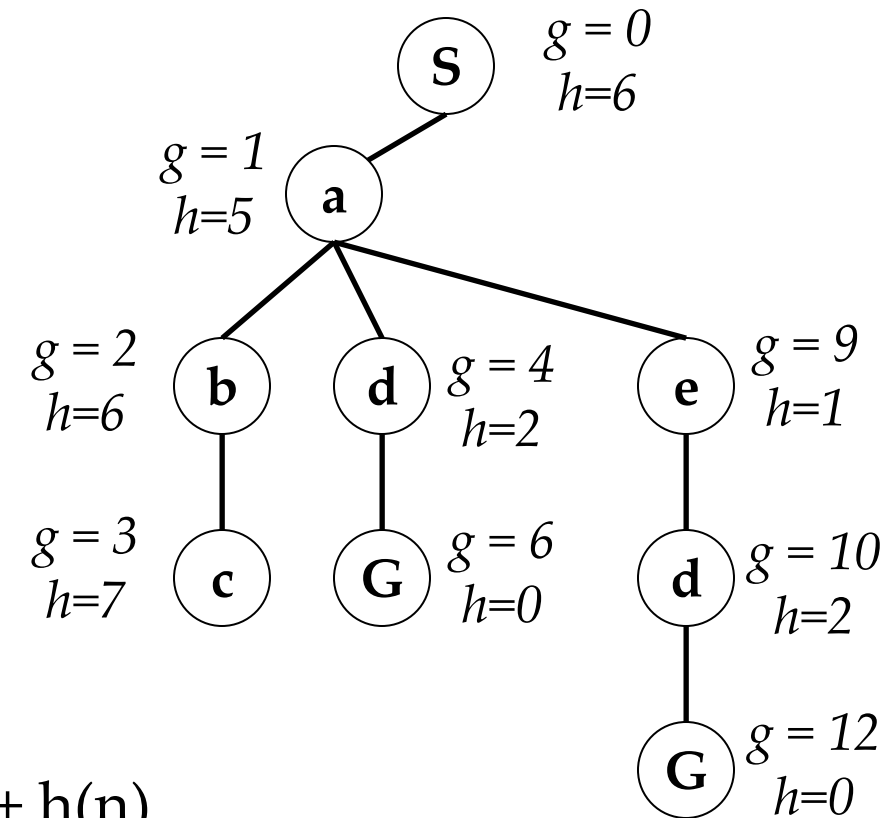
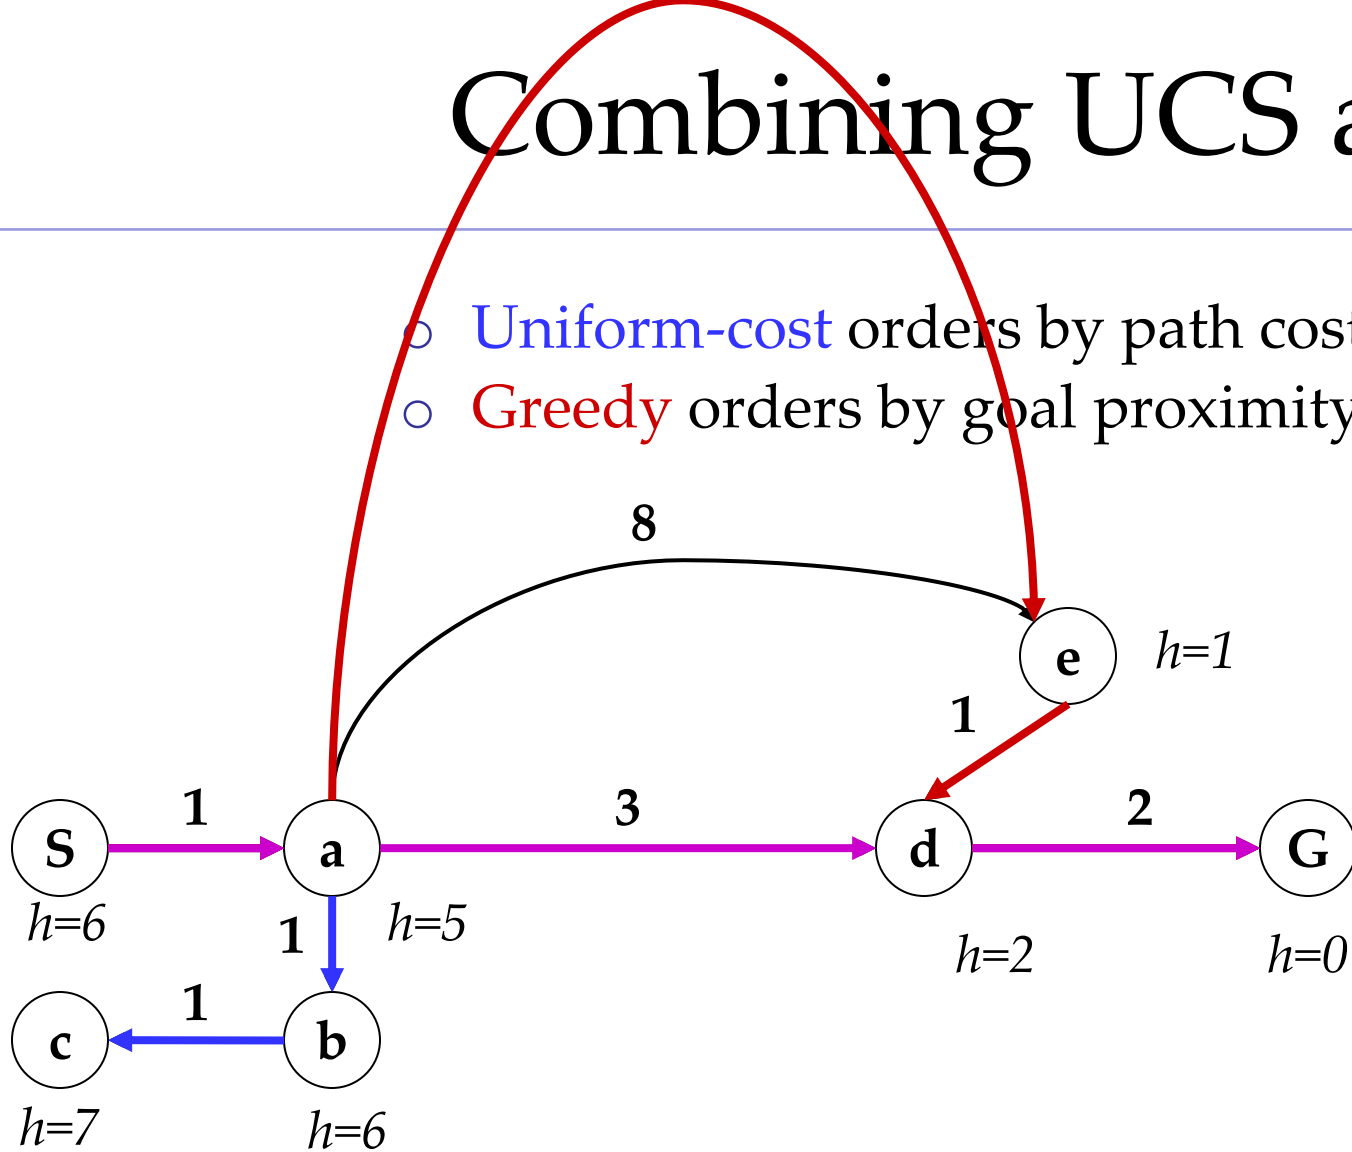
# A\* Search

