

CSE 573: Artificial Intelligence

Autumn 2010

Lecture 4: Adversarial Search

10/12/2009

Luke Zettlemoyer

Based on slides from Dan Klein

Many slides over the course adapted from either Stuart Russell
or Andrew Moore

Announcements

- PS 1 due on Friday
 - Will submit via dropbox, instructions will be posted soon
- PS 2 will be out shortly after

Today

- Proof of Graph Search A* Optimality
- Start Adversarial Search
 - Minimax search
 - α - β search
 - Evaluation functions
 - Expectimax


Recap: Graph Search

- **Search problem:**
 - States (configurations of the world)
 - Successor function: a function from states to lists of (state, action, cost) triples; drawn as a graph
 - Start state and goal test
- **Search tree:**
 - Nodes: represent plans for reaching states
 - Plans have costs (sum of action costs)
- **Search Algorithm:**
 - Systematically builds a search tree
 - Chooses an ordering of the fringe (unexplored nodes)
 - **Graph Search:** only expand each state once

Optimality of A* Graph Search


Proof:

- Main idea: Argue that nodes are popped with non-decreasing f-scores
 - for all n, n' with n' popped after n :
 - $f(n') \geq f(n)$
 - is this enough for optimality?
- Sketch:
 - **assume:** $f(n') \geq f(n)$, for all edges (n, a, n') and all actions a
 - is this true?
 - **proof by induction:** (1) always pop the lowest f-score from the fringe, (2) all new nodes have larger (or equal) scores, (3) add them to the fringe, (4) repeat!



Consistency

- Wait, how do we know parents have better f-values than their successors?



- Consistency** for all edges (n,a,n') :
 - $h(n) \leq c(n,a,n') + h(n')$
- Proof that $f(n') \geq f(n)$,**
 - $f(n') = g(n') + h(n') = g(n) + c(n,a,n') + h(n') \geq g(n) + h(n) = f(n)$

Optimality

- **Tree search:**
 - A* optimal if heuristic is admissible (and non-negative)
 - UCS is a special case ($h = 0$)
- **Graph search:**
 - A* optimal if heuristic is consistent
 - UCS optimal ($h = 0$ is consistent)
- **Consistency implies admissibility**
- **In general, natural admissible heuristics tend to be consistent**

Game Playing State-of-the-Art

- **Checkers:** Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994. Used an endgame database defining perfect play for all positions involving 8 or fewer pieces on the board, a total of 443,748,401,247 positions. Checkers is now solved!
- **Chess:** Deep Blue defeated human world champion Gary Kasparov in a six-game match in 1997. Deep Blue examined 200 million positions per second, used very sophisticated evaluation and undisclosed methods for extending some lines of search up to 40 ply. Current programs are even better, if less historic.
- **Othello:** Human champions refuse to compete against computers, which are too good.
- **Go:** Human champions are beginning to be challenged by machines, though the best humans still beat the best machines. In go, $b > 300$, so most programs use pattern knowledge bases to suggest plausible moves, along with aggressive pruning.
- **Pacman:** unknown

General Game Playing

The IJCAI-09 Workshop on General Game Playing General Intelligence in Game Playing Agents (GIGA'09) Pasadena, CA, USA

Workshop Organizers

Yngvi Björnsson
School of Computer Science
Reykjavik University

Peter Stone
Department of Computer Sciences
University of Texas at Austin

Michael Thielscher
Department of Computer Science
Dresden University of Technology

Program Committee


Yngvi Björnsson,
Reykjavik University

Patrick Doherty,
Linköping University

Artificial Intelligence (AI) researchers have for decades worked on building game-playing agents capable of matching wits with the strongest humans in the world, resulting in several success stories for games like e.g. chess and checkers. The success of such systems has been for a part due to years of relentless knowledge-engineering effort on behalf of the program developers, manually adding application-dependent knowledge to their game-playing agents. Also, the various algorithmic enhancements used are often highly tailored towards the game at hand.

