

CS 188: Artificial Intelligence

Spring 2007

Lecture 7: CSP-II and Adversarial Search

2/6/2007

Srini Narayanan – ICSI and UC Berkeley

Many slides over the course adapted from Dan Klein, Stuart Russell or Andrew Moore

Summary: Consistency

§ Basic solution: DFS / backtracking



- § Add a new assignment
- § Check for violations

§ Forward checking:

- § Pre-filter unassigned domains after every assignment
- § Only remove values which conflict with current assignments

§ Arc consistency

- § We only defined it for binary CSPs
- § Check for impossible values on all pairs of variables, prune them
- § Run (or not) after each assignment before recursing
- § A pre-filter, not search!





Limitations of Arc Consistency

§ After running arc consistency:

§ Can have one solution left

§ Can have multiple solutions left

§ Can have no solutions left (and not know it)



What went wrong here?


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 kth node.



§ Higher k more expensive to compute


§ (You need to know the k=2 algorithm)



Strong K-Consistency

- § Strong k-consistency: also k-1, k-2, ... 1 consistent
- § Claim: strong n-consistency means we can solve without backtracking!
- § Why?
 - § Choose any assignment to any variable
 - § Choose a new variable
 - § By 2-consistency, there is a choice consistent with the first
 - § Choose a new variable
 - § By 3-consistency, there is a choice consistent with the first 2
 - § ...
- § Lots of middle ground between arc consistency and n-consistency! (e.g. path consistency)

K-consistent vs. strong k-consistent



Iterative Algorithms for CSPs


- § Greedy and local methods typically work with “complete” states, i.e., all variables assigned

- § To apply to CSPs:
 - § Allow states with unsatisfied constraints
 - § Operators *reassign* variable values

- § Variable selection: randomly select any conflicted variable

- § Value selection by min-conflicts heuristic:
 - § Choose value that violates the fewest constraints
 - § I.e., hill climb with $h(n)$ = total number of violated constraints

Example: 4-Queens




- § States: 4 queens in 4 columns ($4^4 = 256$ states)
- § Operators: move queen in column
- § Goal test: no attacks
- § Evaluation: $h(n) =$ number of attacks

Performance of Min-Conflicts

- § Given random initial state, can solve n-queens in almost constant time for arbitrary n with high probability (e.g., n = 10,000,000)
- § The same appears to be true for any randomly-generated CSP *except* in a narrow range of the ratio

$$R = \frac{\text{number of constraints}}{\text{number of variables}}$$



Example: Boolean Satisfiability

§ Given a Boolean expression, is it satisfiable?

§ Very basic problem in computer science

$$p_1 \wedge (p_2 \rightarrow p_3) \wedge ((\neg p_1 \wedge \neg p_3) \rightarrow \neg p_2) \wedge (p_1 \vee p_3)$$

§ Turns out you can always express in 3-CNF

$$(p_1) \wedge (\neg p_2 \vee p_3) \wedge (p_1 \vee p_3 \vee \neg p_2) \wedge (p_1 \vee p_2 \vee p_3)$$

§ 3-SAT: find a satisfying truth assignment

Example: 3-SAT

§ Variables: p_1, p_2, \dots, p_n

§ Domains: $\{\text{true}, \text{false}\}$

§ Constraints: $p_i \vee p_j \vee p_k$

$\neg p_{i'} \vee p_{j'} \vee p_{k'}$

\vdots

$p_{i''} \vee \neg p_{j''} \vee \neg p_{k''}$

*Implicitly
conjoined
(all clauses
must be
satisfied)*


CSPs: Queries

§ Types of queries:

§ Legal assignment


§ All assignments

§ Possible values of some query variable(s) given some evidence (partial assignments)




Problem Structure

- § Tasmania and mainland are independent subproblems
- § Identifiable as connected components of constraint graph
- § Suppose each subproblem has c variables out of n total
 - § Worst-case solution cost is $O((n/c)(d^c))$, linear in n
 - § E.g., $n = 80$, $d = 2$, $c = 20$
 - § $2^{80} = 4$ billion years at 10 million nodes/sec
 - § $(4)(2^{20}) = 0.4$ seconds at 10 million nodes/sec



Tree-Structured CSPs




- § Theorem: if the constraint graph has no loops, the CSP can be solved in $O(n d^2)$ time
 - § Compare to general CSPs, where worst-case time is $O(d^n)$

- § This property also applies to logical and probabilistic reasoning: an important example of the relation between syntactic restrictions and the complexity of reasoning.

