


Games, Constraint Satisfaction

CSCI 4511/6511

Joe Goldfrank

Announcements

- Homework 1 is due *today* at 11:55 PM
 - No more than three grace days
- Homework 2 is due on 21 February at 11:55 PM
- Autograder



Algorithms for Games

Adversity

So far:

- The world does not care about us
- This is a simplifying assumption!

Reality:

- The world does not care about us
- ...but it wants things for “itself”
- ...and we don't want the same things

The Adversary

One extreme:

- Single adversary
 - Adversary wants the *exact opposite* from us
 - If adversary “wins,” we lose



Other extreme:

- An entire world of agents with different values
 - They might want some things similar to us
- “Economics”



Simple Games

- Two-player
- Turn-taking
- Discrete-state
- Fully-observable
- Zero-sum
 - This does some work for us!

We Played A Game

- Pick a partner
- Place 11 pieces of candy between you
- Alternating turns, either:
 - Take one piece
 - Take two pieces
- Last person to take a piece wins all of the candy

Max and Min

- Two players want the opposite of each other
- State takes into account both agents
 - Actions depend on whose turn it is

Minimax

- Initial state s_0
- $\text{ACTIONS}(s)$ and $\text{TO-MOVE}(s)$
- $\text{RESULT}(s, a)$
- $\text{IS-TERMINAL}(s)$
- $\text{UTILITY}(s, p)$

Minimax

Minimax

Algorithm Minimax Search

```
1: function MINIMAX-SEARCH(game, state)
2:   player ← game.TO-MOVE(state)
3:   value, move ← MAX-VALUE(game, state)
4:   return move
5:
6: function MAX-VALUE(game, state)
7:   if game.IS-TERMINAL(state) then
8:     return game.UTILITY(state, player), null
9:   v ←  $-\infty$ 
10:  for each a in game.ACTIONS(state) do
11:    v2, a2 ← MIN-VALUE(game, game.RESULT(state, a))
12:    if v2 > v then
13:      v, move ← v2, a
14:  return v, move
15:
16: function MIN-VALUE(game, state)
17:  if game.IS-TERMINAL(state) then
18:    return game.UTILITY(state, player), null
19:  v ←  $\infty$ 
20:  for each a in game.ACTIONS(state) do
21:    v2, a2 ← MAX-VALUE(game, game.RESULT(state, a))
22:    if v2 < v then
23:      v, move ← v2, a
24:  return v, move
```

More Than Two Players

- Two players, two values: v_A, v_B
 - Zero-sum: $v_A = -v_B$
 - Only one value needs to be explicitly represented
- > 2 players:
 - $v_A, v_B, v_C \dots$
 - Value scalar becomes \vec{v}

Society

- > 2 players, only one can win
- Cooperation can be rational!

Example:

- A & B: 30% win probability each
- C: 40% win probability
- A & B cooperate to eliminate C
 - \rightarrow A & B: 50% win probability each

...what about friendship?

Minimax Efficiency

Pruning removes the need to explore the full tree.

- Max and Min nodes alternate
- Once *one* value has been found, we can eliminate parts of search
 - Lower values, for Max
 - Higher values, for Min
- Remember highest value (α) for Max
- Remember lowest value (β) for Min

Pruning

Heuristics 🤔

- In practice, trees are far too deep to completely search
- Heuristic: replace utility with evaluation function
 - Better than losing, worse than winning
 - Represents chance of winning
- Chance? 🎲 🎲
 - Even in deterministic games
 - Why?

More Pruning

- Don't bother further searching bad moves
 - Examples?
- Beam search
 - Lee Sedol's singular win against AlphaGo

Other Techniques

- Move ordering
 - How do we decide?
- Lookup tables
 - For subsets of games

Monte Carlo Tree Search

- Many games are too large even for an efficient α - β search 😞
 - We can still play them
- *Simulate* plays of entire games from starting state
 - Update win probability from each node (for each player) based on result
- “Explore/exploit” paradigm for move selection

Choosing Moves

- We want our search to pick good moves
- We want our search to pick unknown moves
- We *don't* want our search to pick bad moves
 - (Assuming they're actually bad moves)

Select moves based on a heuristic.

Games of Luck

- Real-world problems are rarely deterministic
- Non-deterministic state evolution:
 - Roll a die to determine next position
 - Toss a coin to determine who picks candy first
 - Precise trajectory of kicked football¹
 - Others?

1. Any definition of “football”

Solving Non-Deterministic Games

Previously: Max and Min alternate turns

Now:

- Max
- Chance
- Min
- Chance



We Played Another Game

- Place 11 pieces of candy between you
- Alternating turns:
 - Choose to take one or two pieces
- *Except:*
 - After choosing, flip two coins, record *total* number of heads¹
 - If total is divisible by 3, take one less piece than you chose
 - If total is divisible by 5, take one more piece than you chose
 - If total divisible by 15, take no candy
- Last person to take a piece wins all of the candy

¹ Keep a running total through the game

Expectiminimax

- “Expected value” of next position 
- How does this impact branching factor of the search?



Filled With Uncertainty

What is to be done?

- Pruning is still possible
 - How?
- Heuristic evaluation functions
 - Choose carefully!

Non-Optimal Adversaries

- Is deterministic “best” behavior optimal?
- Are all adversaries rational?

- Expectimax

CSPs

Factored Representation

- Encode relationships between variables and states
- Solve problems with *general* search algorithms
 - Heuristics do not require expert knowledge of problem
 - Encoding problem requires expert knowledge of problem¹

Why?

