CSE 473: Artificial Intelligence

Constraint Satisfaction Problems





Instructor: Luke Zettlemoyer

University of Washington

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All CS188 materials are available at http://ai.berkeley.edu.]

What is Search For?

- Assumptions about the world: a single agent, deterministic actions, fully observed state, discrete state space
- Planning: sequences of actions
 - The path to the goal is the important thing
 - Paths have various costs, depths
 - Heuristics give problem-specific guidance
- Identification: assignments to variables
 - The goal itself is important, not the path
 - All paths at the same depth (for some formulations)
 - CSPs are specialized for identification problems



Constraint Satisfaction Problems



Constraint Satisfaction Problems

- Standard search problems:
 - State is a "black box": arbitrary data structure
 - Goal test can be any function over states
 - Successor function can also be anything
- Constraint satisfaction problems (CSPs):
 - A special subset of search problems
 - State is defined by variables X_i with values from a domain D (sometimes D depends on i)
 - Goal test is a set of constraints specifying allowable combinations of values for subsets of variables
- Simple example of a *formal representation language*
- Allows useful general-purpose algorithms with more power than standard search algorithms





CSP Examples



Example: Map Coloring

- Variables: WA, NT, Q, NSW, V, SA, T
- Domains: D = {red, green, blue}
- Constraints: adjacent regions must have different colors

Implicit: $WA \neq NT$

Explicit: $(WA, NT) \in \{(red, green), (red, blue), \ldots\}$

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

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





Example: N-Queens

- Formulation 1:
 - Variables: X_{ij}
 - Domains: {0, 1}
 - Constraints





 $\begin{aligned} \forall i, j, k \ (X_{ij}, X_{ik}) &\in \{(0, 0), (0, 1), (1, 0)\} \\ \forall i, j, k \ (X_{ij}, X_{kj}) &\in \{(0, 0), (0, 1), (1, 0)\} \\ \forall i, j, k \ (X_{ij}, X_{i+k,j+k}) &\in \{(0, 0), (0, 1), (1, 0)\} \\ \forall i, j, k \ (X_{ij}, X_{i+k,j-k}) &\in \{(0, 0), (0, 1), (1, 0)\} \end{aligned}$

$$\sum_{i,j} X_{ij} = N$$

Example: N-Queens

- Formulation 2:
 - Variables: Q_k
 - Domains: $\{1, 2, 3, \dots N\}$
 - Constraints:

Implicit: $\forall i, j \text{ non-threatening}(Q_i, Q_j)$

Explicit: $(Q_1, Q_2) \in \{(1, 3), (1, 4), \ldots\}$





Constraint Graphs



Constraint Graphs

- Binary CSP: each constraint relates (at most) two variables
- Binary constraint graph: nodes are variables, arcs show constraints
- General-purpose CSP algorithms use the graph structure to speed up search. E.g., Tasmania is an independent subproblem!



[Demo: CSP applet (made available by aispace.org) -- n-queens]

Screenshot of Demo N-Queens



Example: Cryptarithmetic

- Variables:
 - $F T U W R O X_1 X_2 X_3$
- Domains:
 - $\{0,1,2,3,4,5,6,7,8,9\}$
- Constraints:
 - $\operatorname{alldiff}(F, T, U, W, R, O)$
 - $O + O = R + 10 \cdot X_1$

• • •



Example: Sudoku



- Variables:
 - Each (open) square
- Domains:
 - {1,2,...,9}
- Constraints:

9-way alldiff for each column9-way alldiff for each row9-way alldiff for each region

(or can have a bunch of pairwise inequality constraints)

Example: The Waltz Algorithm

- The Waltz algorithm is for interpreting line drawings of solid polyhedra as 3D objects
- An early example of an AI computation posed as a CSP





- Approach:
 - Each intersection is a variable
 - Adjacent intersections impose constraints on each other
 - Solutions are physically realizable 3D interpretations

Example: The Waltz Algorithm

- The Waltz algorithm is for interpreting line drawings of solid polyhedra as 3D objects
- An early example of an AI computation posed as a CSP





Approach:

- Each intersection is a variable
- Adjacent intersections impose constraints on each other
- Solutions are physically realizable 3D interpretations

