

EECS 3401: Intro to AI & LP CSPs & Backtracking Search

- Readings: R & N Chapter 6.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

1

Constraint Satisfaction Problems

- The search algorithms we discussed so far had no knowledge of the states representation (black box). So, we could not take advantage of domain-specific information.
- CSP are a special class of search problems with a **uniform and simple** state representation.
- This allows to design more efficient algorithms.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

2

Constraint Satisfaction Problems

- Many problems can be represented as a search for a vector of feature values.
 - k-features: variables.
 - Each feature has a value. Domain of values for the variables.
 - e.g., height = {short, average, tall}, weight = {light, average, heavy}.
- In these problems the problem is to search for a set of values for the features (variables) so that the values satisfy some conditions (constraints).

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

3

Constraint Satisfaction Problems

- Sudoku:
 - 81 variables, the value in each cell.
 - Values: a fixed value for those cells that are already filled in, the values {1-9} for those cells that are empty.
 - Solution: a value for each cell satisfying the constraints:
 - no cell in the same column can have the same value.
 - no cell in the same row can have the same value.
 - no cell in the same sub-square can have the same value.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

4

Constraint Satisfaction Problems

- Scheduling
 - Want to schedule a time and a space for each final exam so that
 - No student is scheduled to take more than one final at the same time.
 - The space allocated has to be available at the time set.
 - The space has to be large enough to accommodate all of the students taking the exam.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

5

Constraint Satisfaction Problems

- Variables:
 - T_1, \dots, T_m : T_i is a variable representing the scheduled time for the i -th final.
 - Assume domains are fixed to $\{\text{MonAm}, \text{MonPm}, \dots, \text{FriAm}, \text{FriPm}\}$.
 - S_1, \dots, S_m : S_i is the space variable for the i -th final.
 - Domain of S_i are all rooms big enough to hold the i -th final.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

6

Constraint Satisfaction Problems

- Want to find an assignment of values to each variable (times, rooms for each final), subject to the constraints:
 - For all pairs of finals i, j such that there is a student taking both:
 - $T_i \neq T_j$
 - For all pairs of finals i, j
 - $T_i \neq T_j$ or $S_i \neq S_j$
 - either i and j are not scheduled at the same time, or if they are they are not in the same space.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

7

Constraint Satisfaction Problems (CSP)

- More formally.
- A CSP consists of
 - a set of variables V_1, \dots, V_n
 - for each variable a domain of possible values $\text{Dom}[V_i]$.
 - A set of constraints C_1, \dots, C_m .

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

8

Constraint Satisfaction Problems

- Each variable be assigned any value from its domain.
 - $V_i = d$ where $d \in \text{Dom}[V_i]$
- Each constraint C has
 - A set of variables it is over, called its scope: e.g., $C(V_1, V_2, V_4)$.
 - Is a boolean function that maps assignments to these variables to true/false.
 - e.g. $C(V_1=a, V_2=b, V_4=c) = \text{True}$
 - this set of assignments satisfies the constraint.
 - e.g. $C(V_1=b, V_2=c, V_4=c) = \text{False}$
 - this set of assignments falsifies the constraint.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

9

Constraint Satisfaction Problems

- A solution to a CSP is
 - an assignment of a value to all of the variables such that
 - every constraint is satisfied.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

10

Constraint Satisfaction Problems

- Sudoku:
 - $V_{11}, V_{12}, \dots, V_{21}, V_{22}, \dots, V_{91}, \dots, V_{99}$
 - $\text{Dom}[V_{ij}] = \{1-9\}$ for empty cells
 - $\text{Dom}[V_{ij}] = \{k\}$ a fixed value k for filled cells.
 - Row constraints:
 - $\text{CR1}(V_{11}, V_{12}, V_{13}, \dots, V_{19})$
 - $\text{CR2}(V_{21}, V_{22}, V_{23}, \dots, V_{29})$
 - $\dots, \text{CR9}(V_{91}, V_{92}, \dots, V_{99})$
 - Column Constraints:
 - $\text{CC1}(V_{11}, V_{21}, V_{31}, \dots, V_{91})$
 - $\text{CC2}(V_{12}, V_{22}, V_{32}, \dots, V_{92})$
 - $\dots, \text{CC9}(V_{19}, V_{29}, \dots, V_{99})$
 - Sub-Square Constraints:
 - $\text{CSS1}(V_{11}, V_{12}, V_{13}, V_{21}, V_{22}, V_{23}, V_{31}, V_{32}, V_{33})$
 - $\text{CSS1}(V_{14}, V_{15}, V_{16}, \dots, V_{34}, V_{35}, V_{36})$