---



# Combining UCS and Greedy


- **Uniform-cost** orders by path cost, or *backward cost*  $g(n)$
- **Greedy** orders by goal proximity, or *forward cost*  $h(n)$



- **A\* Search** orders by the sum:  $f(n) = g(n) + h(n)$

# When should A\* terminate?

- Should we stop when we enqueue a goal?



	g	h	+
<del>S</del>	<del>0</del>	<del>3</del>	<del>3</del>
<del>S-&gt;A</del>	<del>2</del>	<del>2</del>	<del>4</del>
<del>S-&gt;B</del>	<del>2</del>	<del>1</del>	<del>3</del>
S->B->G	5	0	5
<b>S-&gt;A-&gt;G</b>	<b>4</b>	<b>0</b>	<b>4</b>

- No: only stop when we dequeue a goal



# Admissible Heuristics

- A heuristic  $h$  is *admissible* (optimistic) iff:

$$0 \leq h(n) \leq h^*(n)$$

where  $h^*(n)$  is the true cost to a nearest goal

- Examples:


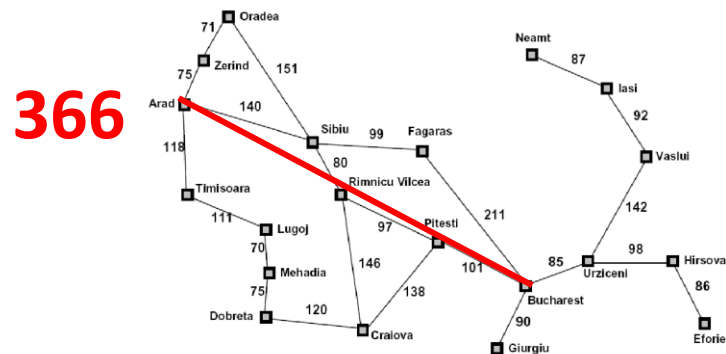


0.0

- Coming up with admissible heuristics is most of what's involved in using  $A^*$  in practice.

# Creating Admissible Heuristics

- Most of the work in solving hard search problems optimally is in coming up with admissible heuristics
- Often, admissible heuristics are solutions to *relaxed problems*, where new actions are available




- Inadmissible heuristics are often useful too



# Graph Search

---




# Graph Search Pseudo-Code

---

```
function GRAPH-SEARCH(problem, fringe) return 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
      for child-node in EXPAND(STATE[node], problem) do
        fringe ← INSERT(child-node, fringe)
      end
    end
  end
```

# Consistency of Heuristics



- Main idea: estimated heuristic costs  $\leq$  actual costs
  - Admissibility: heuristic cost  $\leq$  actual cost to goal
$$h(v) \leq h^*(v) \text{ for all } v \in V$$


Underestimate the true cost to the goal!
  - Consistency: heuristic “arc” cost  $\leq$  actual cost for each arc
$$h(u) - h(v) \leq d(u, v) \text{ for all } (u, v) \in E$$

Underestimate the weight of every edge!
- Consequences of consistency:
  - The f value along a path never decreases
$$h(A) \leq \text{cost}(A \text{ to } C) + h(C)$$
  - A\* graph search is optimal

# Optimality of A\* Search

---

- With a admissible heuristic, Tree A\* is optimal.
- With a consistent heuristic, Graph A\* is optimal.
  - With  $h=0$ , the same proof shows that UCS is optimal.




# Constraint Satisfaction Problems



# Constraint Satisfaction Problems

*N variables*  
*domain D*  
*constraints*



*states*

*partial assignment*

*goal test*


*complete; satisfies constraints*

*successor function*



*assign an unassigned variable*

# Backtracking Search

- Backtracking search is the basic uninformed algorithm for solving CSPs
- Idea 1: One variable at a time
  - Variable assignments are commutative, so fix ordering -> better branching factor!
  - I.e., [WA = red then NT = green] same as [NT = green then WA = red]
  - Only need to consider assignments to a single variable at each step
- Idea 2: Check constraints as you go
  - I.e. consider only values which do not conflict previous assignments
  - Might have to do some computation to check the constraints
  - “Incremental goal test”
- Depth-first search with these two improvements is called *backtracking search* (not the best name)
- Can solve n-queens for  $n \approx 25$



# Backtracking Example



[Demo: coloring -- backtracking]



# Backtracking Search


```
function BACKTRACKING-SEARCH(csp) returns solution/failure
  return RECURSIVE-BACKTRACKING({}, csp)

function RECURSIVE-BACKTRACKING(assignment, csp) returns soln/failure
  if assignment is complete then return assignment
  var ← SELECT-UNASSIGNED-VARIABLE(VARIABLES[csp], assignment, csp)
  for each value in ORDER-DOMAIN-VALUES(var, assignment, csp) do
    if value is consistent with assignment given CONSTRAINTS[csp] then
      add {var = value} to assignment
      result ← RECURSIVE-BACKTRACKING(assignment, csp)
      if result ≠ failure then return result
      remove {var = value} from assignment
  return failure
```

- Backtracking = DFS + variable-ordering + fail-on-violation
- What are the choice points?



