

CSC384:Lecture6

- Lasttime
 - Costs,Heuristics,LCFS,BeFS,A*,IDS
- Today
 - wrapuplastclassslides(Misc Searchtopics)
 - gametreesearch
- Readings:
 - Today:none(youshoulduselectureslides)
 - Nextweek:Ch.8.1,8.2(STRIPS,skimSituation Calculus, **skip** EventCalculus),8.3(uptoSTRIPS planning)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

1

GeneralizingSearchProblems

- Sofar:oursearchproblemshaveassumed agenthascompletecontrolofenvironment
 - agent(courier,robot)hasthe **only** effectonstate
 - straightpathtogoalstateisreasonable
- Assumptionnotalwaysreasonable
 - stochasticenvironment
 - otheragentswhoseinterestsconflictwithyours
- Inthesecases,weneedtogeneralizeourview ofsearchtohandle“uncontrollable”statechange

CSC 384 Lecture Slides (c) 2002, C. Boutilier

2

Two-personZero-SumGames

- **Two-person,zero-sumgames** anextremecase
 - chess,checkers,tic-tac-toe,backgammon,go, Command&Conquer:Renegade,NewAdventures withPooh,travelagents,“findthelastparkingspace”
 - youwanttogetsomewhere(winningposition)and opponentwantsyoutoendupsomewhereelse (losingposition)
- Keyinsight:
 - howyouactdependsonhowotheragentacts(or howyouthinktheywillact)
 - andviceversa

CSC 384 Lecture Slides (c) 2002, C. Boutilier

3

MoreGeneralGames

- Whatmakesomethingagame:
 - therearetwo(ormore)agentsinfluencingstatechange
 - eachagenthastheirrowninterests
 - e.g.,goalstatesaredifferent;orweassigndifferent valuestodifferentpaths/states
- Whatmakesgameshard?
 - howyoushouldplaydependsonhowyouthinkthe otherpersonwillplay;buttheyplaydependson howtheythinkyouwillplay;sohowyoushouldplay dependsonhowyouthinktheythinkyouwillplay;but howtheyplaysoulddependonhowtheythinkyou thinktheythinkyouwillplay;...

CSC 384 Lecture Slides (c) 2002, C. Boutilier

4

MoreGeneralGames

- Zero-sumgamesare“fullycompetitive”
 - ifoneplayerwins,theotherplayerloses
 - e.g.,theamtomoneylwin(lose)atpokeristhe amountofmoneyyoulose(win)
- Moregeneralgamesare“cooperative”
 - someoutcomesarepreferredbybothofus,orat leastourvaluesaren'tdiametricallyopposed
- We'lllookindetailatzero-sumgames
 - butfirst,acouplesimplezero-sumandcooperative gamesforfun

CSC 384 Lecture Slides (c) 2002, C. Boutilier

5

Game1:Rock,PaperScissors

- Scissorscutpaper, papercoversrock,rock smashesscissors
- Representedasa matrix:PlayerIchooses arow,PlayerIIchooses acolumn
- Payofftoeachplayerin eachcell(**Pl.I / Pl.II**)
- 1:win,0:tie,-1:loss
 - soit'szero-sum

		Player II		
		R	P	S
Player I	R	0/0	-1/1	1/-1
	P	1/-1	0/0	-1/1
	S	-1/1	1/-1	0/0

CSC 384 Lecture Slides (c) 2002, C. Boutilier

6

Game 2: Prisoner's Dilemma

- Two prisoner's in separate cells, DA doesn't have enough evidence to convict them
- If one confesses, other doesn't:
 - confessor goes free
 - other sentenced to 4 years
- If both confess (both defect)
 - both sentenced to 3 years
- Neither confess (both cooperate)
 - sentenced to 1 year on minor charge
- Payoff: 4 minus sentence

	C	D
C	3/3	0/4
D	4/0	1/1

CSC 384 Lecture Slides (c) 2002, C. Boutilier

7

Game 3: Battlebots

- Two robots: Blue (Craig's), Red (Fahiem's)
 - one cup of coffee, one tea left
 - both C, F prefer coffee (value 10)
 - tea acceptable (value 8)
- Both robot's go for Cof
 - collide and get no payoff
- Both go for tea: same
- One goes for coffee, other for tea:
 - coffee robot gets 10
 - tea robot gets 8

	C	T
C	0/0	10/8
T	8/10	0/0



CSC 384 Lecture Slides (c) 2002, C. Boutilier

8

Two Player Zero Sum Games

- Key point of previous games: what you should do depends on what other guy does
- Previous games are simple "one shot" games
 - single move each
 - in game theory: *strategic or normal form games*
- Many games extend over multiple moves
 - e.g., chess, checkers, etc.
 - in game theory: *extensive form games*
- We'll focus on the extensive form
 - that's where the computational questions emerge

CSC 384 Lecture Slides (c) 2002, C. Boutilier

9

Two-Player, Zero-Sum Game: Defn

- Two *players* A (Max) and B (Min)
- set of *positions* P (states of game)
- a *starting position* $s \in P$ (where game begins)
- terminal positions* $T \subseteq P$ (where game can end)
- set of directed edges $E_A \subseteq P \times P$ (A's *moves*)
- set of directed edges $E_B \subseteq P \times P$ (B's *moves*)
- utility* or *payoff function* $U : T \rightarrow \mathbb{R}$

