More Constraint Satisfaction

CMPSCI 383 October 13, 2011

Today's lecture

- A review of CSPs
- Local search for CSPs
- Taking advantage of the structure of the CSP
- Some applications

- In CSPs, states are defined by assignments of values to a set of variables X₁...X_n. Each variable X_i has a domain D_i of possible values.
- States are evaluated based on their consistency with a set of constraints C₁...C_m over the values of the variables.
- A goal state is a complete assignment to all variables that satisfies all the constraints.

Local Consistency

- Node Consistency: satisfies all unary constraints
- Arc Consistency: satisfies all binary constraints
- Path Consistency:
- n-consistency: for any consistent assignment to any set of n-1 variables, a consistent value can be found for any n-th variable.

X is <u>arc-consistent with respect to Y</u> if for every value in Dx there is a value in Dy that satisfies the binary constraint on arc (X,Y).

- Note:
 - Not symmetric in general
 - To make X ac with respect to Y, remove values from Dx
 - To make Y ac with respect to X, remove values from Dy

Arc Consistency (slightly different from the book)

function AC-3(*csp*) returns the CSP, possibly with reduced domains inputs: *csp*, a binary CSP with variables $\{X_1, X_2, \ldots, X_n\}$ local variables: *queue*, a queue of arcs, initially all the arcs in *csp*

while queue is not empty do $(X_i, X_j) \leftarrow \text{REMOVE-FIRST}(queue)$ 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

Standard Search Formulation

Let's start with the straightforward approach, then fix it

States are defined by the values assigned so far

- Initial state: the empty assignment { }
- Successor function: assign a value to an unassigned variable that does not conflict with current assignment
 → fail if no legal assignments
- Goal test: the current assignment is complete
- 1. This is the same for all CSPs
- 2. Every solution appears at depth n with n variables \rightarrow use depth-first search
- 3. Path is irrelevant, so can also use complete-state formulation
- 4. b = (n k)d at depth k, hence $n! \cdot d^n$ leaves (d is domain size)

Backtracking Search

Variable assignments are commutative, 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 node \rightarrow b = d and there are dⁿ leaves
- Depth-first search for CSPs with single-variable assignments is called backtracking search
- Backtracking search is the basic uninformed algorithm for CSPs
- Can solve n-queens for n ≈ 25

Backtracking Search (the book's)

function BACKTRACKING-SEARCH(csp) returns a solution, or failure
return BACKTRACK({ }, csp)

function BACKTRACK(assignment, csp) returns a solution, or failure if assignment is complete then return assignment $var \leftarrow SELECT-UNASSIGNED-VARIABLE(csp)$ for each value in ORDER-DOMAIN-VALUES(var, assignment, csp) do if value is consistent with assignment then add {var = value} to assignment inferences \leftarrow INFERENCE(csp, var, value) if inferences \neq failure then add inferences to assignment result \leftarrow BACKTRACK(assignment, csp) if result \neq failure then return result remove {var = value} and inferences from assignment return failure

Improving backtracking efficiency

- Basic question: What next step should our search procedure take?
- Approaches
 - Minimum remaining values heuristic
 - Degree heuristic
 - Least-constraining value heuristic

Minimum remaining values (MRV) heuristic



 Select the most constrained variable (the variable with the smallest number of remaining values)

Degree heuristic



- Select the variable that is involved in the largest number of constraints with other unassigned variables
- The most constraining variable.
- A useful tie breaker (and guide to starting).
- In what order should its values be tried?

Least constraining value



- Given a variable, choose the *least constraining* value — the value that leaves the maximum flexibility for subsequent variable assignments.
- Combining these makes 1000 Queens possible.

Combining Search with Inference

- Basic question Can we provide better information to these heuristics?
- Forward checking
 - Precomputing information needed by MRV
 - Early stopping
- Constraint propagation
 - Arc consistency (2-consistency)
 - n-consistency



- Can we detect inevitable failure early?
 - And avoid it later?
- Yes track remaining legal values for unassigned variables
- Terminate search when any variable has no legal values.



