

CS 188: Artificial Intelligence

Spring 2010

Lecture 6: Adversarial Search

2/4/2010

Pieter Abbeel – UC Berkeley
Many slides adapted from Dan Klein

Announcements

- Project 1 is done!
- Written 1 is out and is due next week Monday
 - Homework policy – see website!

Today

- Finish up Search and CSPs
- Intermezzo on A* and heuristics
- Start on Adversarial Search

3

CSPs: our status

- 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
 - Branching on only one variable per layer in search tree
 - Incremental constraint checks ("Fail fast")
- Heuristics at our points of choice to improve running time:
 - Ordering variables: Minimum Remaining Values and Degree Heuristic
 - Ordering of values: Least Constraining Value
- Today:
 - Filtering: forward checking, arc consistency → enable computation of heuristics
 - Structure: Disconnected and tree-structured CSPs are efficient
 - Iterative improvement: min-conflicts is usually effective in practice

4

Example: Map-Coloring

- Variables: WA, NT, Q, NSW, V, SA, T

- Domain: $D = \{red, green, blue\}$

- Constraints: adjacent regions must have different colors



$$WA \neq NT$$

$$(WA, NT) \in \{(red, green), (red, blue), (green, red), \dots\}$$

- Solutions are assignments satisfying all constraints, e.g.:

$$\{WA = red, NT = green, Q = red, NSW = green, V = red, SA = blue, T = green\}$$

6

Filtering: Forward Checking



- Idea: Keep track of remaining legal values for unassigned variables (using immediate constraints)
- Idea: Terminate when any variable has no legal values

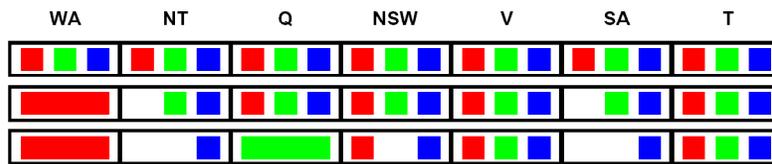


[demo: forward checking animation]

Filtering: Forward Checking



- Forward checking propagates information from assigned to adjacent unassigned variables, but doesn't detect more distant failures:



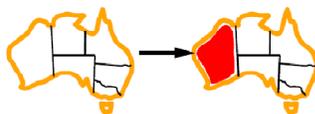
- NT and SA cannot both be blue!
- Why didn't we detect this yet?
- Constraint propagation* repeatedly enforces constraints (locally)

8

Consistency of An 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 tail!



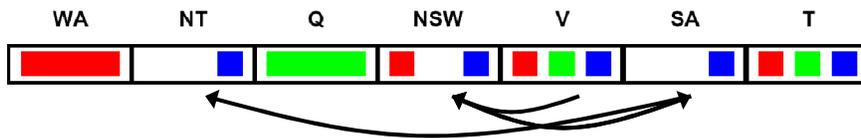
- What happens?
- Forward checking = Enforcing consistency of each arc pointing to the new assignment

9

Arc Consistency of a CSP



- Simplest form of propagation makes each arc *consistent*
 - $X \rightarrow Y$ is consistent iff for every value x there is some allowed y



- If X loses a value, neighbors of X need to be rechecked!
- Arc consistency detects failure earlier than forward checking
- What's the downside of arc consistency?
- Can be run as a preprocessor or after each assignment

