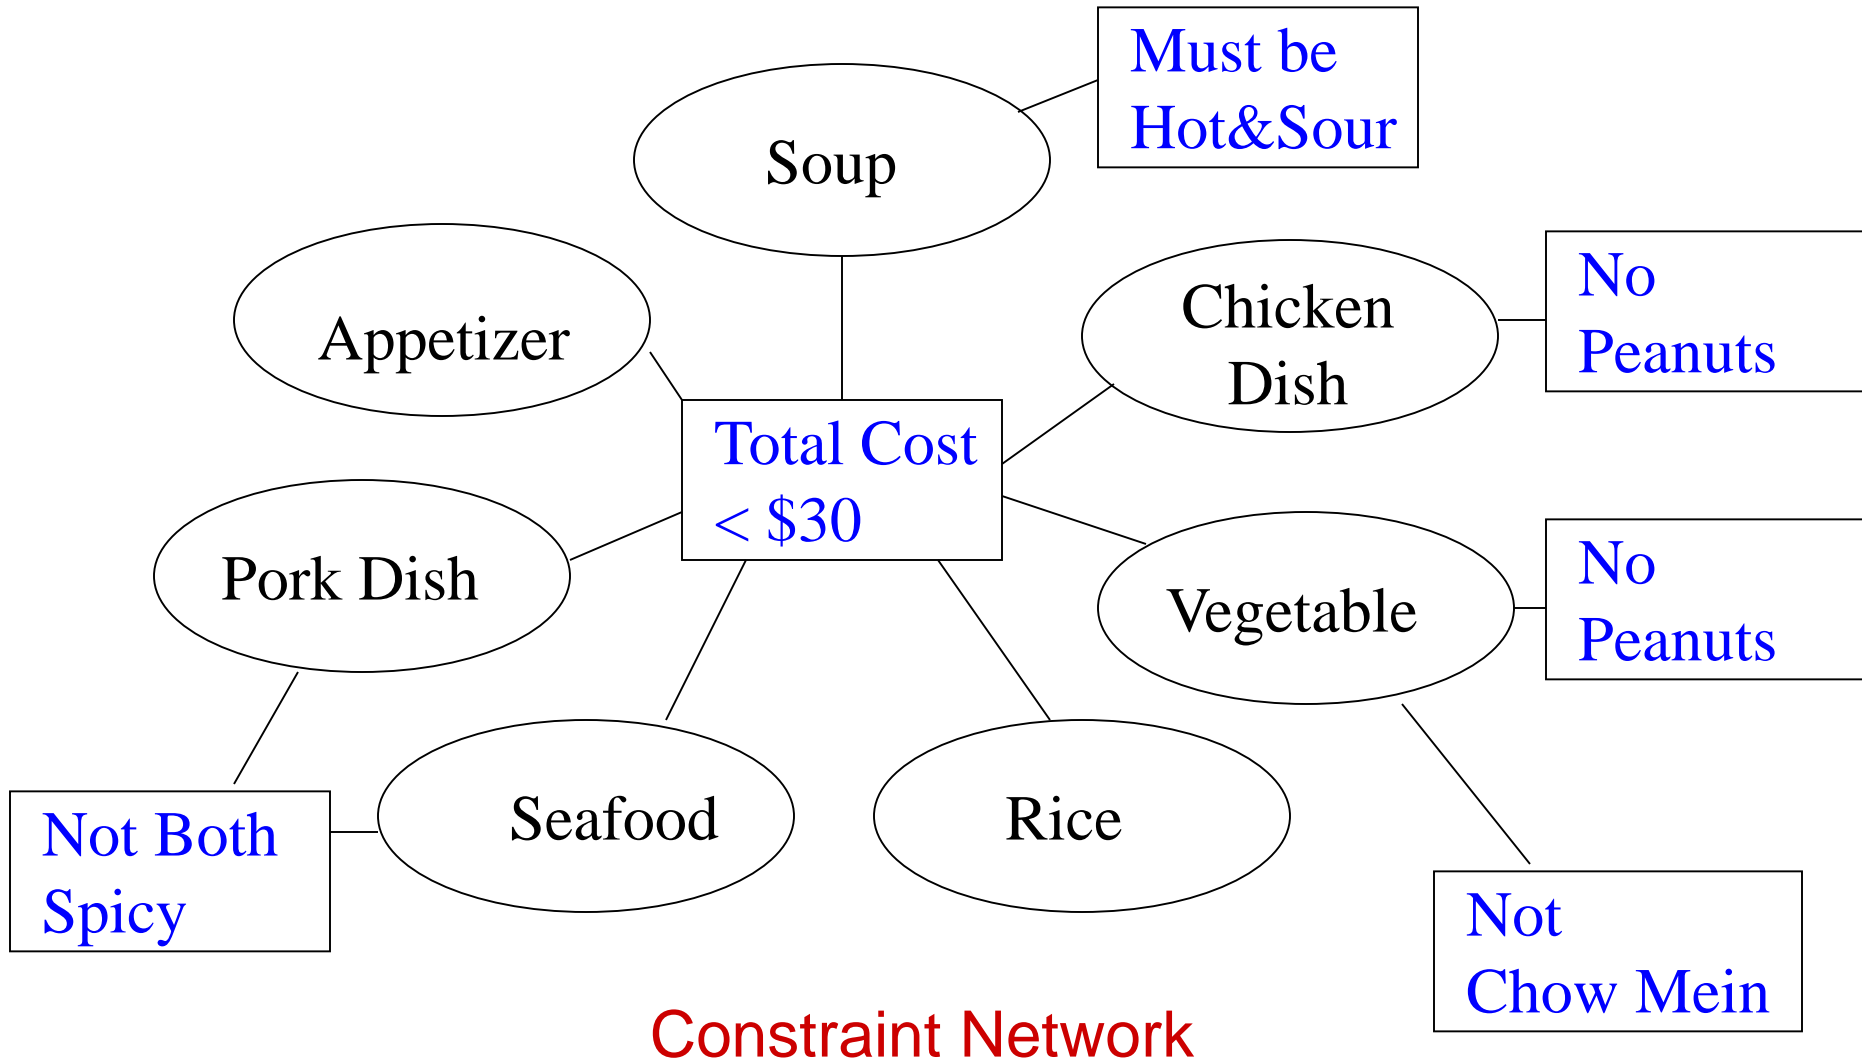


Constraint Satisfaction Problems



Formal Definition of CSP

- A constraint satisfaction problem (**CSP**) is a triple (V, D, C) where
 - V is a set of variables X_1, \dots, X_n .
 - D is the union of a set of domain sets D_1, \dots, D_n , where D_i is the domain of possible values for variable X_i .
 - C is a set of constraints on the values of the variables, which can be pairwise (simplest and most common) or k at a time.