# Filtering: Forward Checking

- Filtering: Keep track of domains for unassigned variables and cross off bad options
- Forward checking: Cross off values that violate a constraint when added to the existing assignment



# Filtering: Constraint Propagation




- Forward checking propagates information from assigned to unassigned variables, but doesn't provide early detection for all failures:



- NT and SA cannot both be blue!
- Why didn't we detect this yet?
- Constraint propagation*: reason from constraint to constraint

# Consistency of A Single Arc

- An arc  $X \rightarrow Y$  is **consistent** iff for *every*  $x$  in the tail there is *some*  $y$  in the head which could be assigned without violating a constraint



*Delete from the tail!*

Forward checking?

Enforcing consistency of arcs pointing to each new assignment

# Enforcing Arc Consistency in a CSP

```
function AC-3(csp) returns the CSP, possibly with reduced domains
inputs: csp, a binary CSP with variables  $\{X_1, X_2, \dots, X_n\}$ 
local variables queue, a queue of arcs, initially all the arcs in csp

while queue is not empty do
   $(X_i, X_j) \leftarrow \text{REMOVE-FIRST}(\textit{queue})$ 
  if REMOVE-INCONSISTENT-VALUES( $X_i, X_j$ ) then
    for each  $X_k$  in NEIGHBORS[ $X_i$ ] do
      add  $(X_k, X_i)$  to queue



---





function REMOVE-INCONSISTENT-VALUES( $X_i, X_j$ ) returns true iff succeeds
  removed  $\leftarrow$  false
  for each  $x$  in DOMAIN[ $X_i$ ] do
    if no value  $y$  in DOMAIN[ $X_j$ ] allows  $(x, y)$  to satisfy the constraint  $X_i \leftrightarrow X_j$ 
      then delete  $x$  from DOMAIN[ $X_i$ ]; removed  $\leftarrow$  true
  return removed
```

- Runtime:  $O(n^2d^3)$ , can be reduced to  $O(n^2d^2)$
- ... but detecting all possible future problems is NP-hard – why?

# K-Consistency

- Increasing degrees of consistency
  - 1-Consistency (Node Consistency): Each single node's domain has a value which meets that node's unary constraints
  - 2-Consistency (Arc Consistency): For each pair of nodes, any consistent assignment to one can be extended to the other
  - K-Consistency: For each k nodes, any consistent assignment to k-1 can be extended to the k<sup>th</sup> node.
- Higher k more expensive to compute
- (You need to know the k=2 case: arc consistency)




# Ordering: Minimum Remaining Values

---

- Variable Ordering: Minimum remaining values (MRV):
  - Choose the variable with the fewest legal left values in its domain





- Why min rather than max?
- Also called “most constrained variable”
- “Fail-fast” ordering



# Ordering: Least Constraining Value

- Value Ordering: Least Constraining Value
  - Given a choice of variable, choose the *least constraining value*
  - I.e., the one that rules out the fewest values in the remaining variables
  - Note that it may take some computation to determine this! (E.g., rerunning filtering)
- Why least rather than most?
- Combining these ordering ideas makes 1000 queens feasible







# Iterative Algorithms for CSPs

---

- Local search methods typically work with “complete” states, i.e., all variables assigned
- To apply to CSPs:
  - Take an assignment with unsatisfied constraints
  - Operators *reassign* variable values
  - No fringe! Live on the edge.
- Algorithm: While not solved,
  - Variable selection: randomly select any conflicted variable
  - Value selection: min-conflicts heuristic:
    - Choose a value that violates the fewest constraints
    - I.e., hill climb with  $h(x)$  = total number of violated constraints






# Hill Climbing




# Tree-Structured CSPs


- Algorithm for tree-structured CSPs:
  - Order: Choose a root variable, order variables so that parents precede children



- Remove backward: For  $i = n : 2$ , apply  $\text{RemoveInconsistent}(\text{Parent}(X_i), X_i)$
  - Assign forward: For  $i = 1 : n$ , assign  $X_i$  consistently with  $\text{Parent}(X_i)$
- Runtime:  $O(n d^2)$  (why?)