10

Establishing Arc Consistency

```

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?

[demo: arc consistency animation]₁

Backtracking with MRV, Degree, LCV, Filtering

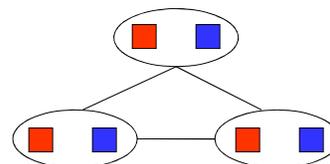
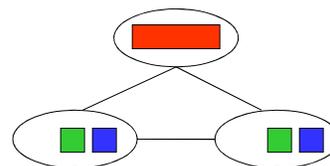
```
function RecursiveBacktracking(pa, fd, vars, constraints)
  if IsComplete(pa) then return pa
  next_var <-- select_MRV_Degree(pa, fd, vars, constraints)
  for each value in fd[next_var] do
    new_fd[value] <-- constraint_prop(pa, fd, vars, constraints)
  for each value in fd[next_var] in order of LCV do
    if no new_fd[value] is empty
      then
        add {var=value} to pa
        result <-- recursive_backtracking(pa, new_fd[value], vars, constraints)
        if (result not equal to failure) return result
  return failure
```

- select_MRV_degree: selects an unassigned variable based on MRV and degree heuristic
- constraint_prop: performs constraint propagation, this could be through forward propagation or through arc consistency
- pa: partial assignment
- fd: filtered domains

12

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)

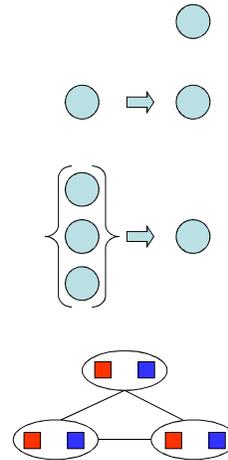


What went wrong here?

14

[[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



15

[[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)

16

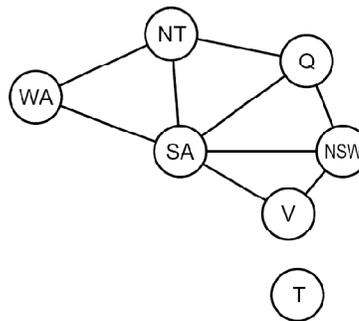
CSPs: our status

- 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
 - Branching on only one variable per layer in search tree
 - Incremental constraint checks (“Fail fast”)
- Heuristics at our points of choice to improve running time:
 - Ordering variables: Minimum Remaining Values and Degree Heuristic
 - Ordering of values: Least Constraining Value
- Today:
 - ✓ Filtering: forward checking, arc consistency → enable computation of heuristics
 - Structure: Disconnected and tree-structured CSPs are efficient
 - Iterative improvement: min-conflicts is usually effective in practice

17

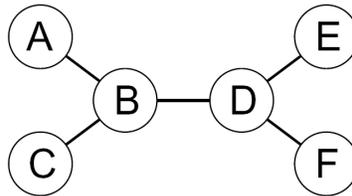
Problem Structure

- 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



18

Tree-Structured CSPs

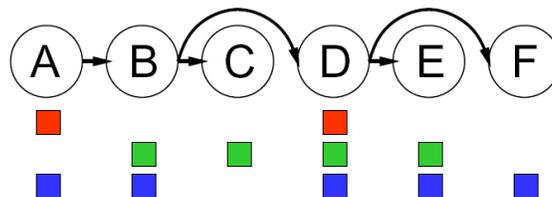
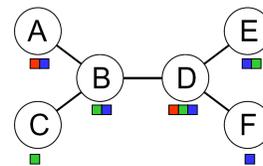


- **Theorem:** if the constraint graph has no loops, the CSP can be solved in $O(n d^2)$ time
 - Compare to general CSPs, where worst-case time is $O(d^n)$
- This property also applies to probabilistic reasoning (later): an important example of the relation between syntactic restrictions and the complexity of reasoning.

19

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

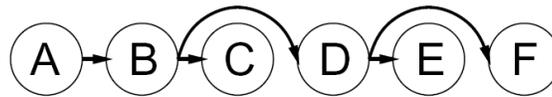


- For $i = n : 2$, apply `RemoveInconsistent(Parent(X_i), X_i)`
- For $i = 1 : n$, assign X_i consistently with `Parent(X_i)`
- Runtime: $O(n d^2)$ (why?)

20

Tree-Structured CSPs

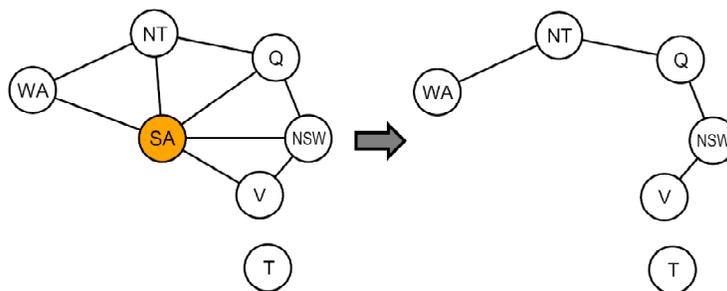
- Why does this work?
- Claim: After each node is processed leftward, all nodes to the right can be assigned in any way consistent with their parent.
- Proof: Induction on position



- Why doesn't this algorithm work with loops?
- Note: we'll see this basic idea again with Bayes' nets

21

Nearly Tree-Structured CSPs

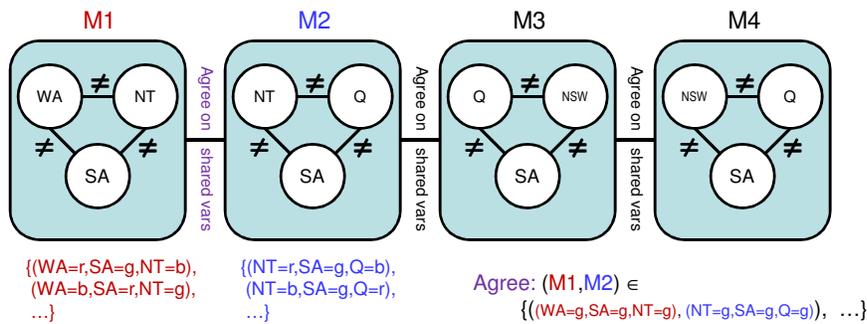
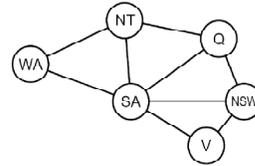


- Conditioning: instantiate a variable, prune its neighbors' domains
- Cutset conditioning: instantiate (in all ways) a set of variables such that the remaining constraint graph is a tree
- Cutset size c gives runtime $O(d^c (n-c) d^2)$, very fast for small c

22

[[Tree Decompositions]**

- Create a tree-structured graph of overlapping subproblems, each is a mega-variable
- Solve each subproblem to enforce local constraints