CSPs vs. Standard Search Problems

- Standard search problem:
 - **state** is a "black box" – any data structure that supports successor function, heuristic function, and goal test
- CSP:
 - **state** is defined by **variables** X_i with **values** from **domain** D_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

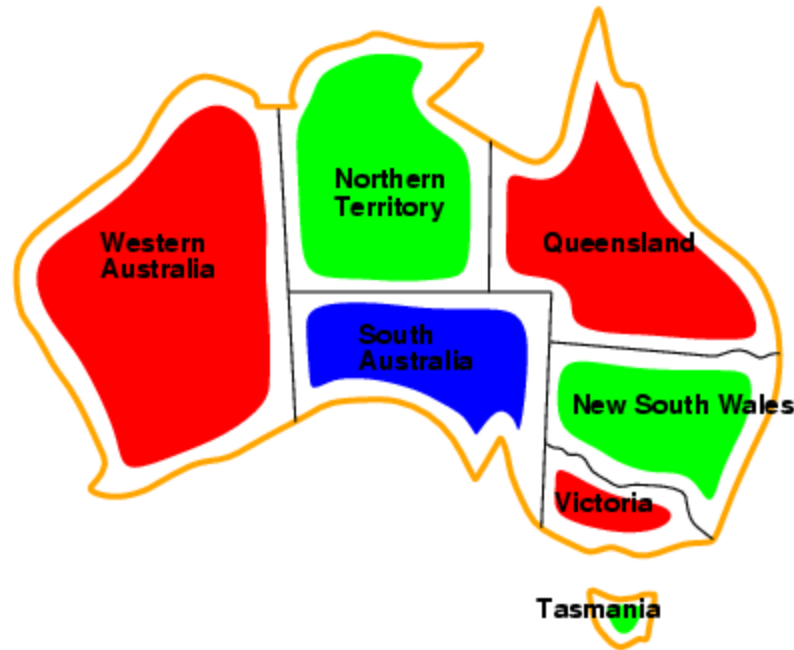
Example: Map-Coloring



memorize
the names

- Variables WA, NT, Q, NSW, V, SA, T
- Domains $D_i = \{\text{red, green, blue}\}$
- Constraints: adjacent regions must have different colors
- e.g., $WA \neq NT$, or (WA, NT) in $\{(\text{red, green}), (\text{red, blue}), (\text{green, red}), (\text{green, blue}), (\text{blue, red}), (\text{blue, green})\}$

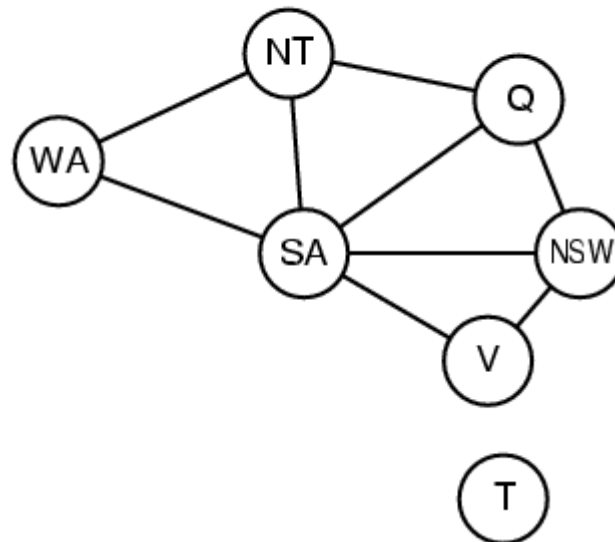
Example: Map-Coloring



- Solutions are **complete** and **consistent** assignments, e.g., WA = red, NT = green, Q = red, NSW = green, V = red, SA = blue, T = green

Constraint graph

- **Binary CSP:** each constraint relates two variables
- **Constraint graph:** nodes are variables, arcs are constraints

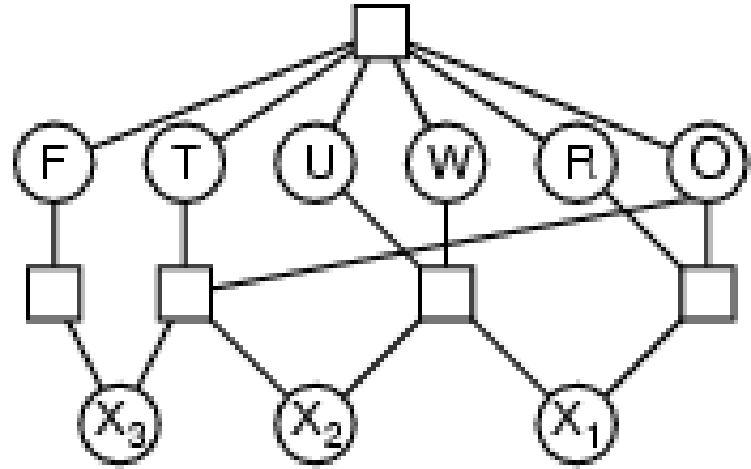


Varieties of constraints

- **Unary** constraints involve a single variable,
 - e.g., $SA \neq \text{green}$
- **Binary** constraints involve pairs of variables,
 - e.g., $\text{value}(SA) \neq \text{value}(WA)$
 - More formally, $R1 \leftrightarrow R2 \rightarrow \text{value}(R1) \leftrightarrow \text{value}(R2)$
- **Higher-order** constraints involve 3 or more variables,
 - e.g., cryptarithmic column constraints

Example: Cryptarithmic





$$\begin{array}{r} \text{ T W O} \\ + \text{ T W O} \\ \hline \text{ F O U R} \end{array}$$



- 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: *Alldiff* (F, T, U, W, R, O)
 - $O + O = R + 10 \cdot X_1$
 - $X_1 + W + W = U + 10 \cdot X_2$
 - $X_2 + T + T = O + 10 \cdot X_3$
 - $X_3 = F, T \neq 0, F \neq 0$

Example: Latin Squares Puzzle

X_{11}	X_{12}	X_{13}	X_{14}
X_{21}	X_{22}	X_{23}	X_{24}
X_{31}	X_{32}	X_{33}	X_{34}
X_{41}	X_{42}	X_{43}	X_{44}

				
red	RT	RS	RC	RO
green	GT	GS	GC	GO
blue	BT	BS	BC	BO
yellow	YT	YS	YC	YO

Variables

Values

Constraints: In each row, each column, each major diagonal, there must be no two markers of the same color or same shape.

How can we formalize this?

V: $\{X_{ij} \mid i=1 \text{ to } 4 \text{ and } j=1 \text{ to } 4\}$

D: $\{(C,S) \mid C \in \{R,G,B,Y\} \text{ and } S \in \{T,S,C,O\}\}$

C: $\text{val}(X_{ij}) \neq \text{val}(X_{in})$ if $j \neq n$ (same row)

$\text{val}(X_{ij}) \neq \text{val}(X_{ni})$ if $i \neq n$ (same col)

$\text{val}(X_{ij}) \neq \text{val}(X_{il})$ if $j \neq l$ (one diag)

$i+l=n+m=5 \rightarrow \text{val}(X_{ij}) \neq \text{val}(X_{nm}), i \neq nm$

Real-world CSPs

- **Assignment problems**
 - e.g., who teaches what class
- **Timetabling problems**
 - e.g., which class is offered when and where?
- **Transportation scheduling**
- **Factory scheduling**

Notice that many real-world problems involve real-valued variables

The Consistent Labeling Problem

- Let $P = (V, D, C)$ be a constraint satisfaction problem.
- An **assignment** is a partial function $f : V \rightarrow D$ that assigns a value (from the appropriate domain) to each variable
- A consistent assignment or **consistent labeling** is an assignment f that satisfies all the constraints.
- A **complete consistent labeling** is a consistent labeling in which every variable has a value.

Standard Search Formulation

- **state:** (partial) assignment
 - **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
(and is a consistent labeling)
1. This is the same for all CSPs regardless of application.
 2. Every solution appears at depth n with n variables
→ we can use depth-first search.
 3. Path is irrelevant, so we can also use complete-state formulation.

What Kinds of Algorithms are used for CSP?