- Assign {*WA=red*}
- Effects on other variables connected by constraints with WA
 - NT can no longer be red
 - SA can no longer be red



- Assign {*Q=green*}
- Effects on other variables connected by constraints with WA
 - NT can no longer be green
 - NSW can no longer be green
 - SA can no longer be green



- If V is assigned blue
- Effects on other variables connected to WA
 - SA is empty
 - NSW can no longer be blue
- FC has detected a partial assignment that is *inconsistent* with the constraints.



- Solving CSPs with combination of heuristics plus forward checking is more efficient than either approach alone.
- FC checking propagates information from assigned to unassigned variables but does not provide detection for all failures.
 - NT and SA cannot be blue!
- Makes each current variable assignment arc consistent, but does not look far enough ahead to detect all inconsistencies (as AC-3 would)



• $X \rightarrow Y$ is consistent iff

for *every* value *x* of *X* there is some allowed *y*

• SA → NSW is consistent iff SA=blue and NSW=red



• $X \rightarrow Y$ is consistent iff

for *every* value *x* of *X* there is some allowed *y*

• $NSW \rightarrow SA$ is consistent iff

NSW=red and SA=blue NSW=blue and SA=???

Arc can be made consistent by removing *blue* from *NSW*



- Arc can be made consistent by removing *blue* from *NSW*
- Recheck *neighbours*
 - Remove red from *V*



- Arc can be made consistent by removing *blue* from *NSW*
- Recheck *neighbours*
 - Remove red from V
- Arc consistency detects failure earlier than forward checking
- Can be run as a preprocessor or after each assignment.
 - Repeated until no inconsistency remains

n-consistency

- Arc consistency does not detect all inconsistencies:
 - Partial assignment {WA=red, NSW=red} is inconsistent.
- Stronger forms of propagation can be defined using the notion of n-consistency.
- A CSP is n-consistent if for any set of n-1 variables and for any consistent assignment to those variables, a consistent value can always be assigned to any nth variable.
 - E.g. 1-consistency or node-consistency
 - E.g. 2-consistency or arc-consistency
 - E.g. 3-consistency or path-consistency

Further improvements

- Checking special constraints
 - Checking *Alldif*(...) constraint
 - Checking *Atmost*(...) constraint
 - Bounds propagation for larger value domains
- Intelligent backtracking
 - Standard form is chronological backtracking i.e. try different value for preceding variable.
 - More intelligent, backtrack to conflict set.
 - Set of variables that caused the failure or set of previously assigned variables that are connected to X by constraints.
 - Backjumping moves back to most recent element of the conflict set.
 - Forward checking can be used to determine conflict set.

Key Ideas so far

- Basic form of a CSP
- Different types of CSPs
- Types of constraints
- Consistent assignment
- Complete assignment
- Constraint graph
- Constraint propagation
- Backtracking search for CSPs

- Heuristics to improve backtracking search
 - MRV
 - Degree heuristic
 - Least-constraining value
- Interleaving search and inference
 - Forward checking
 - Arc consistency
 - Smart backtracking

Can this be a *local* search problem?

- Do we need the path to the solution or only the solution itself?
- Can we apply local search methods?
 - Hillclimbing
 - Simulated annealing
 - Genetic algorithms
- What's a state?



What are good heuristics for choosing moves in local search when solving CSPs?

Min-conflicts heuristic for local search

- To enable local search
 - 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)
- Actions: move queen in column
- Goal test: no attacks; *h*(*n*) = 0
- Evaluation: h(n) = number of attacks



Given random initial state, can solve *n*-queens in almost constant time for arbitrary *n* with high probability (e.g., *n* = 10,000,000). Average of 50 steps for n = 1M.

Why can local search work well?