- Solve the CSP over subproblem mega-variables using our efficient tree-structured CSP algorithm



23

CSPs: our status

- 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
 - Branching on only one variable per layer in search tree
 - Incremental constraint checks ("Fail fast")
- Heuristics at our points of choice to improve running time:
 - Ordering variables: Minimum Remaining Values and Degree Heuristic
 - Ordering of values: Least Constraining Value
- Today:
 - ✓ Filtering: forward checking, arc consistency → enable computation of heuristics
 - ✓ Structure: Disconnected and tree-structured CSPs are efficient
 - Iterative improvement: min-conflicts is usually effective in practice

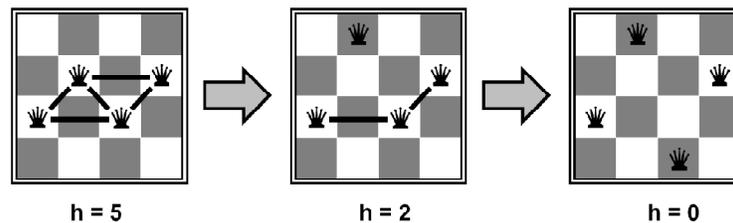
24

Iterative Algorithms for CSPs

- Local search methods: typically work with “complete” states, i.e., all variables assigned
- To apply to CSPs:
 - Start with some assignment with unsatisfied constraints
 - Operators *reassign* variable values
 - No fringe! Live on the edge.
- 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

25

Example: 4-Queens



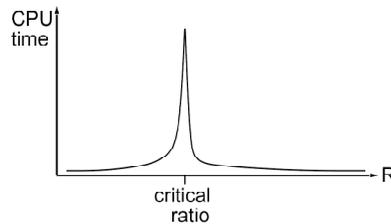
- States: 4 queens in 4 columns ($4^4 = 256$ states)
- Operators: move queen in column
- Goal test: no attacks
- Evaluation: $c(n)$ = number of attacks

27

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

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



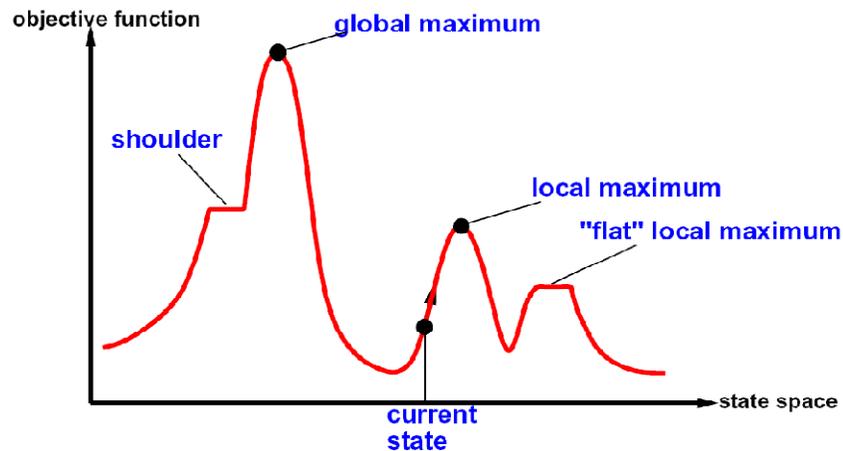
28

Hill Climbing

- **Simple, general idea:**
 - Start wherever
 - Always choose the best neighbor
 - If no neighbors have better scores than current, quit
- **Why can this be a terrible idea?**
 - Complete?
 - Optimal?
- **What's good about it?**

29

Hill Climbing Diagram



- Random restarts?
- Random sideways steps?

30

Simulated Annealing

- Idea: Escape local maxima by allowing downhill moves
 - But make them rarer as time goes on

function SIMULATED-ANNEALING(*problem*, *schedule*) **returns** a solution state

inputs: *problem*, a problem
schedule, a mapping from time to "temperature"

local variables: *current*, a node
next, a node
T, a "temperature" controlling prob. of downward steps

current ← MAKE-NODE(INITIAL-STATE[*problem*])

for *t* ← 1 **to** ∞ **do**

T ← *schedule*[*t*]

if *T* = 0 **then return** *current*

next ← a randomly selected successor of *current*

ΔE ← VALUE[*next*] − VALUE[*current*]

if $\Delta E > 0$ **then** *current* ← *next*

else *current* ← *next* only with probability $e^{\Delta E/T}$

31

CSPs 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 incremental constraint checks
- Ordering: variable and value choice heuristics help significantly
- Filtering: forward checking, arc consistency
 - assist in computing heuristic ordering choices
 - prevent assignments that guarantee later failure
- Structure: Disconnected and tree-structured CSPs are efficient
- Iterative improvement: min-conflicts is usually effective in practice

32

33

Intermezzo: A* heuristics --- 8 puzzle

7	2	4
5		6
8	3	1

Start State

	1	2
3	4	5
6	7	8

Goal State

- What are the states?
- What are the actions?
- What is the cost?

Intermezzo: A* heuristics --- 8 puzzle

- Number of misplaced tiles: Admissible or not?
- What if we had an easier 8-puzzle where any tile could slide any direction at any time, ignoring other tiles and used their total *Manhattan* distance. Admissible or not?
- What if we had a piece of code that could quickly find a sequence of actions that reaches the goal state. Is the number of actions returned by that piece of code an admissible heuristic?

Intermezzo: A* heuristics --- pacman trying to eat all food pellets