- Backtracking Tree Search
- Tree Search with Forward Checking
- Tree Search with Discrete Relaxation (arc consistency, k-consistency)
- Many other variants
- Local Search using Complete State Formulation

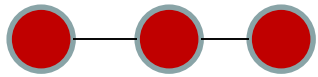
Backtracking Tree Search

- 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.
- **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 $n \approx 25$.

Subgraph Isomorphisms

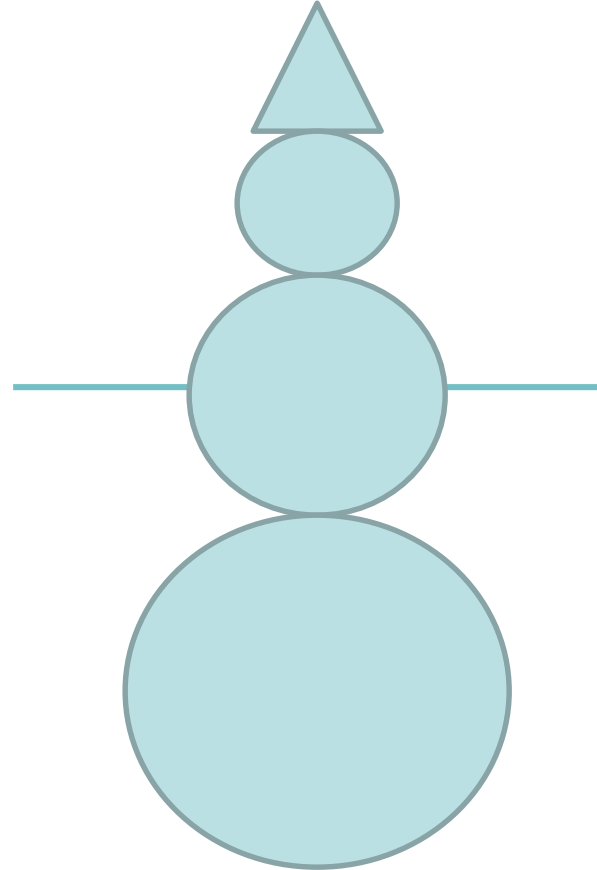
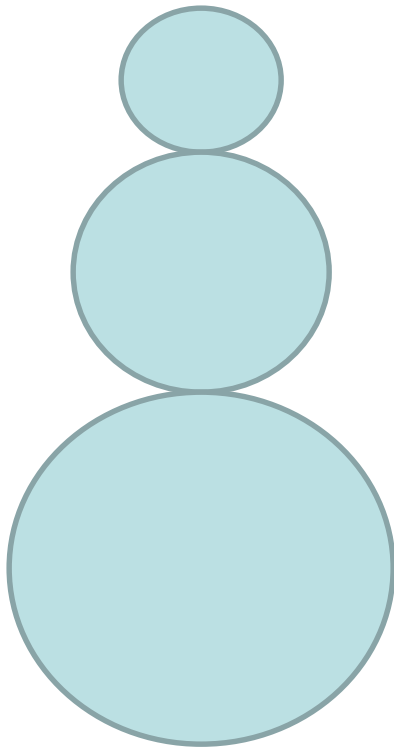
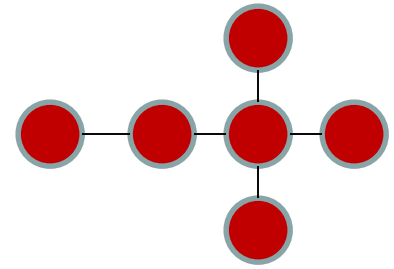
- Given 2 graphs $G1 = (V,E)$ and $G2 = (W,F)$.
- Is there a copy of $G1$ in $G2$?

- V is just itself, the vertices of $G1$
- $D = W$
- $f: V \rightarrow W$
- $C: (v1,v2) \in E \Rightarrow (f(v1),f(v2)) \in F$



adjacency relation

Example

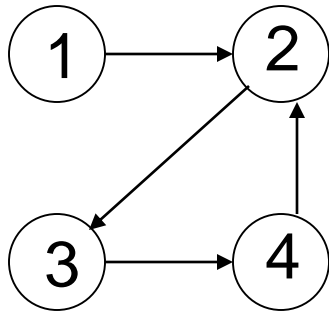


Is there a copy of the snowman on the left in the picture on the right?

Graph Matching Example

Find a subgraph isomorphism from R to S.

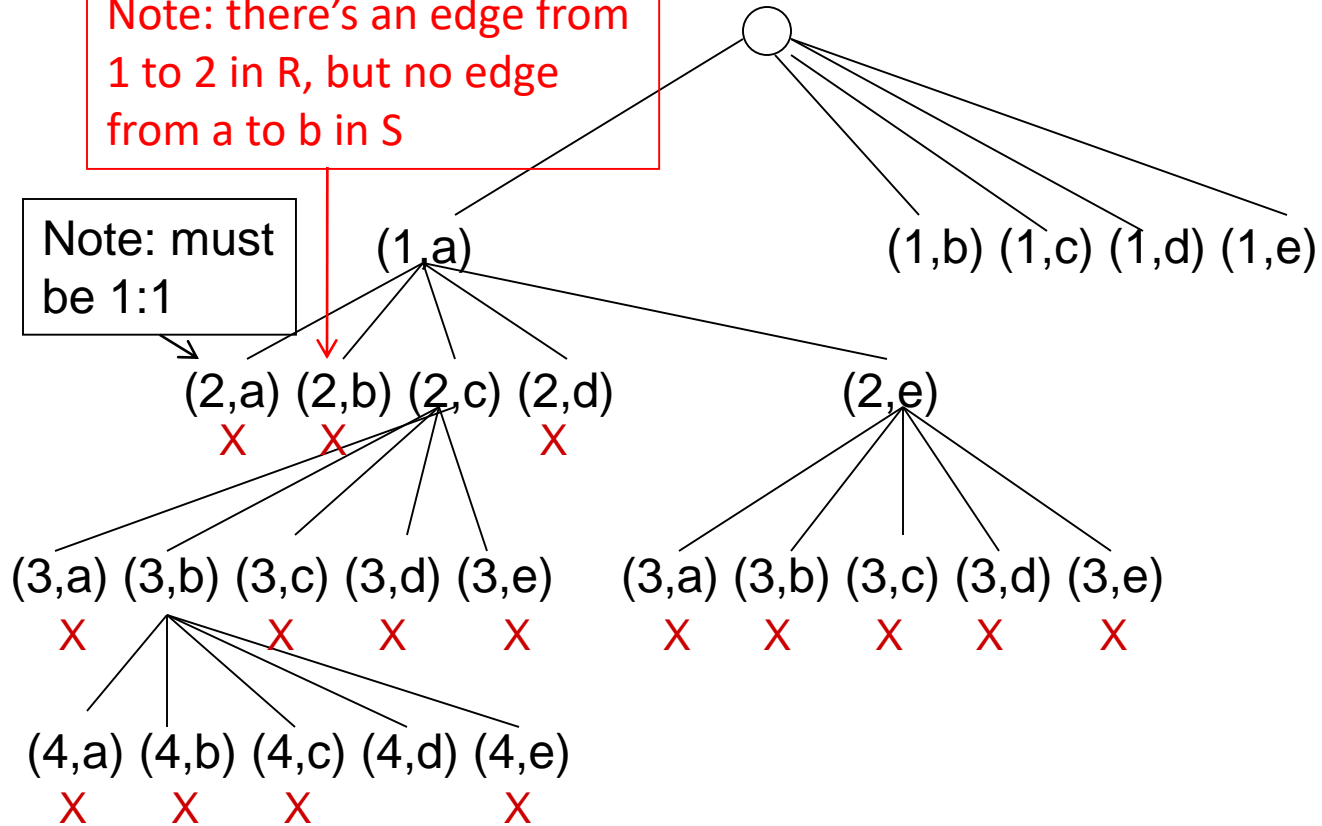
R



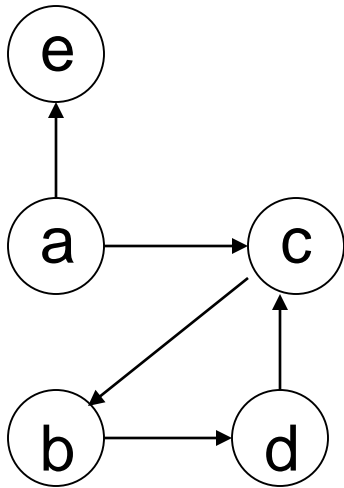
“snowman”