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

11

Constraint Satisfaction Problems

- Sudoku:
 - Each of these constraints is over 9 variables, and they are all the same constraint:
 - Any assignment to these 9 variables such that each variable has a unique value satisfies the constraint.
 - Any assignment where two or more variables have the same value falsifies the constraint.
 - Such constraints are often called ALL-DIFF constraints.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

12

Constraint Satisfaction Problems

- Sudoku:
 - Thus Sudoku has 3x9 ALL-Diff constraints, one over each set of variables in the same row, one over each set of variables in the same column, and one over each set of variables in the same sub-square.
 - Note also that an ALL-Diff constraint over k variables can be equivalently represented by k choose 2 not-equal constraints over each pair of these variables.
 - e.g. $CSS1(V_{11}, V_{12}, V_{13}, V_{21}, V_{22}, V_{23}, V_{31}, V_{32}, V_{33}) = NEQ(V_{11}, V_{12}), NEQ(V_{11}, V_{13}), NEQ(V_{11}, V_{21}) \dots, NEQ(V_{32}, V_{33})$
 - NEQ is a not-equal constraint.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

13

Constraint Satisfaction Problems

- Exam Scheduling
 - constraints:
 - For all pairs of finals i, j such that there is a student taking both:
 - NEQ(Ti, Tj)
 - For all pairs of finals i, j
 - C(Ti, Tj, Si, Sj)
 - This constraint is satisfied
 - by any set of assignments in which $T_i \neq T_j$.
 - any set of assignments in which $S_i \neq S_j$.
 - Falsified by any set of assignments in which $T_i = T_j$ as well as $S_i = S_j$.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

14

Solving CSPs

- CSPs can be solved by a specialized version of depth first search.
- Key intuitions:
 - We can build up to a solution by searching through the space of partial assignments.
 - Order in which we assign the variables does not matter---eventually they all have to be assigned.
 - If during the process of building up a solution we falsify a constraint, we can immediately reject all possible ways of extending the current partial assignment.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

15

Backtracking Search

- These ideas lead to the backtracking search algorithm

```

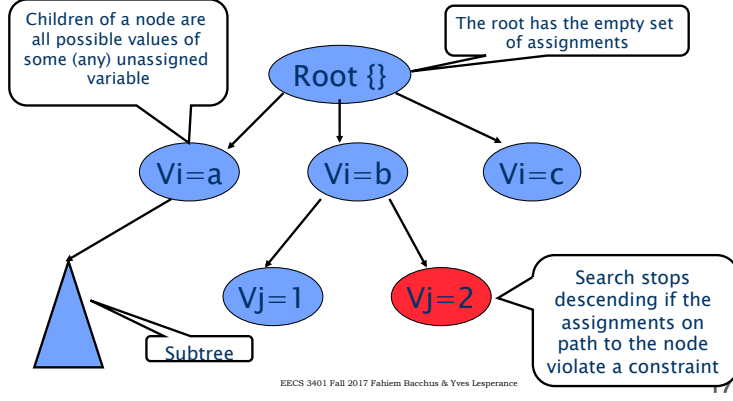
Algorithm BT (Backtracking)
BT(Level)
  If all variables assigned
    PRINT value of each variable
    RETURN or EXIT (RETURN for more solutions)
    (EXIT for only one solution)
  V := PickUnassignedVariable()
  Variable[Level] := V
  Assigned[V] := TRUE
  for d := each member of Domain(V)
    Value[V] := d
    OK := TRUE
    for each constraint C such that
      V is a variable of C
      and all other variables of C are assigned.
      if C is not satisfied by the current set of assignments
        OK := FALSE
    if(OK)
      BT(Level+1)
  return
    
```

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

16

Solving CSPs

- The algorithm searches a tree of partial assignments.