# Game Playing: Search with other agents

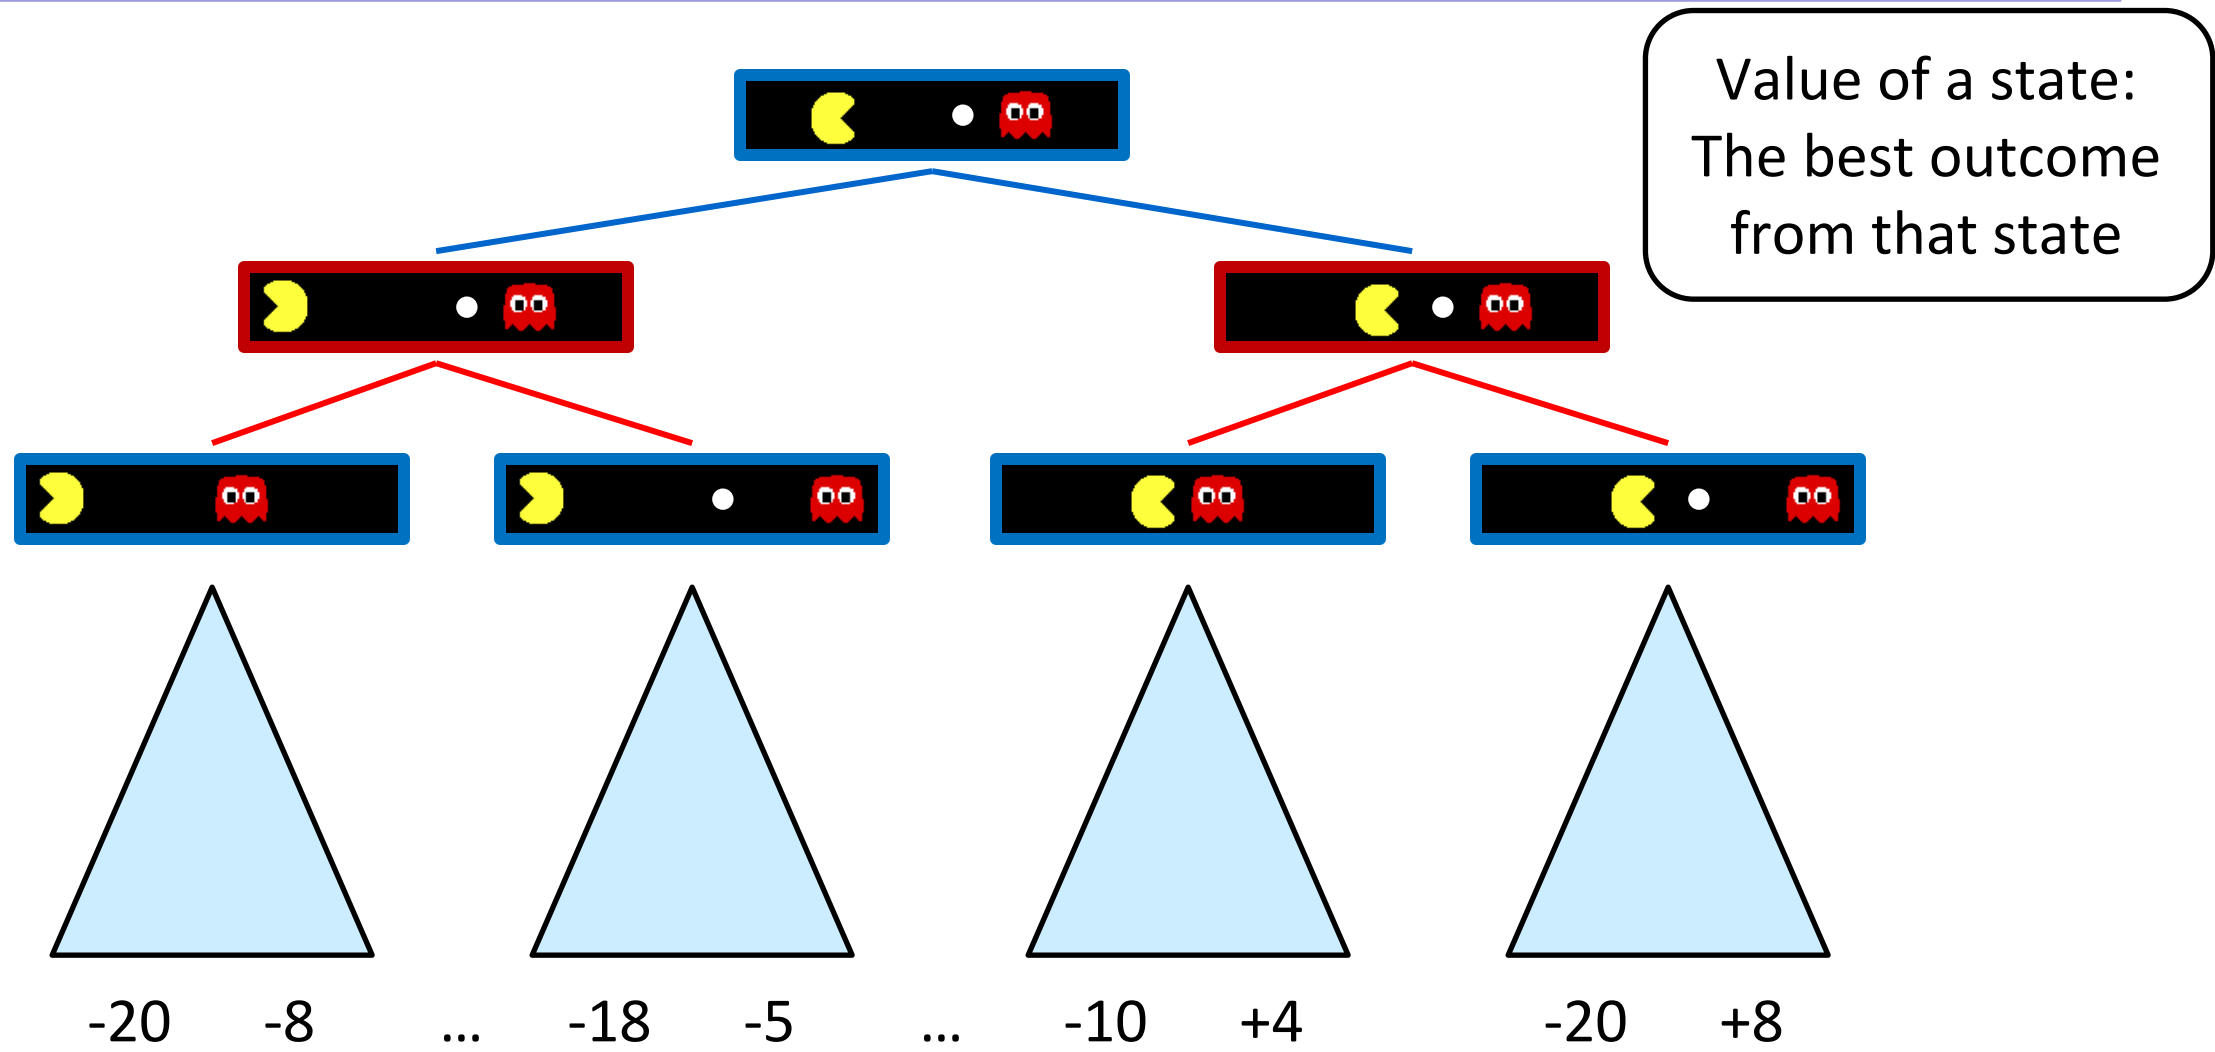


# Adversarial Search

---



# Adversarial Game Trees




# Minimax Values

States Under Agent's Control:

$$V(s) = \max_{s' \in \text{successors}(s)} V(s')$$

States Under Opponent's Control:

$$V(s') = \min_{s \in \text{successors}(s')} V(s)$$



Terminal States:

$$V(s) = \text{known}$$

# Minimax Implementation (Dispatch)

---

def value(state):

if the state is terminal: return the state's utility

if the next agent is MAX: return max-value(state)

if the next agent is MIN: return min-value(state)

def max-value(state):

initialize  $v = -\infty$

for each successor of state:

$v = \max(v, \text{value}(\text{successor}))$

return v


def min-value(state):

initialize  $v = +\infty$

for each successor of state:

$v = \min(v, \text{value}(\text{successor}))$


return v





# Game Tree Pruning

---



# Alpha-Beta Implementation

---

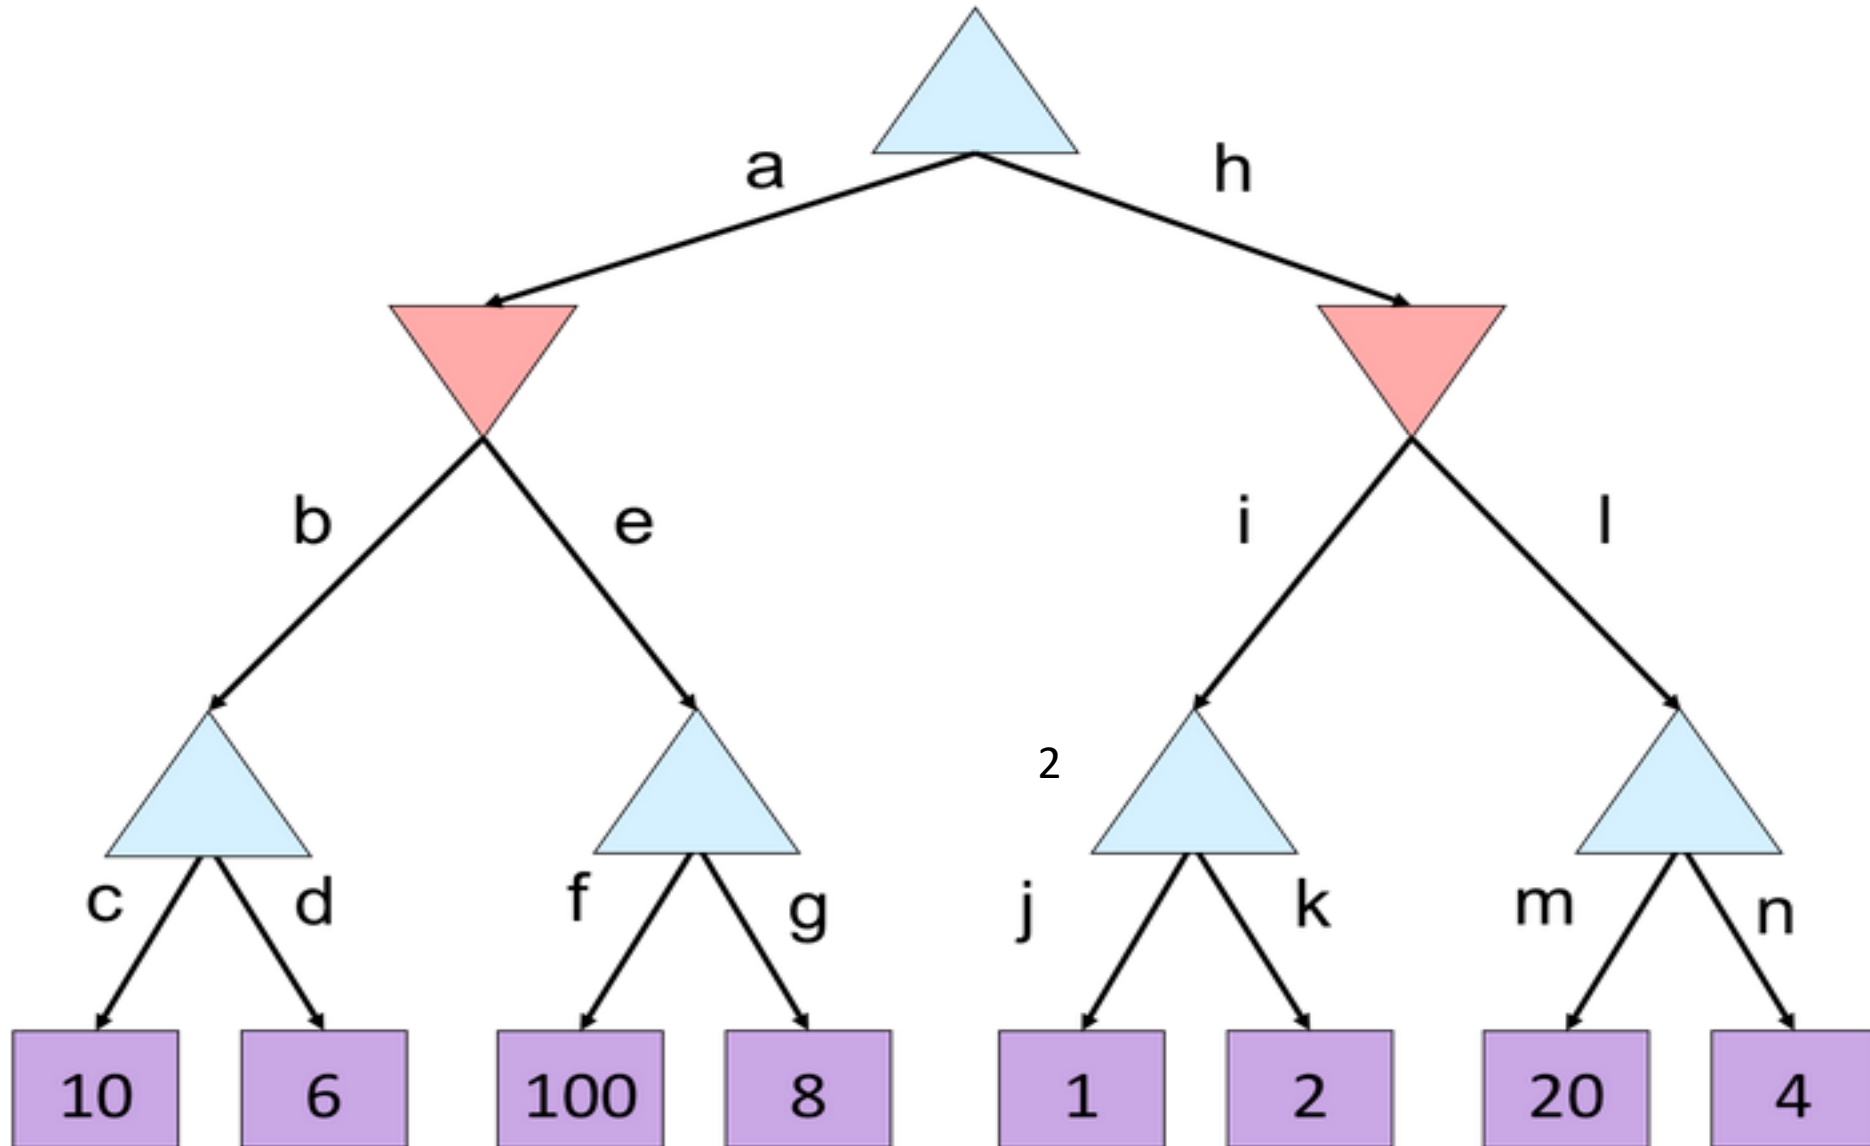
$\alpha$ : MAX's best option on path to root  
 $\beta$ : MIN's best option on path to root

```
def max-value(state,  $\alpha$ ,  $\beta$ ):  
    initialize  $v = -\infty$   
    for each successor of state:  
         $v = \max(v, \text{value}(\text{successor}, \alpha, \beta))$   
        if  $v \geq \beta$  return  $v$   
         $\alpha = \max(\alpha, v)$   
    return  $v$ 
```


```
def min-value(state,  $\alpha$ ,  $\beta$ ):  
    initialize  $v = +\infty$   
    for each successor of state:  
         $v = \min(v, \text{value}(\text{successor}, \alpha, \beta))$   
        if  $v \leq \alpha$  return  $v$   
         $\beta = \min(\beta, v)$   
    return  $v$ 
```