Note: there's an edge from 1 to 2 in R, but no edge from a to b in S

Note: must be 1:1



S



“snowman with hat and arms”

Backtracking Search

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

function RECURSIVE-BACKTRACKING(assignment, csp) returns a solution, or
failure
  if assignment is complete then return assignment
  1. var ← SELECT-UNASSIGNED-VARIABLE(Variables[csp], assignment, csp)
  2. for each value in ORDER-DOMAIN-VALUES(var, assignment, csp) do
  3.   if value is consistent with assignment according to 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
```

1. One variable at each tree level
2. Try all values for that variable (depth first)
3. Check for consistency, backup if not consistent

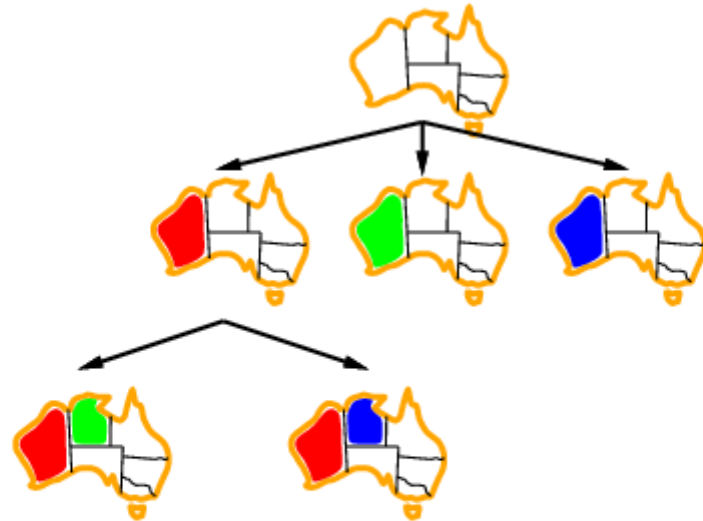
Backtracking Example



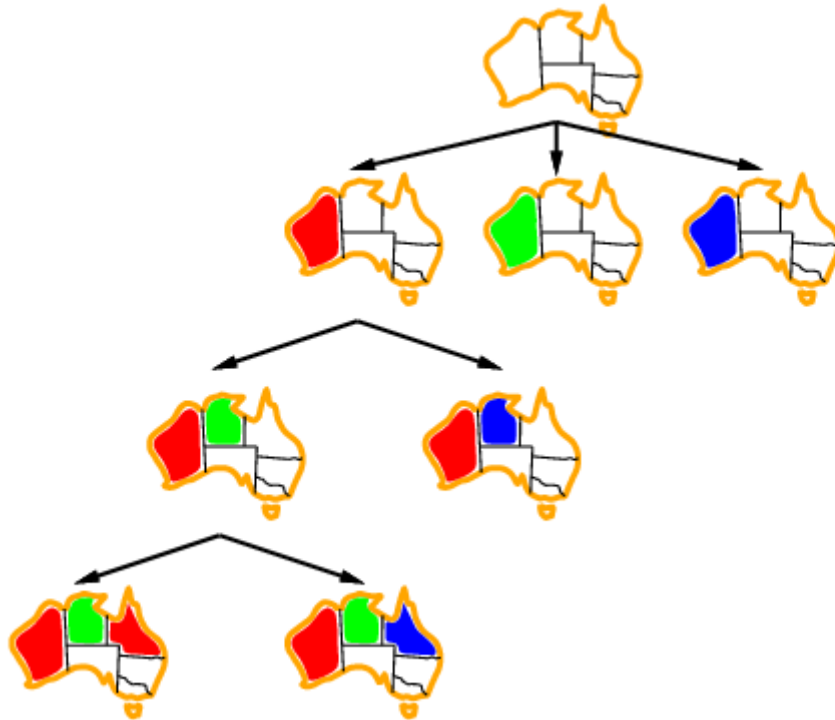
Backtracking Example



Backtracking Example



Backtracking Example

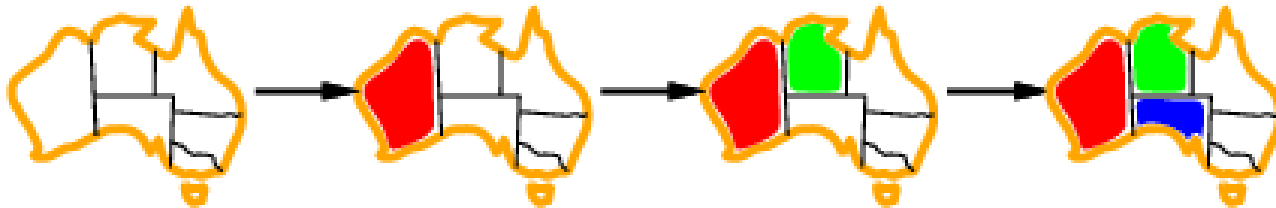


Improving Backtracking Efficiency

- **General-purpose** methods can give huge gains in speed:
 - Which variable should be assigned next?
 - In what order should its values be tried?
 - Can we detect inevitable failure early?

Most Constrained Variable

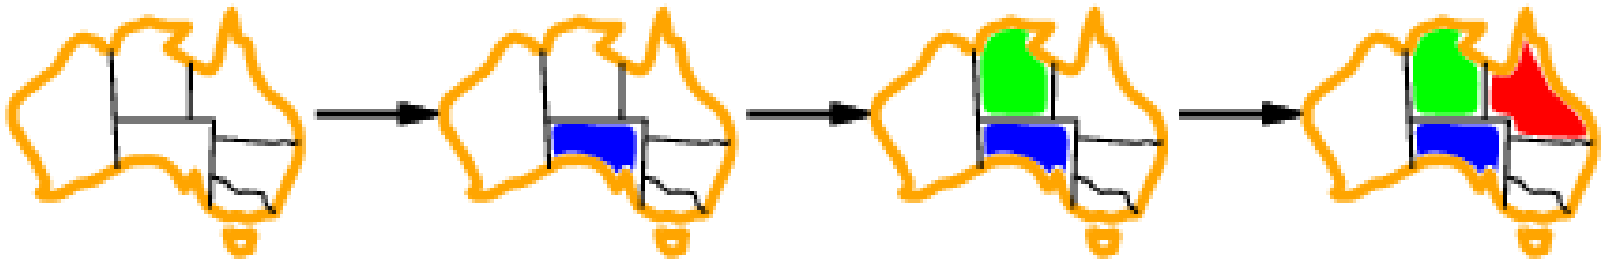
- Most constrained variable:
choose the variable with the fewest legal values



- a.k.a. **minimum remaining values (MRV)**
heuristic

Most Constraining Variable

- Tie-breaker among most constrained variables
- Most constraining variable:
 - choose the variable with the most constraints on remaining variables



Least Constraining Value

- Given a variable, choose the least constraining value:
 - the one that rules out the fewest values in the remaining variables



- Combining these heuristics makes 1000 queens feasible

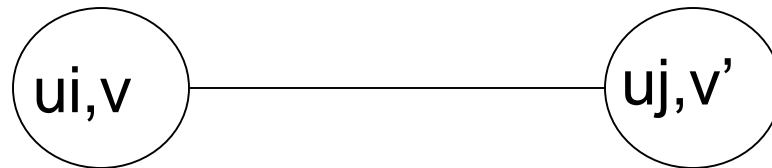