CSC 384 Lecture Slides (c) 2002, C. Boutilier

10


Intuitions

- Players alternate moves (starting with Max)
 - Game ends when some terminal $p \in T$ is reached
- A game *state*: a position-player pair
 - tells us what position we're in, whose move it is
- Utility function and terminals replace goals
 - Max wants to maximize the terminal payoff
 - Min wants to minimize the terminal payoff
- Think of it as:
 - Max gets $U(t)$, Min gets $-U(t)$ for terminal node t
 - This is why it's called zero (or constant) sum

CSC 384 Lecture Slides (c) 2002, C. Boutilier

11


Tic-tac-toe: States



CSC 384 Lecture Slides (c) 2002, C. Boutilier

12

Tic-tac-toe: Game Tree



CSC 384 Lecture Slides (c) 2002, C. Boutilier

13

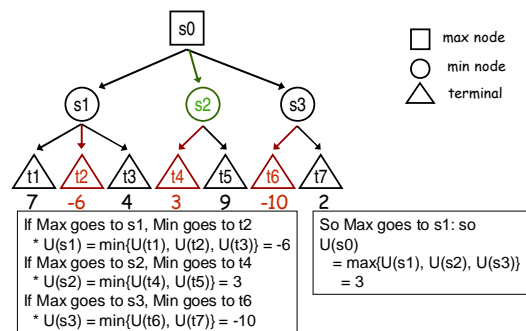
Game Tree

- Game tree looks like a search tree
 - Layers reflect the alternating moves
- But Max doesn't decide where to go alone
 - after Max moves to state **a**, Mins decides whether to move to state **b**, **c**, or **d**
- Thus Max must have a **strategy**
 - must know what to do next no matter what move Min makes (**b**, **c**, or **d**)
 - a sequence of moves will not suffice: Max may want to do something different in response to **b**, **c**, or **d**
- What is a **reasonable strategy**?

CSC 384 Lecture Slides (c) 2002, C. Boutilier

14

Minimax Strategy: Intuitions



CSC 384 Lecture Slides (c) 2002, C. Boutilier

15

Minimax Strategy

- Build full game tree (all leaves are terminals)
 - root is start state, edges are possible moves, etc.
 - label terminal nodes with utilities
- Back values **up** the tree
 - $U(t)$ is defined for all terminals (part of input)
 - $U(n) = \min\{U(c) : c \text{ a child of } n\}$ if n is a min node
 - $U(n) = \max\{U(c) : c \text{ a child of } n\}$ if n is a max node
- Max chooses, at any (max) state, the action that leads to the highest utility child

CSC 384 Lecture Slides (c) 2002, C. Boutilier

16

Depth-first Implementation on MMX

```

utility(N,L,U) :- terminal(L), utility(N,U).
utility(N,L,U) :- maxLevel(L, children(N,CList),
    maxValue(CList, U).
utility(N,L,U) :- minLevel(L, children(N,CList),
    minValue(CList, U).
    
```

- Depth-first evaluation of game tree
 - $\text{maxLevel}(L)$ holds if level L is Max's move; similarly for $\text{minLevel}(L)$ and $\text{terminal}(L)$.
 - utility of terminals must be specified as input
 - if game isn't strictly alternating, then move might be associated with the node (or you might need to record whose move it is explicitly)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

17

Depth-first Implementation of MMX

```

utility(N,L,U) :- maxLevel(L, children(N,CList),
    maxValue(CList, U).
utility(N,L,U) :- minLevel(L, children(N,CList),
    minValue(CList, U).
    
```

- maxValue recursively calls utility on each child and sets U to max of all values over these children (ditto for minValue)
 - note: can't return until all children evaluated
- Notice that the game tree has to have finite depth for this to work
- Advantage of DF implementation: space efficient

CSC 384 Lecture Slides (c) 2002, C. Boutilier

18

Once s17 eval'd, no need to store tree: s16 only needs its value.
Once s24 value computed, we can evaluate s16

```
graph TD; s0[s0] --> s1((s1)); s0 --> s13((s13)); s0 --> s16((s16)); s1 --> s2[s2]; s1 --> s6[s6]; s2 --> t3[t3]; s2 --> t4[t4]; s2 --> t5[t5]; s6 --> s7((s7)); s6 --> s10((s10)); s7 --> t18[t18]; s7 --> t19[t19]; s10 --> t11[t11]; s10 --> t12[t12]; s10 --> t13[t13]; s13 --> t14[t14]; s13 --> t15[t15]; s16 --> s17[s17]; s16 --> s24[s24]; s17 --> s18((s18)); s17 --> s21((s21)); s18 --> t19[t19]; s18 --> t20[t20]; s18 --> t21[t21]; s21 --> t22[t22]; s21 --> t23[t23]; s24 --> t25[t25]; s24 --> t26[t26];
```