Backtracking Search

- Heuristics are used to determine which variable to assign next “PickUnassignedVariable”.
- The choice can vary from branch to branch, e.g.,
 - under the assignment $V1=a$ we might choose to assign $V4$ next, while under $V1=b$ we might choose to assign $V5$ next.
- This “dynamically” chosen variable ordering has a tremendous impact on performance.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

18

Example.


- N-Queens. Place N Queens on an N X N chess board so that no Queen can attack any other Queen.
 - Variables, one per row.
 - Value of V_i is the column the Queen in row i is place.
 - Constraints.
 - $V_i \neq V_j$ for all $i \neq j$ (can put two Queens in same column)
 - $|V_i - V_j| \neq i - j$ (Diagonal constraint)
 - (i.e., the difference in the values assigned to V_i and V_j can't be equal to the difference between i and j .)

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

19

Example.

- 4X4 Queens

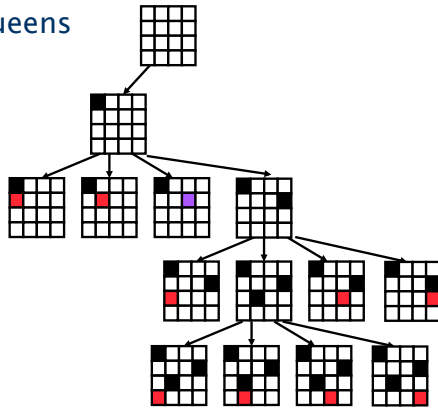


EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

20

Example.

• 4X4 Queens




EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

21

Example.

• 4X4 Queens




EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

22

Backtracking Search

- **Unary** Constraints (over one variable)
 - e.g. $C(X): X=2$ $C(Y): Y>5$
- **Binary** Constraints (over two variables)
 - e.g. $C(X,Y): X+Y<6$
 - Can be represented by **Constraint Graph**
 - Nodes are variables, arcs show constraints.
 - E.g. 4-Queens:



- **Higher-order** constraints: over 3 or more variables
 - We can convert any constraint into a set of binary constraints (may need some auxiliary variables)

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

23

Problems with plain backtracking.


1	2	3						
					4	5	6	
		7						
		8						
		9						

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

24

Constraint Satisfaction Problems

- Sudoku:
 - The 3,3 cell has no possible value. But in the backtracking search we don't detect this until all variables of a row/column or sub-square constraint are assigned. So we have the following situation



EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

25

Constraint Propagation

- Constraint propagation refers to the technique of “looking ahead” in the search at the as yet unassigned variables.
- Try to detect if any obvious failures have occurred.
- “Obvious” means things we can test/detect efficiently.
- Even if we don't detect an obvious failure we might be able to eliminate some possible part of the future search.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

26

Constraint Propagation

- Propagation has to be applied during search. Potentially at every node of the search tree.
- If propagation is slow, this can slow the search down to the point where using propagation actually slows search down!
- There is always a tradeoff between searching fewer nodes in the search, and having a higher nodes/second processing rate.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

27

Forward Checking

- Forward checking is an extension of backtracking search that employs a “modest” amount of propagation (lookahead).
- When a variable is instantiated we check all constraints that have **only one uninstantiated variable** remaining.
- For that uninstantiated variable, we check all of its values, pruning those values that violate the constraint.

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

28

Forward Checking

```

FCCheck(C, x)
  // C is a constraint with all
  // its variables already
  // assigned, except for variable x.
  for d := each member of CurDom[x]
    if making x = d together with
    previous assignments to
    variables in scope C falsifies C
    then
      remove d from CurDom[V]
  if CurDom[V] = {} then return DWO (Domain Wipe Out)
  return ok
  