Research into general game playing (GGP) aims at taking this approach to the next level: to build intelligent software agents that can, given the rules of any game, automatically learn a strategy for playing that game at an expert level without any human intervention. On contrary to software systems designed to play one specific game, systems capable of playing arbitrary unseen games cannot be provided with game-specific domain knowledge a priori. Instead they must be endowed with high-level abilities to learn strategies and make abstract reasoning. Successful realization of this poses many interesting research challenges for a wide variety of artificial-intelligence sub areas including (but not limited to):

Adversarial Search



Game Playing


- Many different kinds of games!
- Axes:
 - Deterministic or stochastic?
 - One, two, or more players?
 - Perfect information (can you see the state)?
- Want algorithms for calculating a **strategy** (**policy**) which recommends a move in each state

Deterministic Games


- Many possible formalizations, one is:
 - States: S (start at s_0)
 - Players: $P=\{1\dots N\}$ (usually take turns)
 - Actions: A (may depend on player / state)
 - Transition Function: $S \times A \rightarrow S$
 - Terminal Test: $S \rightarrow \{t,f\}$
 - Terminal Utilities: $S \times P \rightarrow R$
- Solution for a player is a policy: $S \rightarrow A$

Deterministic Two-Player


- E.g. tic-tac-toe, chess, checkers
- Zero-sum games
 - One player maximizes result
 - The other minimizes result
- **Minimax search**
 - A state-space search tree
 - Players alternate
 - Choose move to position with highest **minimax value** = best achievable utility against best play



Tic-tac-toe Game Tree



Minimax Example




Minimax Search

function MAX-VALUE(*state*) *returns a utility value*
 if TERMINAL-TEST(*state*) **then return** UTILITY(*state*)
 $v \leftarrow -\infty$
 for *a, s* in SUCCESSORS(*state*) **do** $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s))$
 return *v*


function MIN-VALUE(*state*) *returns a utility value*
 if TERMINAL-TEST(*state*) **then return** UTILITY(*state*)
 $v \leftarrow \infty$
 for *a, s* in SUCCESSORS(*state*) **do** $v \leftarrow \text{MIN}(v, \text{MAX-VALUE}(s))$
 return *v*

Minimax Properties


- Optimal against a perfect player. Otherwise?
- Time complexity?
 - $O(b^m)$
- Space complexity?
 - $O(bm)$
- For chess, $b \approx 35$, $m \approx 100$
 - Exact solution is completely infeasible
 - But, do we need to explore the whole tree?



α - β Pruning Example




α - β Pruning Example



α - β Pruning

- General configuration
 - α is the best value that MAX can get at any choice point along the current path
 - If n becomes worse than α , MAX will avoid it, so can stop considering n 's other children
 - Define β similarly for MIN




α - β Pruning Pseudocode

```
function MAX-VALUE(state) returns a utility value
  if TERMINAL-TEST(state) then return UTILITY(state)
   $v \leftarrow -\infty$ 
  for  $a, s$  in SUCCESSORS(state) do  $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s))$ 
  return  $v$ 
```

```
function MAX-VALUE(state,  $\alpha$ ,  $\beta$ ) returns a utility value
  inputs: state, current state in game
          $\alpha$ , the value of the best alternative for MAX along the path to state
          $\beta$ , the value of the best alternative for MIN along the path to state

  if TERMINAL-TEST(state) then return UTILITY(state)
   $v \leftarrow -\infty$ 
  for  $a, s$  in SUCCESSORS(state) do
     $v \leftarrow \text{MAX}(v, \text{MIN-VALUE}(s, \alpha, \beta))$ 
    if  $v \geq \beta$  then return  $v$ 
     $\alpha \leftarrow \text{MAX}(\alpha, v)$ 
  return  $v$ 
```




α - β Pruning Properties

- Pruning has **no effect** on final result
- Good move ordering improves effectiveness of pruning
- With “perfect ordering”:
 - Time complexity drops to $O(b^{m/2})$
 - Doubles solvable depth
 - Full search of, e.g. chess, is still hopeless!