- Consider an algorithm that takes the distance to the closest food pellet, say at (x,y) . Then it adds the distance between (x,y) and the closest food pellet to (x,y) , and continues this process until no pellets are left, each time calculating the distance from the last pellet. Is this heuristic admissible?
- What if we used the Manhattan distance rather than distance in the maze in the above procedure?

36

Intermezzo: A* heuristics

- A particular procedure to quickly find a perhaps suboptimal solution to the search problem is in general not admissible.
 - It is only admissible if it always finds the optimal solution (but then it is already solving the problem we care about, hence not that interesting as a heuristic).
- A particular procedure to quickly find a perhaps suboptimal solution to a relaxed version of the search problem need not be admissible.
 - It will be admissible if it always finds the *optimal* solution to the relaxed problem.

37

[[Advanced Heuristics]]**

- Fast algorithm to find solution to problem that is at most a factor k more expensive than the optimal solution
 - Can use (cost fast algorithm solution) / k
- Often problem can be formulated as an integer programming problem
 - Relaxation to linear programming problem provides admissible heuristic

38

39

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

40

GamesCrafters

The GamesCrafters research and development group was formed in 2001 as a "nesting hole" to gather and engage top undergraduates as they explore the fertile area of computational game theory. At the core of the project is GAMESMAN, an open-source AI architecture developed for solving, playing and analyzing two-person abstract strategy games (e.g., Tic-Tac-Toe or Chess). Given the description of a game as input, our system generates a command-line interface and Tcl/Tk graphical application that will solve it (in the strong sense), and then play it perfectly. Programmers can easily prototype a new game with multiple rule variants, learn the strategy via color-coded value moves (win = go = green, tie = caution = yellow, lose = stop = red), and perform extended analysis.

The group is accessible to undergraduates at all levels. Those not yet ready to dive into code can create graphics, find bugs, or research the history of games for our website. Programmers can easily prototype a new game with multiple rule variants, design a fun interface, and perform extended analysis. Advanced students are encouraged to tinker with the software core, and optimize the solves, evaluations, hash functions, reordering, user experience, etc.

Since this is not a class, but directed group study, students can re-register as often as they like, most stay for two or three semesters. This allows for a real continuity to be formed, with veterans providing continuity and mentoring as project leads, as well as allowing for more ambitious multi-term projects. Our alumni have told us how valuable this experience has been for them, growing from a nurturing environment to mature as researchers, developers, and leaders.

Over the past six years, over two hundred undergraduates have implemented more than sixty-five games and several advanced software engineering projects. Our future research direction is "hunting big game", i.e., implementing, solving, and analyzing large games whose perfect strategy is yet unknown.

This semester (Fall 2008), we're meeting in 606 Soda Hall, Mondays from 6-9PM. It is a "Directed Group Study" course worth two units led by Dr. Dan Garcia.

<http://gamescrafters.berkeley.edu/>

Dan Garcia.

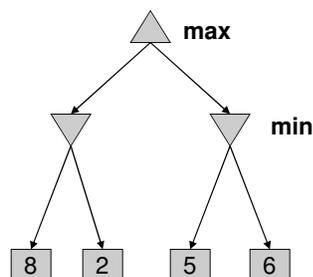
41

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

43

Simple two-player game example



44