Graph coloring (color the graph or report impossible): NP complete but "almost always easy"





Graph coloring (color the graph or report impossible): NP complete but "almost always easy"

computational cost using backtrack algorithm (with MRV, breaking ties with degree heuristic)



avg. constraints per variable

Exploiting the structure of CSPs

- Decompose into independent problems
- Tree-structured CSPs can be solved in linear time
- Reduce problems to tree-structured CSPs
 - Cycle cutset conditioning Remove nodes to create trees
 - Tree decomposition Decompose problem into a tree-structured set of subproblems



Cycle cutset conditioning

- Want to create a tree
 - What is a tree?
 - Why do we want to create one?
 - Tree-structured CSPs solvable in linear time
- Create a tree by deleting nodes
 - How can you delete nodes in CSPs?
 - Set value and restrict domains
- Does this always work well?
 - No, what can we do about that?
 - Step through possible settings
- What's the payoff?
 - Big efficiency gains O(d^c•(n-c)d²)





Tree decomposition

- Again, want to create a tree
 - What's another way of creating a tree?
 - Merging nodes
- What are the rules for doing this?
 - Every variable in \geq 1 subproblems
 - All connected variable pairs, and assoc. constraints, in ≥1 subproblems
 - If a variable appears in 2 subproblems, it must appear in all subproblems on the path connecting the two subproblems
- Now, how can we solve this new problem?



Key Ideas for CSPs

- Basic form of a CSP
- Different types of CSPs
- Types of constraints
- Consistent assignment
- Complete assignment
- Constraint graph
- Constraint propagation
- Backtracking search for CSPs
- Heuristics to improve backtracking search
 - MRV
 - Degree heuristic
 - Least-constraining value

- Interleaving search and inference
 - Forward checking
 - Arc consistency
 - Smart backtracking
- Local search
 - Min-conflicts heuristic
- Using problem structure
 - Decomposing into independent subproblems
 - Turn into a tree structured problem
 - Cutset conditioning
 - Tree decomposition

Thanks for Andrew Moore...

Note to other teachers and users of these slides. Andrew would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: http://www.cs.cmu.edu/~awm/tutorials . Comments and corrections gratefully received.

The Waltz algorithm

One of the earliest examples of a computation posed as a CSP. The Waltz algorithm is for interpreting line drawings of solid polyhedra.



an external convex intersection?

Adjacent intersections impose constraints on each other. Use CSP to find a unique set of labelings. Important step to "understanding" the image.

Waltz Alg. on simple scenes

Assume all objects:

- Have no shadows or cracks
- Three-faced vertices
- "General position": no junctions change with small movements of the eye.

Then each line on image is one of the following:

- Boundary line (edge of an object) (<) with right hand of arrow denoting "solid" and left hand denoting "space"
- Interior convex edge (+)
- Interior concave edge (-)





Given a representation of the diagram, label each junction in one of the above manners.

The junctions must be labeled so that lines are labeled consistently at both ends.

Can you formulate that as a CSP? FUN FACT: Constraint Propagation always works perfectly. Slide 36

Waltz Examples



Constraint satisfaction via a Neural Network



	R	un 1	Run 2		Run 3	
	LEFT SUBNET	R I GHT SUBNET	LEFT SUBNET	RIGHT SUBNET	LEFT SUBNET	R I GHT SUBNET
	•••	· · · ·	· · · ·	• •	· · · · ·	· · · · · ·
	• •	•••	•••	• •	• •	· •
	• •	• •	• •	• •	• •	•
	• •	• • • •	• •	• •	• • • • • • • • • • • • • • • • • • •	•
Time	•	•••	· · ·	•	• •	•
	• • •	• •	• •	• • •	•••	••••
		• •	· · · · · · · · · · · · · · · · · · ·		•	•••
		· · · ·	· · · · · · · · · · · · · · · · · · ·	•••	••	• • •
		•••				•••
						\sum

Interpretations

T: 40.40

Scheduling

A very big, important use of CSP methods.