Forward Checking

(Haralick and Elliott, 1980)

Variables: $U = \{u_1, u_2, \dots, u_n\}$

Values: $V = \{v_1, v_2, \dots, v_m\}$

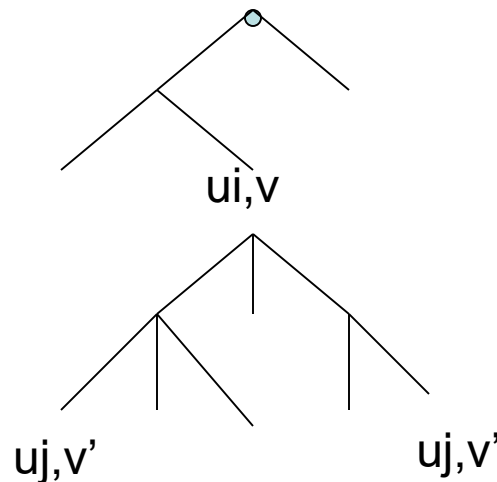
Constraint Relation: $R = \{(u_i, v, u_j, v') \mid u_i \text{ having value } v \text{ is compatible with } u_j \text{ having label } v'\}$



If (u_i, v, u_j, v') is not in R , they are incompatible, meaning if u_i has value v , u_j cannot have value v' .

Forward Checking

Forward checking is based on the idea that once variable u_i is assigned a value v , then certain future variable-value pairs (u_j, v') become impossible.



Instead of finding this out at many places on the tree, we can rule it out in advance.

Data Structure for Forward Checking

Future error table (FTAB)

One per level of the tree (ie. a stack of tables)

	v1	v2	...	vm
u1				
u2				
:				
un				

What does it mean if a whole row becomes 0?

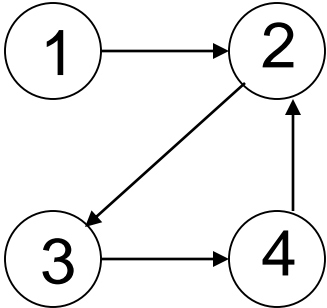
At some level in the tree,

for future (unassigned) variables u

$FTAB(u,v) = 1$ if it is still possible to assign v to u
 0 otherwise

Graph Matching Example

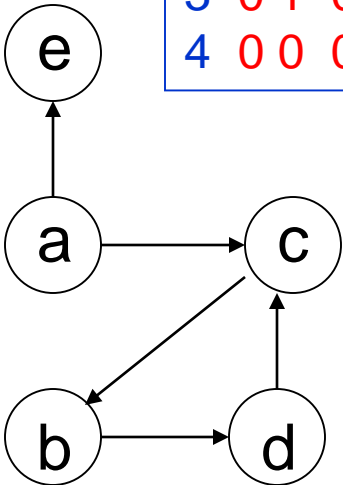
R



	a	b	c	d	e
2	0	0	1	0	1
3	0	1	1	1	1
4	0	1	1	1	1

	a	b	c	d	e
1	1	1	1	1	1
2	1	1	1	1	1
3	1	1	1	1	1
4	1	1	1	1	1

S



	a	b	c	d	e
3	0	1	0	0	0
4	0	0	0	1	0

	a	b	c	d	e
3	0	0	0	0	0
4		X			

(1,a)

(1,b) (1,c) (1,d) (1,e)

(2,c)

(2,e)

(3,b)

(4,d)

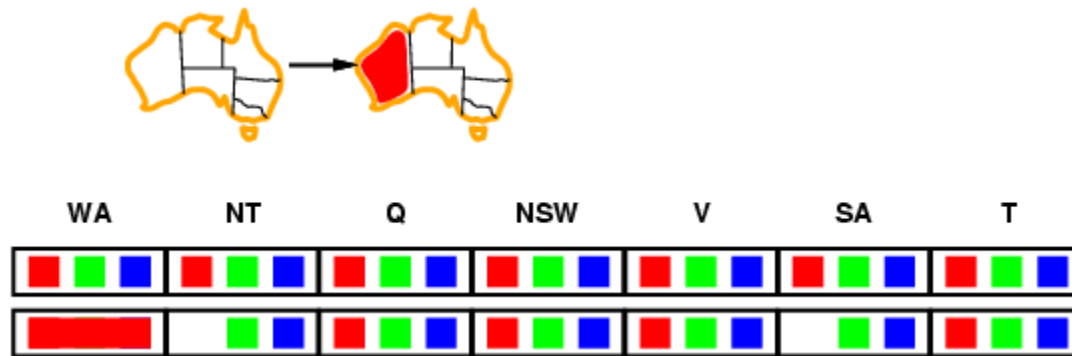
Book's Forward Checking Example

- Idea:
 - Keep track of remaining legal values for unassigned variables
 - Terminate search when any variable has no legal values



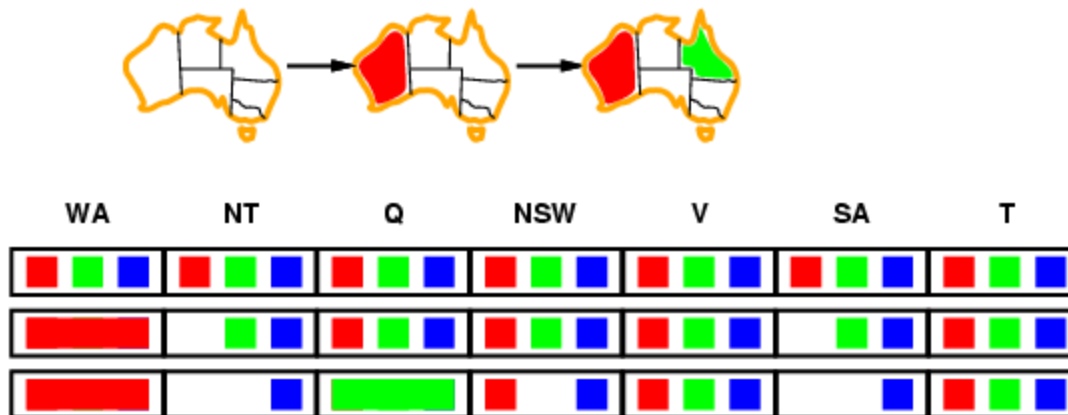
Forward Checking

- Idea:
 - Keep track of remaining legal values for unassigned variables
 - Terminate search when any variable has no legal values



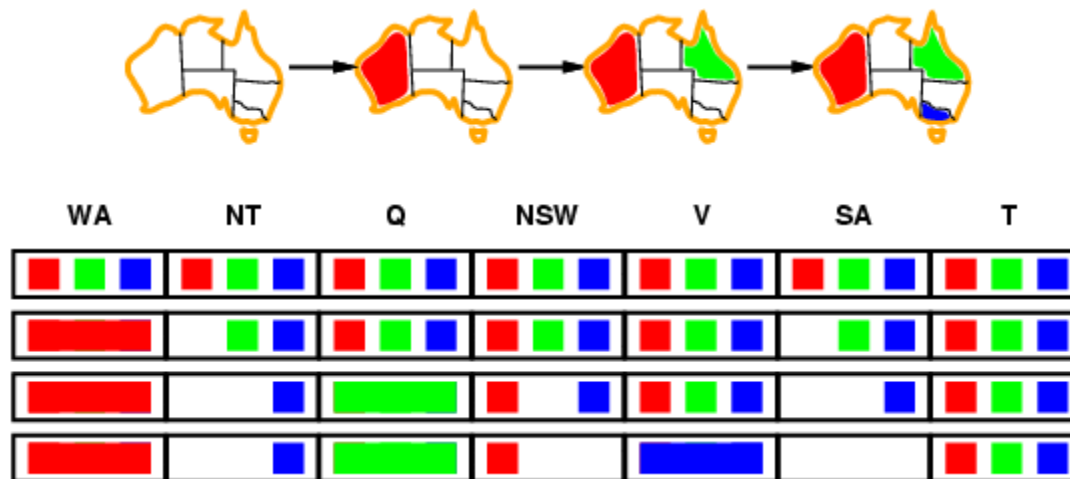