# Alpha-Beta Example


---



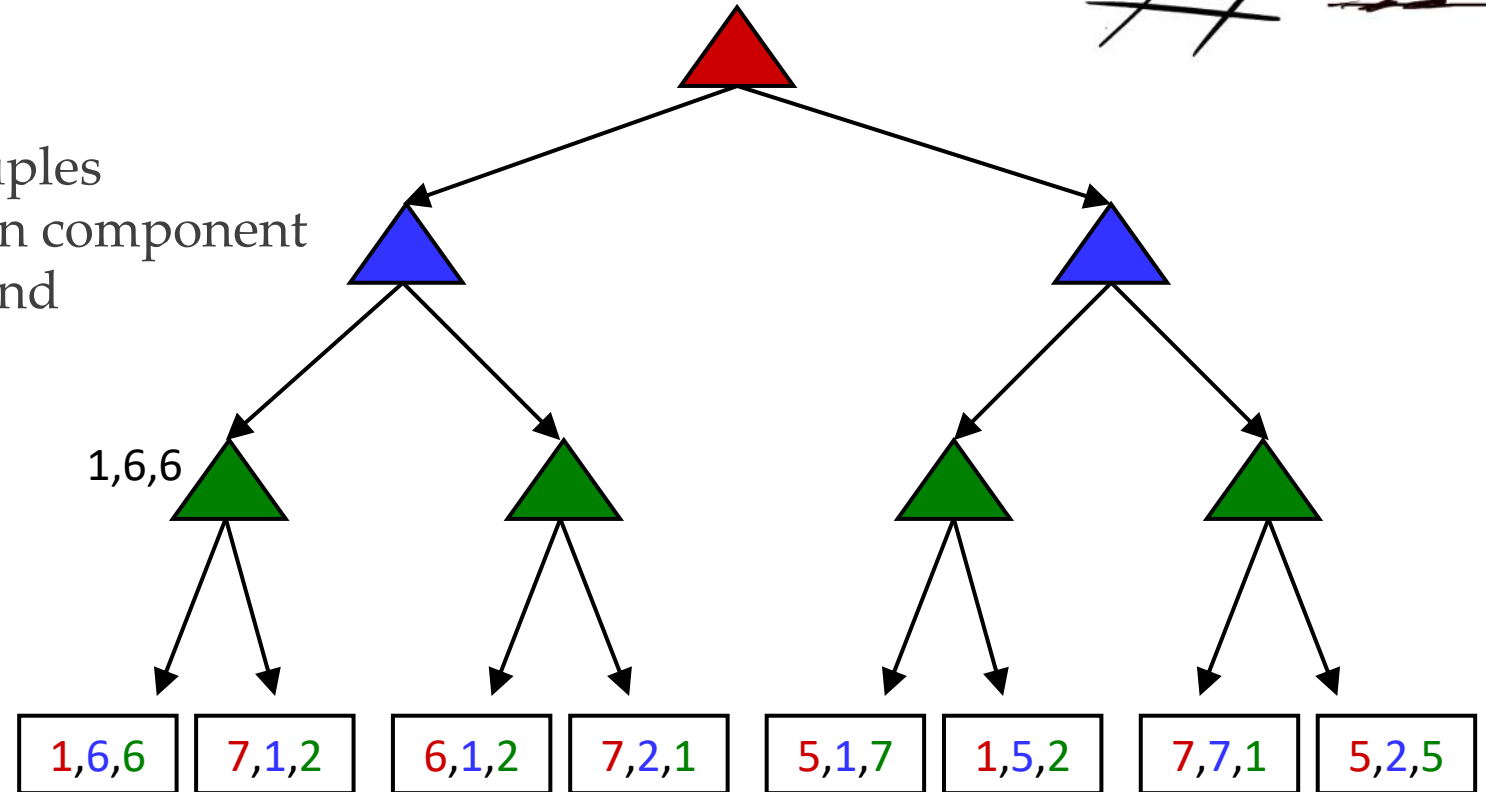

# Alpha-Beta Quiz 2



# Multi-Agent Utilities




- What if the game is not zero-sum, or has multiple players?
- Generalization of minimax:
  - Terminals have utility tuples
  - Node values are also utility tuples
  - Each player maximizes its own component
  - Can give rise to cooperation and competition dynamically...