Varieties of CSPs and Constraints



Varieties of CSPs

- Discrete Variables
 - Finite domains
 - Size *d* means O(*dⁿ*) complete assignments
 - E.g., Boolean CSPs, including Boolean satisfiability (NPcomplete)
 - Infinite domains (integers, strings, etc.)
 - E.g., job scheduling, variables are start/end times for each job
 - Linear constraints solvable, nonlinear undecidable
- Continuous variables
 - E.g., start/end times for Hubble Telescope observations
 - Linear constraints solvable in polynomial time by LP methods (see cs170 for a bit of this theory)





Varieties of Constraints

- Varieties of Constraints
 - Unary constraints involve a single variable (equivalent to reducing domains), e.g.:

 $SA \neq green$

Binary constraints involve pairs of variables, e.g.:

 $SA \neq WA$

- Higher-order constraints involve 3 or more variables: e.g., cryptarithmetic column constraints
- Preferences (soft constraints):
 - E.g., red is better than green
 - Often representable by a cost for each variable assignment
 - Gives constrained optimization problems
 - (We'll ignore these until we get to Bayes' nets)



Real-World CSPs

- Assignment problems: e.g., who teaches what class
- Timetabling problems: e.g., which class is offered when and where?
- Hardware configuration
- Transportation scheduling
- Factory scheduling
- Circuit layout
- Fault diagnosis
- Interpretended in the second secon



Many real-world problems involve real-valued variables...

Solving CSPs



Standard Search Formulation

- Standard search formulation of CSPs
- States defined by the values assigned so far (partial assignments)
 - Initial state: the empty assignment, {}
 - Successor function: assign a value to an unassigned variable
 - Goal test: the current assignment is complete and satisfies all constraints
- We'll start with the straightforward, naïve approach, then improve it



Search Methods

What would BFS do?
What would DFS do?

What problems does naïve search have?



.

Video of Demo Coloring ---DFS



Backtracking Search



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
 - 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 constraint
 - "Incremental goal test"
- Depth-first search with these two improvements is called *backtracking search* (not the best name)
- Can solve n-queens for n ≈ 25



Backtracking Example



Backtracking Search



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

[Demo: coloring -- backtracking]

Auton's Graphics

 \mathbf{x}

Video of Demo Coloring – Backtracking

 \mathbf{x}

Video of Demo Coloring – Backtracking



Improving Backtracking

- General-purpose ideas give huge gains in speed
 - ... but it's all still NP-hard
- Filtering: Can we detect inevitable failure early?
- Ordering:
 - Which variable should be assigned next? (MRV)
 - In what order should its values be tried? (LCV)
- Structure: Can we exploit the problem structure?





Arc Consistency and Beyond



Forward Checking



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







Auton's Graphics

Are We Done?



Constraint Propagation

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

NT

SA

WA

Ω

NSW



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

Arc Consistency of an Entire CSP

• A simple form of propagation makes sure all arcs are simultaneously consistent:



- Arc consistency detects failure earlier than forward checking
- Important: If X loses a value, neighbors of X need to be rechecked!
- Must rerun after each assignment!

Remember: Delete from the tail!


Constraint Propagation Are We Done?



Limitations of Arc Consistency

- After enforcing arc consistency:
 - Can have one solution left
 - Can have multiple solutions left
 - Can have no solutions left (and not know it)

Arc consistency still runs inside a backtracking search!





What went wrong here?

Ordering: Minimum Remaining Values

- Minimum remaining values (MRV):
 - Choose the variable with the fewest legal values



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

Ordering: Degree Heuristic

- Tie-breaker among MRV variables
- Degree heuristic:
 - Choose the variable participating in the most constraints on remaining variables



Why most rather than fewest constraints?

Ordering: Least Constraining Value

- Given a choice of variable:
 - Choose the *least constraining value*
 - The one that rules out the fewest values in the remaining variables
 - Note that it may take some computation to determine this!
- Why least rather than most?
- Combining these heuristics makes 1000 queens feasible





K-Consistency



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 case: arc consistency)



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. k=3, called path consistency)

Structure



Problem Structure

- Extreme case: independent subproblems
 - Example: Tasmania and mainland do not interact
- Independent subproblems are identifiable as connected components of constraint graph
- Suppose a graph of n variables can be broken into subproblems of only c variables:
 - Worst-case solution cost is O((n/c)(d^c)), linear in n
 - E.g., n = 80, d = 2, c = 20
 - 2⁸⁰ = 4 billion years at 10 million nodes/sec
 - (4)(2²⁰) = 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 d²) time
 - Compare to general CSPs, where worst-case time is O(dⁿ)
- This property also applies to probabilistic reasoning (earlier): an example of the relation between syntactic restrictions and the complexity of reasoning

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 RemoveInconsistent(Parent(X_i),X_i)
- Assign forward: For i = 1 : n, assign X_i consistently with Parent(X_i)
- Runtime: O(n d²) (why?)