```

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

29

Forward Checking

```

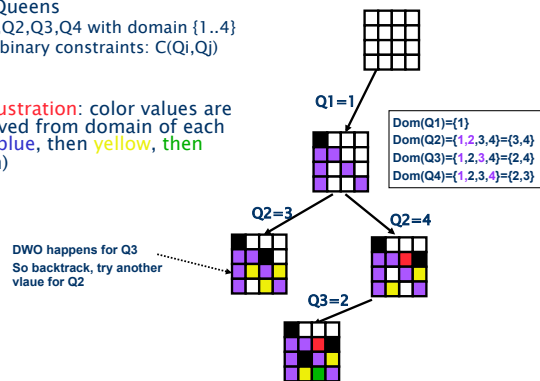
FC(Level) (Forward Checking)
  If all variables are assigned
  PRINT value of each variable
  RETURN or EXIT (RETURN for more solutions)
  (EXIT for only one solution)
  v := PickAnUnassignedVariable()
  Variable[Level] := v
  Assigned[v] := TRUE
  for d := each member of CurDom(v)
    value[v] := d
    for each constraint c over v that has one
    unassigned variable in its scope X.
      val := Fccheck(C,x)
      if(val != DWO)
        FC(Level+1)
  RestoreAllValuesPrunedByFccheck()
  return;
  
```

EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

30

FC Example.

- 4X4 Queens
 - Q1, Q2, Q3, Q4 with domain {1..4}
 - All binary constraints: $C(Q_i, Q_j)$
- FC illustration: color values are removed from domain of each row (blue, then yellow, then green)




EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

31

Example.

- 4X4 Queens continue...



EECS 3401 Fall 2017 Fahiem Bacchus & Yves Lesperance

32

Restoring Values

- After we backtrack from the current assignment (in the for loop) we must restore the values that were pruned as a result of that assignment.
- Some bookkeeping needs to be done, as we must remember which values were pruned by which assignment (FCCheck is called at every recursive invocation of FC).

Minimum Remaining Values

- FC also gives us for free a very powerful heuristic
 - Always branch on a variable with the smallest remaining values (smallest CurDom).
 - If a variable has only one value left, that value is forced, so we should propagate its consequences immediately.
 - This heuristic tends to produce skinny trees at the top. This means that more variables can be instantiated with fewer nodes searched, and thus more constraint propagation/DWO failures occur with less work.

Empirically

- FC often is about 100 times faster than BT
- FC with MRV (minimum remaining values) often 10000 times faster.
- But on some problems the speed up can be much greater
 - Converts problems that are not solvable to problems that are solvable.


Arc Consistency (2-consistency)

- Another form of propagation is to make each arc **consistent**.
- $C(X,Y)$ is consistent iff for **every** value of X there is **some** value of Y that satisfies C .
- Can remove values from the domain of variables:
 - E.G. $C(X,Y): X > Y$ $\text{Dom}(X) = \{1, 5, 11\}$ $\text{Dom}(Y) = \{3, 8, 15\}$
 - For $X=1$ there is no value of Y s.t. $1 > Y$ => remove 1 from domain X
 - For $Y=15$ there is no value of X s.t. $X > 15$, so remove 15 from domain Y
 - We obtain $\text{Dom}(X) = \{5, 11\}$ and $\text{Dom}(Y) = \{3, 8\}$.
- Removing a value from a domain may trigger further inconsistency, so we have to repeat the procedure until everything is consistent.
 - For efficient implementation, we keep track of inconsistent arcs by putting them in a Queue (See AC3 algorithm in the book).
- This is stronger than forward checking. Why?

Backjumping

- Standard backtracking backtracks to the most recent variable (1 level up).
- Trying different values for this variable may have no effect:
 - E.g. $C(X,Y,Z): X \neq Y \ \& \ Z > 3$ and $C(W): W \bmod 2 = 0$
 - $\text{Dom}(X)=\text{Dom}(Y)=\{1..5\}$, $\text{Dom}(Z)=\{3,4,5\}$ $\text{Dom}(W)=\{10..99\}$

After assigning $X=1, Y=1$, and $W=10$, every value of Z fails. So we backtrack to W . But trying different values of W is useless, X and Y are sources of failure!



We should backtrack to Y !

- More intelligent: **Simple Backjumping** backtracks to the last variable among the set of variables that caused the failure, called the conflict set. Conflict set of variable V is the set of previously assigned variables that share a constraint with V . Can be shown that FC is stronger than simple backjumping.
- Even a more efficient approach: **Conflict-Directed-Backjumping**: a more complex notion of conflict set is used: When we backjump to Y from Z , we update the conflict set of Y : $\text{conf}(Y)=\text{conf}(Y) \cup \text{Conf}(Z)-\{Z\}$