# Chance Nodes

- We don't know what the result of an action will be:
  - Explicit randomness: rolling dice
  - Unpredictable opponents
  - Actions can fail
- Values should now reflect average-case (expectimax) outcomes, not worst-case (minimax) outcomes
- **Expectimax search**: compute the average score under optimal play
  - Max nodes as in minimax search
  - Chance nodes: calculate **expected utilities**



# Expectimax Pseudocode

```
def value(state):
```

```
    if the state is a terminal state: return the state's utility  
    if the next agent is MAX: return max-value(state)  
    if the next agent is EXP: return exp-value(state)
```


```
def max-value(state):
```

```
    initialize v =  $-\infty$   
    for each successor of state:  
        v = max(v, value(successor))  
    return v
```

```
def exp-value(state):
```


```
    initialize v = 0  
    for each successor of state:  
        p = probability(successor)  
        v += p * value(successor)  
    return v
```

# Bayesian Networks






# Probability



T	W	P
hot	sun	0.4
hot	rain	0.1
cold	sun	0.2
cold	rain	0.3

Summing Out

$$P(h) = P(h, s) + P(h, \sim s)$$




Normalization

$$P(s|h) = \frac{P(s, h)}{P(h, s) + P(h, \sim s)}$$

Bayes' Rule/ Def. of Conditional Probability


$$P(s|h) = \frac{P(s, h)}{P(h)}$$



Chain Rule

$$P(s, h) = P(s|h) * P(h)$$

# Conditional Independence





- $X$  and  $Y$  are **independent** iff

$$\forall x, y \quad P(x, y) = P(x)P(y) \quad X \perp\!\!\!\perp Y$$

- Given  $Z$ , we say  $X$  and  $Y$  are **conditionally independent** iff

$$\forall x, y, z \quad P(x, y|z) = P(x|z)P(y|z) \quad \text{---} \rightarrow \quad X \perp\!\!\!\perp Y|Z$$



- (Conditional) independence is a property of a distribution



# Bayesian Networks


- A directed acyclic graph (DAG), one node per random variable
- A conditional probability table (CPT) for each node
  - Probability of  $X$ , given a combination of values for parents.  
$$P(X|a_1 \dots a_n)$$
- Bayes nets implicitly encode joint distributions as a product of local conditional distributions
  - To see what probability a BN gives to a full assignment, multiply all the relevant conditionals together:

$$P(x_1, x_2, \dots, x_n) = \prod_{i=1}^n P(x_i | \text{parents}(X_i))$$




# Independence Assumptions

- Definition: Each node, given its parents, is conditionally independent of all its non-descendants in the graph



Each node, given its MarkovBlanket, is conditionally independent of all other nodes in the graph



MarkovBlanket refers to the parents, children, and children's other parents.

# Inference by Enumeration

- General case:

- Evidence variables:  $E_1 \dots E_k = e_1 \dots e_k$
  - Query\* variable:  $Q$
  - Hidden variables:  $H_1 \dots H_r$
- }  $X_1, X_2, \dots, X_n$   
All variables



- We want:

$$P(Q|e_1 \dots e_k)$$

- Step 1: Select the entries consistent with the evidence

- Step 2: Sum out H to get joint of Query and evidence

- Step 3: Normalize




$$\times \frac{1}{Z}$$

$$P(Q, e_1 \dots e_k) = \sum_{h_1 \dots h_r} P(Q, \underbrace{h_1 \dots h_r}_{X_1, X_2, \dots, X_n}, e_1 \dots e_k)$$

$$Z = \sum_q P(Q, e_1 \dots e_k)$$

$$P(Q|e_1 \dots e_k) = \frac{1}{Z} P(Q, e_1 \dots e_k)$$

# Traffic Domain



$$P(L) = ?$$

- Inference by Enumeration
- Variable Elimination

$$= \sum_t \left[ \sum_r P(L|t) P(r) P(t|r) \right]$$

Eliminate t

Join on t

Join on r

Eliminate r

$$= \sum_t P(L|t) \left[ \sum_r P(r) P(t|r) \right]$$


Join on t

Eliminate r

Eliminate t

# Marginalizing Early (Variable Elimination)

---



# Variable Elimination



$P(R)$

T	L
+r	0.1
-r	0.9

$P(T|R)$

+r	+t	0.8
+r	-t	0.2
-r	+t	0.1
-r	-t	0.9


$P(L|T)$

+t	+l	0.3
+t	-l	0.7
-t	+l	0.1
-t	-l	0.9

Join R

$P(R, T)$

+r	+t	0.08
+r	-t	0.02
-r	+t	0.09
-r	-t	0.81




$P(L|T)$

+t	+l	0.3
+t	-l	0.7
-t	+l	0.1
-t	-l	0.9

Sum out R

+t	0.17
-t	0.83



$P(L|T)$

+t	+l	0.3
+t	-l	0.7
-t	+l	0.1
-t	-l	0.9

Join T



$P(T, L)$

+t	+l	0.051
+t	-l	0.119
-t	+l	0.083
-t	-l	0.747

Sum out T



+l	0.134
-l	0.866




# General Variable Elimination

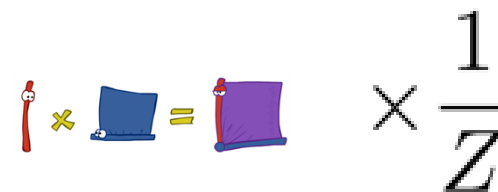
- Query:  $P(Q|E_1 = e_1, \dots, E_k = e_k)$
- Start with initial factors:
  - Local CPTs (but instantiated by evidence)
- While there are still hidden variables (not Q or evidence):
  - Pick a hidden variable H
  - Join all factors mentioning H
  - Eliminate (sum out) H
- Join all remaining factors and normalize



x	P(x)
-3	0.05
-1	0.25
0	0.07
1	0.2
5	0.01

2 0.15




$$\text{stick} \times \text{blue square} = \text{purple square} \times \frac{1}{Z}$$

# Independence Assumptions in a Bayes Net

---

- Assumptions we are required to make to define the Bayes net when given the graph:

$$P(x_i | x_1 \cdots x_{i-1}) = P(x_i | \text{parents}(X_i))$$

- Important for modeling: understand assumptions made when choosing a Bayes net graph



# Active / Inactive Paths

- Question: Are X and Y conditionally independent given evidence variables {Z}?