Tree-Structured CSPs

- Claim 1: After backward pass, all root-to-leaf arcs are consistent
- Proof: Each X→Y was made consistent at one point and Y's domain could not have been reduced thereafter (because Y's children were processed before Y)



- Claim 2: If root-to-leaf arcs are consistent, forward assignment will not backtrack
- Proof: Induction on position
- Why doesn't this algorithm work with cycles in the constraint graph?
- Note: same basic idea as variable elimination in Bayes' nets

Improving Structure



Nearly Tree-Structured CSPs



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) d²), very fast for small c

Cutset Conditioning



Cutset Quiz

Find the smallest cutset for the graph below.



Tree Decomposition*

Idea: create a tree-structured graph of mega-variables NT Q Each mega-variable encodes part of the original CSP WA Subproblems overlap to ensure consistent solutions SA NSW M2 M3 M1 M4 Agree Agree Agree NS NS NT NT WA Q Q W W Ő **0** Р ¥ ¥ ¥ ¥ shared vars shared vars shared vars SA SA SA SA {(WA=r,SA=g,NT=b), {(NT=r,SA=g,Q=b), Agree: $(M1, M2) \in$ (WA=b,SA=r,NT=g), (NT=b,SA=g,Q=r), {((WA=g,SA=g,NT=g), (NT=g,SA=g,Q=g)), ...} ...} ...}

Iterative Improvement



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(n) = total number of violated constraints



Example: 4-Queens



- States: 4 queens in 4 columns (4⁴ = 256 states)
- Operators: move queen in column
- Goal test: no attacks
- Evaluation: c(n) = number of attacks

[Demo: n-queens – iterative improvement (L5D1)] [Demo: coloring – iterative improvement]

Video of Demo Iterative Improvement – n Queens



Video of Demo Iterative Improvement – Coloring



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





Summary: CSPs

- CSPs are a special kind of search problem:
 - States are partial assignments
 - Goal test defined by constrained
- Basic solution: backtracking sea
- Speed-ups:
 - Ordering
 - Filtering
 - Structure



Iterative min-conflicts is often effective in practice

Local Search



Local Search

- Tree search keeps unexplored alternatives on the fringe (ensures completeness)
- Local search: improve a single option until you can't make it better (no fringe!)
- New successor function: local changes



Generally much faster and more memory efficient (but incomplete and suboptimal)

Hill Climbing

- Simple, general idea:
 - Start wherever
 - Repeat: move to the best neighboring state
 - If no neighbors better than current, quit
- What's bad about this approach?
 - Complete?
 - Optimal?
- What's good about it?



Hill Climbing Diagram



Hill Climbing Quiz



Starting from X, where do you end up ?

Starting from Y, where do you end up ?

Starting from Z, where do you end up ?

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 \leftarrow MAKE-NODE(INITIAL-STATE[problem])
for t \leftarrow 1 to \infty do
     T \leftarrow schedule[t]
     if T = 0 then return current
     next \leftarrow a randomly selected successor of current
     \Delta E \leftarrow \text{VALUE}[next] - \text{VALUE}[current]
     if \Delta E > 0 then current \leftarrow next
     else current \leftarrow next only with probability e^{\Delta E/T}
```



Simulated Annealing

- Theoretical guarantee:
 - Stationary distribution:

$$p(x) \propto e^{rac{E(x)}{kT}}$$

- If T decreased slowly enough, will converge to optimal state!
- Is this an interesting guarantee?
- Sounds like magic, but reality is reality:
 - The more downhill steps you need to escape a local optimum, the less likely you are to ever make them all in a row
 - People think hard about *ridge operators* which let you jump around the space in better ways



Genetic Algorithms



- Genetic algorithms use a natural selection metaphor
 - Keep best N hypotheses at each step (selection) based on a fitness function
 - Also have pairwise crossover operators, with optional mutation to give variety
- Possibly the most misunderstood, misapplied (and even maligned) technique around
Example: N-Queens



- Why does crossover make sense here?
- When wouldn't it make sense?
- What would mutation be?
- What would a good fitness function be?