

CSE 473: Intro. to Artificial Intelligence

Constraint Satisfaction Problems




Presenter: Galen Andrew

[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>.]

Announcements

- Prof. Weld away today and Wednesday
 - I will be beginning the lecture series on Constraint Satisfaction Problems (CSPs)
 - Prof. Luke Zettlemoyer will continue on Wednesday.
- Project 1: Search
 - Due next week, Monday 10/13 at 11:59 PM.
 - Start early and ask questions. It's longer than most!
- Come to TA office hours with questions or general help
 - Galen: Wed 1:00-3:00
 - Nao: Tue 1:30-2:30, Thu 1:00-2:00
 - Travis: Fri 3:30-4:30
 - Jeff: Wed 10:30-11:30

Constraint Satisfaction Problems




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
 - Assume little about problem structure
- Identification: assignments to variables
 - The goal itself is important, not the path
 - All paths at the same depth (for some formulations)
 - CSPs are *structured* identification 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 (CSP):
 - 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
- Making use of CSP formulation allows for optimized algorithms
 - Typical example of trading generality for utility (in this case, speed)




CSP Examples




Example: Map Coloring

- Variables: WA, NT, Q, NSW, V, SA, T
- Domains: $D = \{\text{red, green, blue}\}$
- Constraints: adjacent regions must have different colors
 - Implicit: $WA \neq NT$
 - Explicit: $(WA, NT) \in \{(\text{red, green}), (\text{red, blue}), \dots\}$
- Solutions are assignments satisfying all constraints, e.g.:
 - $\{WA=\text{red}, NT=\text{green}, Q=\text{red}, NSW=\text{green}, V=\text{red}, SA=\text{blue}, T=\text{green}\}$



Example: N-Queens

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



$$\forall i, j, k \quad (X_{ij}, X_{ik}) \in \{(0, 0), (0, 1), (1, 0)\}$$

$$\forall i, j, k \quad (X_{ij}, X_{kj}) \in \{(0, 0), (0, 1), (1, 0)\}$$


$$\forall i, j, k \quad (X_{ij}, X_{i+k, j+k}) \in \{(0, 0), (0, 1), (1, 0)\}$$

$$\forall i, j, k \quad (X_{ij}, X_{i+k, j-k}) \in \{(0, 0), (0, 1), (1, 0)\}$$


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

Example: N-Queens

- Formulation 2:
 - Variables: Q_k
 - Domains: $\{1, 2, 3, \dots, N\}$
 - Constraints:
 - Implicit: $\forall i, j$ non-threatening(Q_i, Q_j)
 - Explicit: $(Q_1, Q_2) \in \{(1, 3), (1, 4), \dots\}$




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!




Example: Cryptarithmic

- 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:
 - $\text{alldiff}(F, T, U, W, R, O)$
 - $O + O = R + 10 \cdot X_1$
 - ...



Example: Sudoku




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

9-way all diff for each column
 9-way all diff for each row
 9-way all diff 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

Varieties of CSPs



Varieties of CSP Variables

- Discrete Variables
 - Finite domains
 - Size d means $O(d^n)$ complete assignments
 - E.g., Boolean CSPs, including Boolean satisfiability (NP-complete)
 - 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 linear program methods (see CSE 521 for a bit of LP theory)




Varieties of CSP Constraints

- Varieties of Constraints
 - Unary constraints involve a single variable (equivalent to reducing domains), e.g.:
 - $SA \neq \text{green}$
 - Binary constraints involve pairs of variables, e.g.:
 - $SA \neq WA$
 - Higher-order constraints involve 3 or more variables:
 - e.g., cryptarithmic 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 (VLSI layout)
- Transportation scheduling
- Factory scheduling
- Circuit layout
- Fault diagnosis
- ... lots more!




- 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 DFS do?
- What would BFS do?
- What problems does naïve search have?




[Demo: coloring -- dfs]

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*
- Can solve n-queens for $n \approx 25$



Backtracking Example



Backtracking Search

```

function BACKTRACKING-SEARCH(csp) returns solution/failure
  return RECURSIVE-BACKTRACKING({}, csp)
function RECURSIVE-BACKTRACKING(assignment, csp) returns soln/failure
  if assignment is complete then return assignment
  var ← SELECT-UNASSIGNED-VARIABLE(VARIABLES[csp], assignment, csp)
  for each value in ORDER-DOMAIN-VALUES(var, assignment, csp) do
    if value is consistent with assignment given CONSTRAINTS[csp] then
      add {var = value} to assignment
      result ← RECURSIVE-BACKTRACKING(assignment, csp)
      if result ≠ failure then return result
      remove {var = value} from assignment
  return failure
    
```

- What are the choice points?

[Demo: coloring – backtracking]

Video of Demo Coloring – Backtracking



Improving Backtracking

- General-purpose ideas give huge gains in speed
- Ordering:
 - Which variable should be assigned next?
 - In what order should its values be tried?
- Filtering: Can we detect inevitable failure early?
- Structure: Can we exploit the problem structure?




Filtering



Filtering: Forward Checking

- Filtering: Keep track of domains for unassigned variables and cross off bad options
- Forward checking: Cross off values that violate a constraint when added to the existing assignment




[Demo: coloring -- forward checking]

Video of Demo Coloring – Backtracking with Forward Checking



Filtering: Constraint Propagation


- Forward checking only propagates information from assigned to unassigned
- It doesn't catch when two unassigned variables have no consistent assignment:



- NT and SA cannot both be blue!
- Why didn't we detect this yet?
- Constraint propagation: reason from constraint to constraint

Consistency of a Single 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



- Forward checking: Enforcing consistency of arcs pointing to each new assignment

Arc Consistency of an Entire CSP

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



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

Remember: Delete from the tail!

AC-3 algorithm for 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
     $(X_i, X_j) \leftarrow \text{REMOVE-FIRST}(\text{queue})$ 
    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 ← 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 ← 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: CSP applet (made available by aisspace.org) – n-queens]

Video of Demo Arc Consistency – CSP Applet – n Queens



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?

[Demo: coloring -- forward checking]
[Demo: coloring -- arc consistency]

Video of Demo Coloring – Backtracking with Forward Checking – Complex Graph



Video of Demo Coloring – Backtracking with Arc Consistency – Complex Graph



Ordering



Ordering: Minimum Remaining Values

- Variable Ordering: Minimum remaining values (MRV):
 - Choose the variable with the fewest legal left values in its domain



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



Ordering: Maximum Degree

- Tie-breaker among MRV variables
 - What is the very first state to color? (All have 3 values remaining.)
- Maximum degree heuristic:
 - Choose the variable participating in the most constraints on remaining variables



- Why most rather than fewest constraints?

Ordering: Least Constraining Value

- Value Ordering: Least Constraining Value

- Given a choice of variable, choose the *least constraining value*
- I.e., the one that rules out the fewest values in the remaining variables
- Note that it may take some computation to determine this! (E.g., rerunning filtering)



- Why least rather than most?

- Combining these ordering ideas makes 1000 queens feasible



Rationale for MRV, MD, LCV

- We want to enter the most promising branch, but we also want to detect failure quickly

- MRV+MD:

- Choose the variable that is most likely to cause failure
- It must be assigned at some point, so if it is doomed to fail, better to find out soon

- LCV:

- We hope our early value choices do not doom us to failure
- Choose the value that is most likely to succeed