- Yes, if X and Y “d-separated” by Z
- Consider all (undirected) paths from X to Y
- No active paths = independence!

- A path is active if each triple is active:




- Causal chain  $A \rightarrow B \rightarrow C$  where B is unobserved (either direction)
- Common cause  $A \leftarrow B \rightarrow C$  where B is unobserved
- Common effect (aka v-structure)  
 $A \rightarrow B \leftarrow C$  where B or one of its descendants is observed

- All it takes to block a path is a single inactive segment

Active Triples



Inactive Triples




# D-Separation

- Query:  $X_i \perp\!\!\!\perp X_j \mid \{X_{k_1}, \dots, X_{k_n}\} ?$
- Check all (undirected!) paths between  $X_i$  and  $X_j$ 
  - If one or more active paths, then independence not guaranteed

$$X_i \not\perp\!\!\!\perp X_j \mid \{X_{k_1}, \dots, X_{k_n}\}$$


- Otherwise (i.e. if all paths are inactive), then independence is guaranteed

$$X_i \perp\!\!\!\perp X_j \mid \{X_{k_1}, \dots, X_{k_n}\}$$



# Another Perspective: Bayes Ball


An undirected path is active if a Bayes ball travelling along it never encounters the “stop” symbol:  $\rightarrow \perp$



If there are no active paths from  $X$  to  $Y$  when  $\{Z_1, \dots, Z_k\}$  are shaded, then  $X \perp\!\!\!\perp Y \mid \{Z_1, \dots, Z_k\}$ .


# Topology Limits Distributions

- Given some graph topology  $G$ , only certain joint distributions can be encoded
- The graph structure guarantees certain (conditional) independences
- (There might be more independence)
- Adding arcs increases the set of distributions, but has several costs
- Full conditioning can encode any distribution




# Approximate Inference: Sampling

---



# Prior Sampling


- For  $i = 1, 2, \dots, n$  in topological order
  - Sample  $x_i$  from  $P(X_i \mid \text{Parents}(X_i))$
- Return  $(x_1, x_2, \dots, x_n)$






# Rejection Sampling

- Input: evidence instantiation
- For  $i = 1, 2, \dots, n$  in topological order
  - Sample  $x_i$  from  $P(X_i \mid \text{Parents}(X_i))$
  - If  $x_i$  not consistent with evidence
    - Reject: return – no sample is generated in this cycle
- Return  $(x_1, x_2, \dots, x_n)$



# Likelihood Weighting


- Input: evidence instantiation
- $w = 1.0$
- for  $i = 1, 2, \dots, n$  in topological order
  - if  $X_i$  is an evidence variable
    - $X_i = \text{observation } x_i$  for  $X_i$
    - Set  $w = w * P(x_i \mid \text{Parents}(X_i))$
  - else
    - Sample  $x_i$  from  $P(X_i \mid \text{Parents}(X_i))$
- return  $(x_1, x_2, \dots, x_n), w$



# Gibbs Sampling



- Step 1: Fix evidence

- $R = +r$




- Step 2: Initialize other variables

- Randomly



- Steps 3: Repeat:

- Choose a non-evidence variable  $X$
  - Resample  $X$  from  $P(X \mid \text{MarkovBlanket}(X))$



Sample from  $P(S \mid +c, -w, +r)$

Sample from  $P(C \mid +s, -w, +r)$

Sample from  $P(W \mid +s, +c, +r)$


# Hidden Markov Models



# Markov Chains

- Value of  $X$  at a given time is called the **state**

$P(X_0)$	
sun	rain
1	0.0




$$P(X_1)$$

$$P(X_t|X_{t-1})$$


$$P(X_t) = ?$$

$X_{t-1}$	$X_t$	$P(X_t X_{t-1})$
sun	sun	0.9
sun	rain	0.1
rain	sun	0.3
rain	rain	0.7

State Transition Diagram  
(Flow Graph)




State Trellis



# Mini-Forward Algorithm


- Question: What's  $P(X)$  on some day  $t$ ?



$$P(x_1) = \text{known}$$

$$P(x_t) = \sum_{x_{t-1}} P(x_{t-1}, x_t)$$
$$= \sum_{x_{t-1}} P(x_t \mid x_{t-1}) P(x_{t-1})$$

← *Forward simulation*



# Stationary Distribution

---

- For most chains:

- Influence of the initial distribution gets less and less over time.
- The distribution we end up in is independent of the initial distribution

- Stationary distribution:


- The distribution we end up with is called the **stationary distribution**  $P_\infty$  of the chain
- It satisfies

$$P_\infty(X) = P_{\infty+1}(X) = \sum_x P(X|x)P_\infty(x)$$



# Hidden Markov Models


- Markov chains not so useful for most agents
  - Need observations to update your beliefs
- Hidden Markov models (HMMs)
  - Underlying Markov chain over states  $X_i$
  - You observe outputs (effects) at each time step






# Inference tasks


Prediction:  $P(X_{t+k} | e_{1:t})$




Filtering:  $P(X_t | e_{1:t})$



Smoothing:  $P(X_k | e_{1:t}), k < t$




Explanation:  $P(X_{1:t} | e_{1:t})$




# Inference: Find State Given Evidence

---

- We are given evidence at each time and want to know  $P(X_t | e_{1:t})$
- Idea: start with  $P(X_1)$  and derive  $P(X_t | e_{1:t})$  in terms of  $P(X_{t-1} | e_{1:t-1})$
- Two steps: Passage of time + Incorporate Evidence





Kidramat 2012