Forward Checking

- Idea:
 - Keep track of remaining legal values for unassigned variables
 - Terminate search when any variable has no legal values



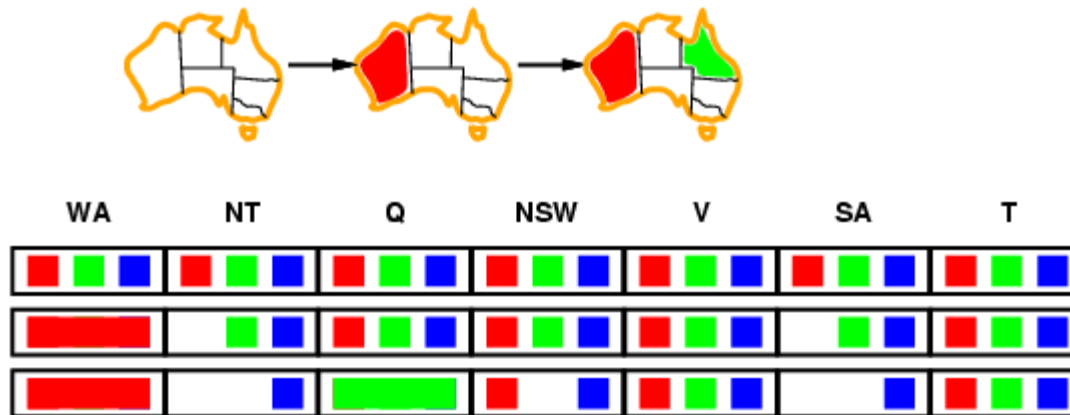
Forward Checking

- Idea:
 - Keep track of remaining legal values for unassigned variables
 - Terminate search when any variable has no legal values



Constraint Propagation

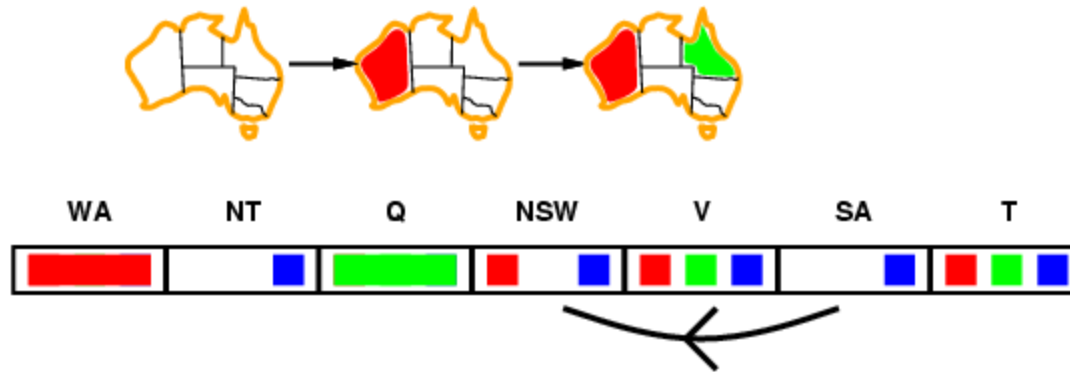
- Forward checking propagates information from assigned to unassigned variables, but doesn't provide early detection for all failures:



- NT and SA cannot both be blue!
- Constraint propagation** repeatedly enforces constraints locally

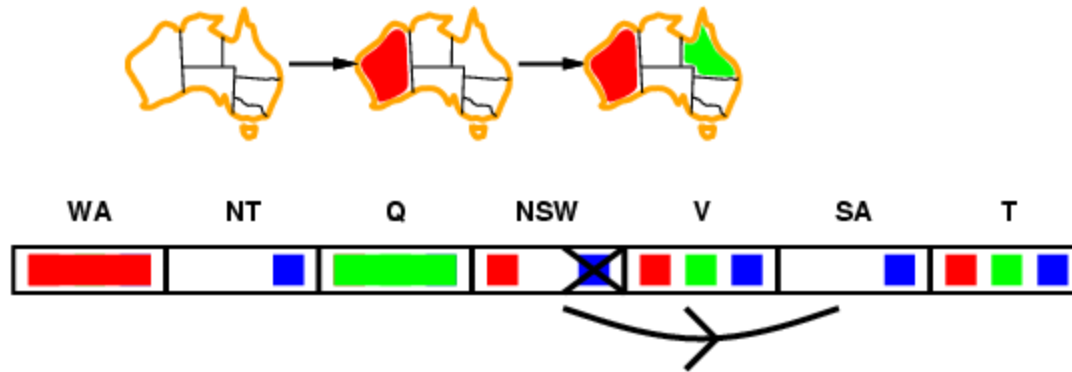
Arc Consistency

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



Arc Consistency

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



Putting It All Together

- backtracking tree search
- with forward checking
- add arc-consistency
 - For each pair of future variables (u_i, u_j) that constrain one another
 - Check each possible remaining value v of u_i
 - Is there a compatible value w of u_j ?
 - If not, remove v from possible values for u_i
(set $FTAB(u_i, v)$ to 0)

Comparison of Methods

- Backtracking tree search is a blind search.
- Forward checking checks constraints between the current variable and all future ones.
- Arc consistency then checks constraints between all pairs of future (unassigned) variables.

- What is the complexity of a backtracking tree search?
- How do forward checking and arc consistency affect that?

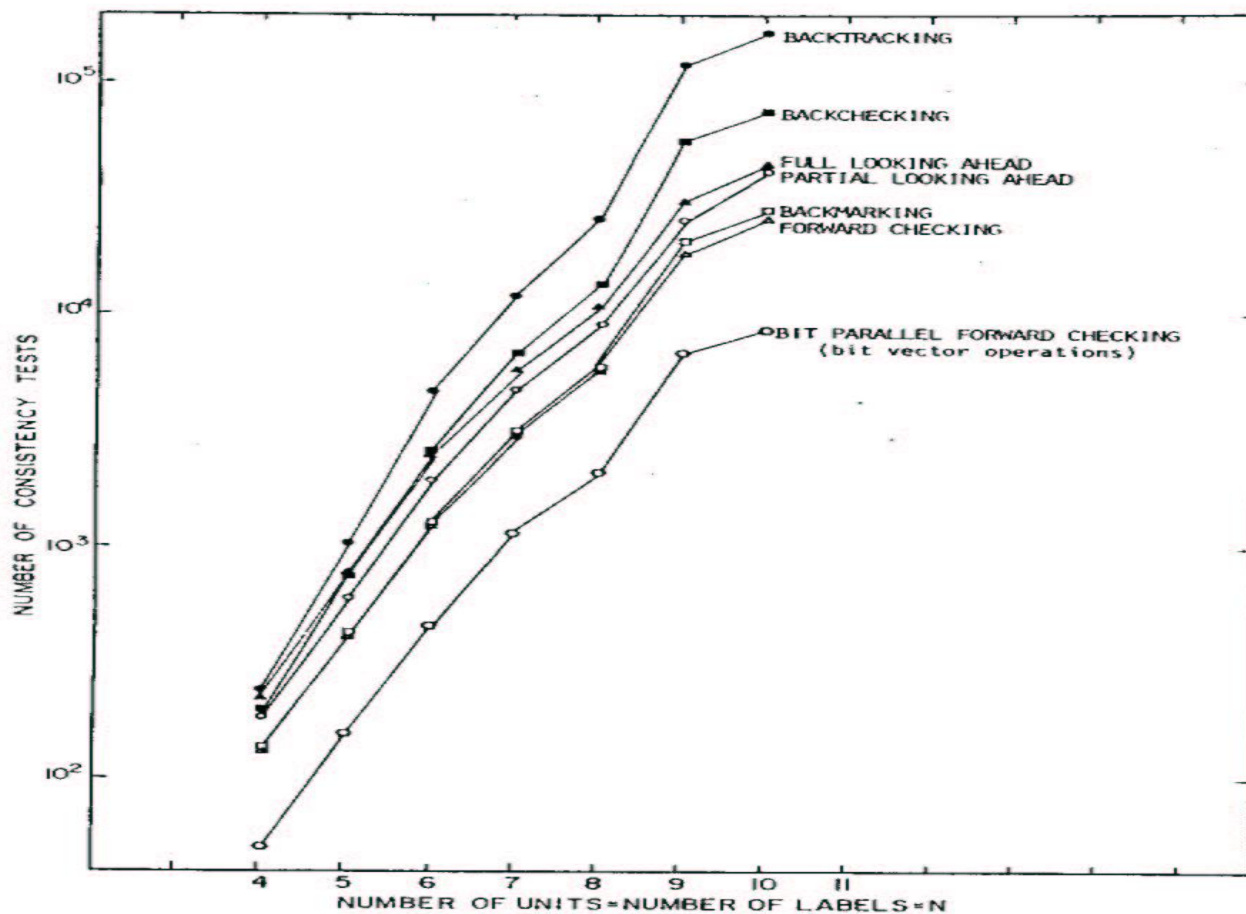


FIG. 8. The number of consistency tests in the average of 5 runs of the indicated programs. Relations are random with consistency check probability $p = 0.65$ and number of units = number of labels = N . Each random relation is tested on all 6 methods, using the same 5 different relations generated for each N . The number of bit-vector operations in bit parallel forward checking is also shown for the same relations.