Resource Limits

- Cannot search to leaves
- Depth-limited search
 - Instead, search a limited depth of tree
 - Replace terminal utilities with an eval function for non-terminal positions
- Guarantee of optimal play is gone
- Example:
 - Suppose we have 100 seconds, can explore 10K nodes / sec
 - So can check 1M nodes per move
 - α - β reaches about depth 8 – decent chess program



Evaluation Functions

- Function which scores non-terminals



$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$

- Ideal function: returns the utility of the position
- In practice: typically weighted linear sum of features:
 - e.g. $f_1(s) = (\text{num white queens} - \text{num black queens})$, etc.

Evaluation for Pacman




What features would be good for Pacman?

$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$


Which algorithm?

α - β , depth 4, simple eval fun



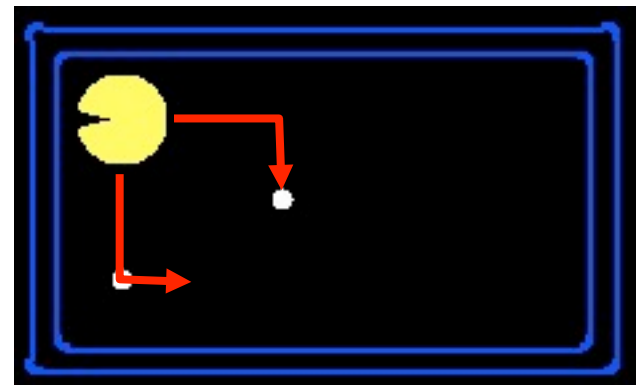

Which algorithm?

α - β , depth 4, better eval fun



Why Pacman Starves

- He knows his score will go up by eating the dot now
- He knows his score will go up just as much by eating the dot later on
- There are no point-scoring opportunities after eating the dot
- Therefore, waiting seems just as good as eating




Iterative Deepening

Iterative deepening uses DFS as a subroutine:

1. Do a DFS which only searches for paths of length 1 or less. (DFS gives up on any path of length 2)
2. If “1” failed, do a DFS which only searches paths of length 2 or less.
3. If “2” failed, do a DFS which only searches paths of length 3 or less.
.....and so on.


Why do we want to do this for multiplayer games?



Multi-Player Non-Zero-Sum Games


- Similar to minimax:

- Utilities are now tuples
- Each player maximizes their own entry at each node
- Propagate (or back up) nodes from children




Stochastic Single-Player

- What if we don't know what the result of an action will be? E.g.,
 - In solitaire, shuffle is unknown
 - In minesweeper, mine locations
- Can do **expectimax search**
 - Chance nodes, like actions except the environment controls the action chosen
 - Max nodes as before
 - Chance nodes take average (expectation) of value of children



Which Algorithms?

Expectimax




Minimax



3 ply look ahead, ghosts move randomly

Stochastic Two-Player

- E.g. backgammon
- Expectiminimax (!)
 - Environment is an extra player that moves after each agent
 - Chance nodes take expectations, otherwise like minimax



if state is a MAX node then

*return the highest EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)*

if state is a MIN node then


*return the lowest EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)*

if state is a chance node then

*return average of EXPECTIMINIMAX-VALUE of SUCCESSORS(*state*)*

Stochastic Two-Player

- Dice rolls increase b : 21 possible rolls with 2 dice
 - Backgammon \approx 20 legal moves
 - Depth 4 = $20 \times (21 \times 20)^3 = 1.2 \times 10^9$
- As depth increases, probability of reaching a given node shrinks
 - So value of lookahead is diminished
 - So limiting depth is less damaging
 - But pruning is less possible...
- TDGammon uses depth-2 search + very good eval function + reinforcement learning: world-champion level play



What's Next?

- Make sure you know what:
 - Probabilities are
 - Expectations are
- Next topics:
 - Dealing with uncertainty
 - How to learn evaluation functions
 - Markov Decision Processes