Once s17 eval'd, no need to store tree: s16 only needs its value.
Once s24 value computed, we can evaluate s16

19

- It is not necessary to examine entire tree to make correct minimax decision
- Assume depth-first generation of tree
 - After generating value for only *some* of node n 's children we can prove that we'll never reach n
 - So we needn't generate or evaluate any further children of n !
- Two types of cuts:
 - α -cuts: pruning of max nodes
 - β -cuts: pruning of min nodes

CSC 384 Lecture Slides (c) 2002, C. Boutilier


20

- At a Max node n :
 - α is best value of n 's children examined so far (dynamic: changes as children examined)
 - β is worst value of n 's parent's children examined so far (fixed when evaluating n)
- At a Min node n :
 - β is worst value of n 's children examined so far (dynamic: changes as children examined)
 - α is best value of n 's parent's children examined so far (fixed when evaluating n)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

21


- Once alpha-value at max node n reaches its beta-value (i.e., its parent P 's beta-value), we can stop expansion of n
 - Min will never choose to move to n from P since it can guarantee the lower beta-value already



CSC 384 Lecture Slides (c) 2002, C. Boutilier

22

- Once beta-value at min node n reaches its alpha-value (i.e., its parent P 's alpha-value), we can stop expansion of n
 - Max will never choose to move to n from P since it can guarantee the greater alpha-value already



CSC 384 Lecture Slides id 2002, C. Boutilier

23

Pseudo-code algorithm that associates a value with each node. Strategy extracted by moving to Max node (if you are player Max) at each step.

```
Evaluate(startNode):
/* assume Max moves first */
MavEval(start, -infnty, +infnty)
```

```

MaxEval(node, alpha, beta):
  If terminal(node), return U(n)
  For each c in childlist(n)
    val ← MinEval(c, alpha, beta)
    alpha ← max(alpha, val)
    If alpha ≥ beta, return beta
  Return alpha

```

```
MinEval(node, alpha, beta):
  If terminal(node), return U(n)
  For each c in childlist(n)
    val ← MaxEval(c, alpha, beta)
    beta ← min(beta, val)
  If alpha ≥ beta, return alpha
  Return beta
```

CSC 384 Lecture Slides (c) 2002, C. Boutilier

24

Rational Opponents

- This all assumes that your opponent is rational
 - e.g., will choose moves that minimize your max score
- Storing your strategy is a potential issue:
 - you must store "decisions" for each node you can reach by playing optimally
 - if your opponent has unique rational choices, this is a single branch through game tree
 - if there are "ties", opponent could choose any one of the "tied" moves: must store strategy for each subtree
- What if your opponent doesn't play rationally?
 - will it affect quality of outcome?
 - what to do if you haven't stored a full strategy?

CSC 384 Lecture Slides (c) 2002, C. Boutilier

25

Practical Matters

- All "real" games are too large to enumerate tree
 - e.g., chess branching factor is roughly 35
 - Depth 10 tree: 2,700,000,000,000,000 nodes
 - Yikes! Even alpha-beta pruning won't help here!
- We must limit depth of search tree
 - can't expand all the way to terminal nodes
 - we must make *heuristic estimates* about the values of the (nonterminal) states at the leaves of the tree
 - *evaluation function* is an often used term
 - evaluation functions are often learned
- Depth-first expansion almost always used for game trees because of sheer size of trees

CSC 384 Lecture Slides (c) 2002, C. Boutilier

26

Heuristics

- Think of a few games and suggest some heuristics for estimating the "goodness" of a position
 - chess?
 - checkers?
 - your favorite video game?
 - "find the last parking spot"?

CSC 384 Lecture Slides (c) 2002, C. Boutilier

27

Some Interesting Games

- Tesauro's TD-Gammon
 - champion backgammon player which learned evaluation function; stochastic component (dice)
- Checker's (Samuel, 1950s; Schaeffer)
- Chess (which you all know about)
- Bridge, Poker, etc.
- Check out Jonathan Schaeffer's Web page:
 - www.cs.ualberta.ca/~games
 - they've studied lots of games (you can play too)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

28


An Aside on Large Search Problems

- Issue: inability to expand tree to terminal nodes is relevant even in standard search
 - often we can't expect A* to reach a goal by expanding full frontier
 - so we often limit our lookahead, and make moves before we actually know the true path to the goal
 - sometimes called *online* or *realtime* search
- In this case, we use the heuristic function not just to guide our search, but also commits us to moves we actually make
 - in general, guarantees of optimality are lost, but we reduce computational/memory expense dramatically

CSC 384 Lecture Slides (c) 2002, C. Boutilier

29

Realtime Search Graphically



1. We run A* (or our favorite search algorithm) until we are forced to make a move or run out of memory. Note: no leaves are goals yet.
2. We use evaluation function $f(n)$ to decide which path *looks* best (let's say it is the *red* one).
3. We take the first step along the best path (red), by actually *making that move*.
4. We restart search at the node we reach by making that move. (We may actually cache the results of the relevant part of first search tree if it's hanging around, as it would with A*).

CSC 384 Lecture Slides (c) 2002, C. Boutilier

30