- · Used in many industries. Makes many multi-million dollar decisions.
- · Used extensively for space mission planning.
- Military uses.

People really care about improving scheduling algorithms!

Problems with phenomenally huge state spaces. But for which solutions are needed very quickly.

Many kinds of scheduling problems e.g.:

- Job shop: Discrete time; weird ordering of operations possible; set of separate jobs.
- Batch shop: Discrete or continuous time; restricted operation of ordering; grouping is important.
- Manufacturing cell: Discrete, automated version of open job shop.

Job Shop scheduling

At a job-shop you make various products. Each product is a "job" to be done. E.G.

Job₁ = Make a polished-thing-with-a-hole

Job₂ = Paint and drill a hole in a widget

Each job requires several operations. E.G.

Operations for Job₁: Polish, Drill Operations for Job₂: Paint, Drill

Each operation needs several resources. E.G.

Polishing needs the Polishing machine Polishing needs Pat (a Polishing expert) Drilling needs the Drill Drilling needs Pat (also a Drilling expert) Or Drilling can be done by Chris

Some operations need to be done in a particular order (e.g. Paint after you've Drilled)

Job Shop Formalized

A Job Shop problem is a pair (J, RES) J is a set of jobs $J = \{j_1, j_2, ..., j_n\}$ RES is a set of resources RES = $\{R_1 ... R_m\}$

Each job j_l is specified by:

- a set of operations O' = {O'₁ O'₂ ... O'_{n(l)} }
- and must be carried out between release-date rd₁ and due-date dd₁.
- and a partial order of operations: (O'_i before O'_i), (O'_i before O'_i), etc...

Each operation O_i^l has a variable start time st_i^l and a fixed duration du_i^l and requires a set of resources. e.g.: O_i^l requires { R_{i1}^l , R_{i2}^l ... }.

Each resource can be accomplished by one of several possible physical resources, e.g. R_{i1}^l might be accomplished by any one of $\{r_{ij1}^l, r_{ij2}^l, ...\}$. Each of the r_{ijk}^l s are a member of *RES*.

Job Shop Example

 $j_1 = polished-hole-thing = \{O_1^1, O_2^1\}$ j_2 = painted-hole-widget = { O_1^2 , O_2^2 } RES = { Pat, Chris, Drill, Paint, Drill, Polisher } O¹₁ = polish-thing: need resources... $\{R_{11}^{1} = Pat, R_{12}^{1} = Polisher\}$ $O_2^1 = drill-thing: need resources...$ $\{R_{21}^{1} = (r_{211}^{1} = Pat \text{ or } r_{212}^{1} = Chris), R_{22}^{1} = Drill \}$ O²₁ = paint-widget: need resources... $\{ R^{2}_{11} = Paint \}$ O_{2}^{2} = drill-widget : need resources... $\{R_{21}^2 = (r_{211}^2 = Pat \text{ or } r_{212}^2 = Chris), R_{222}^2 = Drill\}$ Precedence constraints : O_2^2 before O_4^2 . All operations take one time unit du_i^1 = 1 forall *i*,*l*. Both jobs have release-date *rd*^{*l*} = 0 and due-date *dd*^{*l*} = 1.

Job-shop: the Variables and Constraints

Variables

- The operation state times stⁱ
- The resources R^I_{ij} (usually these are obvious from the definition of O^I_i. Only need to be assigned values when there are alternative physical resources available, e.g. *Pat* or *Chris* for operating the *drill*). Constraints:
- Precedence constraints. (Some O'_is must be before some other O'_is).
- Capacity constraints. There must never be a pair of operations with overlapping periods of operation that use the same resources.

Non-challenging question. Can you schedule our Job-shop?

Next Class: Tuesday Oct 18

- Problem set 2 due
- Review for midterm:
 Everything except Constraint Satisfaction