Tree-Structured CSPs

§ Choose a variable as root, order variables from root to leaves such that every node's parent precedes it in the ordering



§ For $i = n : 2$, apply $\text{RemoveInconsistent}(\text{Parent}(X_i), X_i)$

§ For $i = 1 : n$, assign X_i consistently with $\text{Parent}(X_i)$


§ Runtime: $O(n d^2)$ (why?)

Tree-Structured CSPs

§ Why does this work?

§ Claim: After each node is processed leftward, all nodes to the right can be assigned in any way consistent with their parent.


§ Proof: Induction on position



§ Why doesn't this algorithm work with loops?

§ Note: we'll see this basic idea again with Bayes' nets and call it belief propagation

Nearly Tree-Structured CSPs



- § Conditioning: instantiate a variable, prune its neighbors' domains
- § Cutset conditioning: instantiate (in all ways) a set of variables such that the remaining constraint graph is a tree
- § Cutset size c gives runtime $O((d^c)(n-c)d^2)$, very fast for small c

CSP Summary

- § CSPs are a special kind of search problem:
 - § States defined by values of a fixed set of variables
 - § Goal test defined by constraints on variable values
- § Backtracking = depth-first search with one legal variable assigned per node
- § Variable ordering and value selection heuristics help significantly
- § Forward checking prevents assignments that guarantee later failure
- § Constraint propagation (e.g., arc consistency) does additional work to constrain values and detect inconsistencies
- § The constraint graph representation allows analysis of problem structure
- § Tree-structured CSPs can be solved in linear time
- § Iterative min-conflicts is usually effective in practice

Games: Motivation

§ Games are a form of *multi-agent environment*

- § What do other agents do and how do they affect our success?
- § Cooperative vs. competitive multi-agent environments.
- § Competitive multi-agent environments give rise to adversarial search a.k.a. *games*

§ Why study games?

- § Games are fun!
- § Historical role in AI
- § Studying games teaches us how to deal with other agents trying to foil our plans
- § *Huge* state spaces – Games are *hard!*
- § Nice, clean environment with clear criteria for success

Game Playing

§ 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 Single-Player?

- § Deterministic, single player, perfect information:
 - § Know the rules
 - § Know what actions do
 - § Know when you win
 - § E.g. Freecell, 8-Puzzle, Rubik's cube
- § ... it's just search!
- § Slight reinterpretation:
 - § Each node stores the best outcome it can reach
 - § This is the maximal outcome of its children
 - § Note that we don't store path sums as before
- § After search, can pick move that leads to best node



Deterministic Two-Player

§ E.g. tic-tac-toe, chess, checkers

§ Minimax search

§ A state-space search tree

§ Players alternate


§ Each layer, or ply, consists of a round of moves

§ Choose move to position with highest **minimax value** = best achievable utility against best play


§ Zero-sum games

§ One player maximizes result


§ The other minimizes result



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?



Resource Limits

§ Cannot search to leaves

§ Limited search

§ Instead, search a limited depth of the tree

§ Replace terminal utilities with an eval function for non-terminal positions

§ Guarantee of optimal play is gone


§ More plies makes a BIG difference

§ 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



§ Ideal function: returns the utility of the position

§ In practice: typically weighted linear sum of features:


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

§ e.g. $f_1(s) = (\text{num white queens} - \text{num black queens})$, etc.

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.




This works for single-agent search as well!

Why do we want to do this for multiplayer games?


Alpha-Beta Pruning

- § A way to improve the performance of the Minimax Procedure
- § Basic idea: *“If you have an idea which is surely bad, don’t take the time to see how truly awful it is”* ~ Pat Winston




- We don't need to compute the value at this node.
- No matter what it is it can't effect the value of the root node.

α - β Pruning Example




Pruning in Minimax Search



α - β Pruning

§ General configuration

- § α is the best value the Player can get at any choice point along the current path
- § If n is worse than α , MAX will avoid it, so prune n 's branch
- § 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!
- § A simple example of **metareasoning**, here reasoning about which computations are relevant


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, don't know where the mines are

- § Can do **expectimax search**


- § Chance nodes, like actions except the environment controls the action chosen
- § Calculate utility for each node
- § Max nodes as in search
- § Chance nodes take average (expectation) of value of children

- § Later, we'll learn how to formalize this as a **Markov Decision Process**



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*)

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.
- § **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.
- § **Othello:** human champions refuse to compete against computers, which are too good.
- § **Go:** human champions refuse to compete against computers, which are too bad. In go, $b > 300$, so most programs use pattern knowledge bases to suggest plausible moves.

Stochastic Two-Player

- § Dice rolls increase b : 21 possible rolls with 2 dice
 - § Backgammon \approx 20 legal moves
 - § Depth 4 = $20 \times (21 \times 20)^3 \approx 1.2 \times 10^9$
- § 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