Inexact Matching

- All of this expects things to be perfect.
- But that's not the case.
- Real data is “noisy”.
- Especially in computer vision or in speech recognition.
- What's the simplest thing we can do?
- Allow some error.

Inexact Matching

- An object model can be made up of its parts and their relationships (like the snowman).
- Not all the parts show up every time.
- Not all the relationships are detected correctly.

Example: Characters

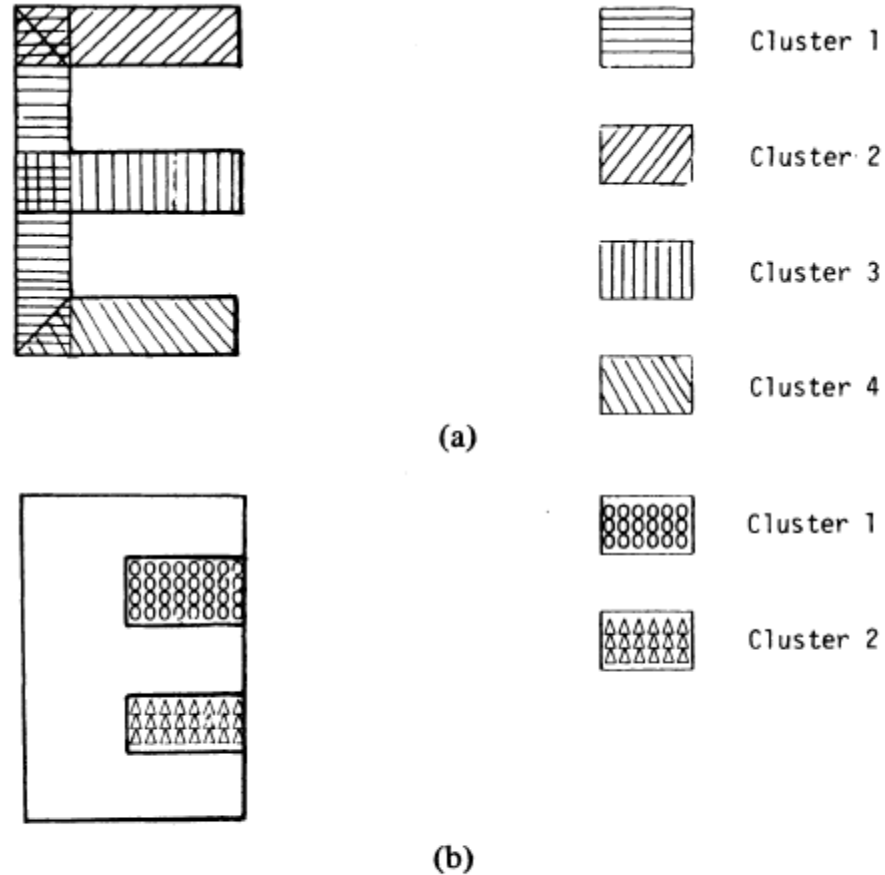
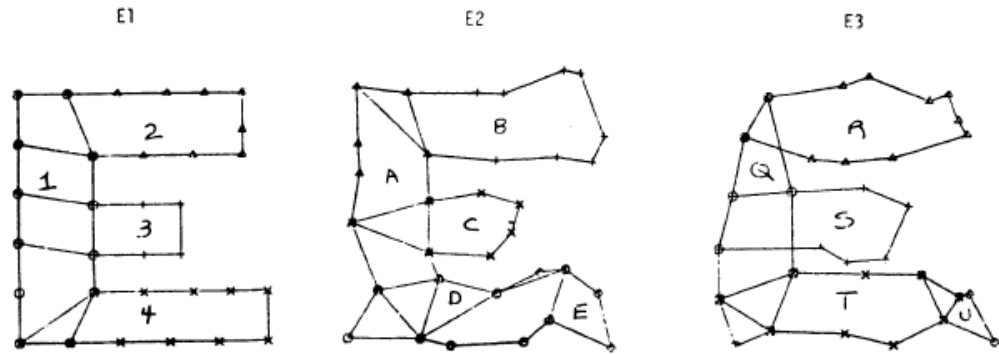
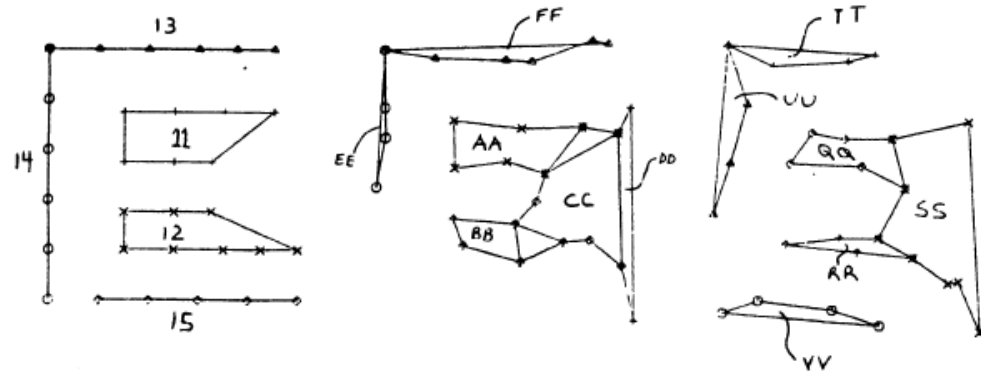


Fig. 1. Illustrates the decomposition of a shape into simple parts and intrusions. (a) Shows the simple parts or interior clusters of a hand-printed letter E. (b) Shows the intrusions or exterior clusters of the hand-printed letter E.

Samples



Decomposition into simple parts



Decomposition into intrusions

Subgraph Isomorphism Extension

- Given 2 graphs $G1 = (V,E)$ and $G2 = (W,F)$.
- Is there a copy of $G1$ in $G2$?
- V is just itself, the vertices of $G1$
- $D = W$
- $f: V \rightarrow W$
- $C: (v1,v2) \in E \Rightarrow (f(v1),f(v2)) \in F$
- **Inexact:** $\sum_{(v1,v2) \in E \text{ and } (f(v1),f(v2)) \notin F} \text{weight}((v1,v2)) < \text{a threshold}$

Algorithms

- Can we still do backtracking tree search?
- YES. Just keep going down a path as long as the error does not exceed the threshold.
- Can we still do forward checking?
- YES. I published the algorithm myself. It still keeps the FTAB, but now the FTAB keeps track of the error so far.

One Step Further

- Relational distance allows us to just find the BEST match between two graphs instead of restricting it to a specific error threshold.
- Given two graphs (or higher-dimensional relational descriptions of objects), what is the least error mapping f from one to the other?
- The error of that mapping is the relational distance between the two graphs.

Relational Distance

A relational description is a data structure that may be used to describe two dimensional shape models, three-dimensional object models, regions on an image, and so on.

Definition A **relational description** DP is a sequence of relations $DX = \{R_1, \dots, R_I\}$.

Let $DA = \{R_1, \dots, R_I\}$ be a relational description with part set A and $DB = \{S_1, \dots, S_I\}$ be a relational description with part set B .

Let f be any one-one, onto mapping from A to B.

The **structural error** of f for the i th pair of corresponding relation (R_i and S_i) in DA and DB is given by

$$| f(R_i) - S_i | + | f^{-1}(S_i) - R_i |$$

The structural error indicates how many tuples in R_i are *not* mapped by f to tuples in S_i and how many tuples in S_i are *not* mapped by f^{-1} to tuples in R_i .

The **total error of f** with respect to DA and DB is the sum of the structural errors for each pair of corresponding relations.

The **relational distance** is the minimal total error obtained for any one-one, onto mapping f from A to B , and that mapping is called the **best mapping**.

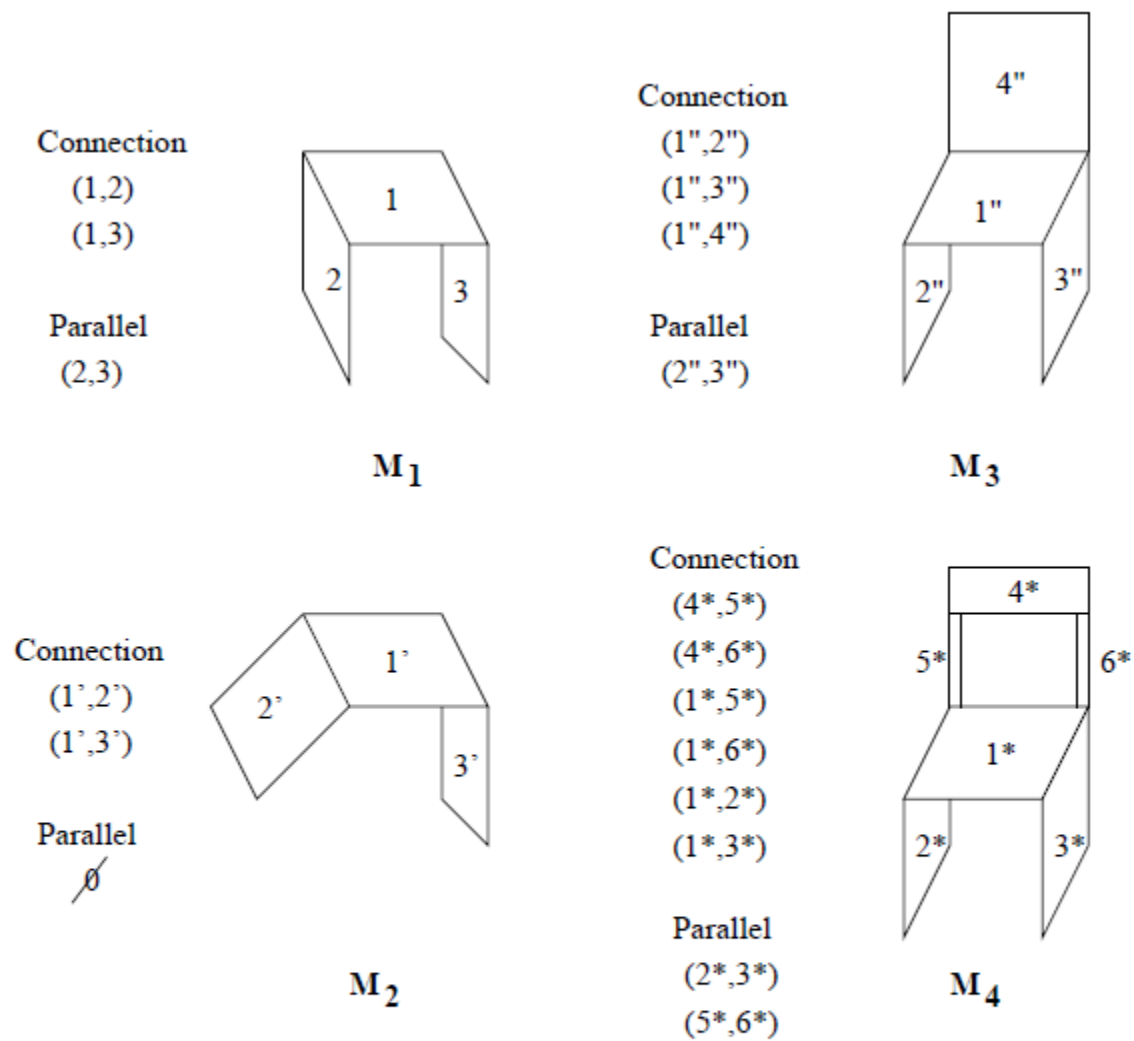


Figure 11.22: Four object models. The relational distance of model M_1 to M_2 and M_1 to M_3 is 1. The relational distance of model M_3 to M_4 is 6.

Summary

- CSPs are a special kind of 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 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
- Searches are still worst case exponential, but pruning keeps the time down.