1. But it always does.

Constraint Satisfaction

- Express problem in terms of state variables
 - Constrain state variables
- Begin with all variables unassigned
- Progressively assign values to variables
- Assignment of values to state variables that “works:” *solution*

More Formally

- State variables: X_1, X_2, \dots, X_n
- State variable domains: D_1, D_2, \dots, D_n
 - The domain specifies which values are permitted for the state variable
 - Domain: set of allowable variables (or permissible range for continuous variables)¹
 - Some constraints C_1, C_2, \dots, C_m restrict allowable values

1. Or a hybrid. such as a union of ranges of continuous variables.

Constraint Types

- Unary: restrict single variable
 - Can be rolled into domain
 - Why even have them?
- Binary: restricts two variables
- Global: restrict “all” variables


Constraint Examples

- X_1 and X_2 both have real domains, i.e. $X_1, X_2 \in \mathbb{R}$
 - A constraint could be $X_1 < X_2$
- X_1 could have domain {red, green, blue} and X_2 could have domain {green, blue, orange}
 - A constraint could be $X_1 \neq X_2$
- $X_1, X_2, \dots, X_n \in \mathbb{R}$
 - Constraint: exactly four of X_i equal 12
 - Rewrite as binary constraint?

Assignments

- Assignments must be to values in each variable's domain
- Assignment violates constraints?
 - Consistency
- All variables assigned?
 - Complete


Yugoslavia¹




1. One of the most difficult problems of the 20th century

Four-Colorings

Two possibilities:



Formulate as CSP?




Graph Representations

- Constraint graph:
 - Nodes are variables
 - Edges are constraints
- Constraint hypergraph:
 - Variables are nodes
 - Constraints are nodes
 - Edges show relationship

Why have two different representations?


Graph Representation I

Constraint graph: edges are constraints



Graph Representation II

Constraint hypergraph: constraints are nodes



How To Solve It

- We can search!
 - ...the space of consistent assignments
- Complexity $O(d^n)$
 - Domain size d , number of nodes n
- Tree search for node assignment
 - Inference to reduce domain size
- Recursive search

How To Solve It

Algorithm Backtracking Search

```
1: function BACKTRACKING-SEARCH( $CSP$ )
2:   return BACKTRACK( $CSP$ , {})
3:
4: function BACKTRACK( $CSP$ ,  $assignment$ )
5:   if  $assignment$  is complete then
6:     return  $assignment$ 
7:    $var \leftarrow$  SELECT-UNASSIGNED-VARIABLE( $CSP$ ,  $assignment$ )
8:   for each  $value$  in ORDER-DOMAIN-VARIABLES( $CSP$ ,  $var$ ,  $assignment$ ) do
9:     if  $value$  is consistent with  $assignment$  then
10:       $assignment.ADD(var = value)$ 
11:       $inferences \leftarrow$  INFERENCE( $CSP$ ,  $var$ ,  $assignment$ )
12:      if  $inferences \neq failure$  then
13:         $CSP.ADD(inferences)$ 
14:         $result \leftarrow$  BACKTRACK( $CSP$ ,  $assignment$ )
15:        if  $result \neq failure$  then
16:          return  $result$ 
17:         $CSP.REMOVE(inferences)$ 
18:       $assignment.REMOVE(var = value)$ 
```

What Even Is Inference

- Constraints on one variable restrict others:
 - $X_1 \in \{A, B, C, D\}$ and $X_2 \in \{A\}$
 - $X_1 \neq X_2$
 - Inference: $X_1 \in \{B, C, D\}$
- If an unassigned variable has no domain...
 - Failure

Inference

- Arc consistency
 - Reduce domains for pairs of variables
- Path consistency
 - Assignment to two variables
 - Reduce domain of third variable

AC-3

Algorithm AC-3

```
1: function AC-3(CSP)
2:   queue  $\leftarrow$  all arcs in CSP
3:   while queue is not empty do
4:      $(X_i, X_j) \leftarrow \text{POP}(\textit{queue})$ 
5:     if REVERSE(CSP,  $X_i, X_j$ ) then
6:       for each  $X_k$  in  $X_i.\text{NEIGHBORS} - \{X_j\}$  do
7:         queue.ADD( $(X_i, X_j)$ )
8:   return True
9:
10: function REVERSE(CSP,  $X_i, X_j$ )
11:   revised  $\leftarrow$  False
12:   for each  $x$  in  $D_i$  do
13:     if  $\mathcal{C}(X_i = x, X_j)$  not satisfied for any value in  $D_j$  then
14:        $D_i.\text{REMOVE}(x)$ 
15:       revised  $\leftarrow$  True
16:   return revised
```

How To Solve It (Again)

Backtracking search:


- Similar to DFS
- Variables are *ordered*
 - Why?
- Constraints checked each step
- Constraints optionally *propagated*

How To Solve It (Again)

Algorithm Backtracking Search

```
1: function BACKTRACKING-SEARCH( $CSP$ )
2:   return BACKTRACK( $CSP, \{\}$ )
3:
4: function BACKTRACK( $CSP, assignment$ )
5:   if  $assignment$  is complete then
6:     return  $assignment$ 
7:    $var \leftarrow$  SELECT-UNASSIGNED-VARIABLE( $CSP, assignment$ )
8:   for each  $value$  in ORDER-DOMAIN-VARIABLES( $CSP, var, assignment$ ) do
9:     if  $value$  is consistent with  $assignment$  then
10:       $assignment.ADD(var = value)$ 
11:       $inferences \leftarrow$  INFERENCE( $CSP, var, assignment$ )
12:      if  $inferences \neq failure$  then
13:         $CSP.ADD(inferences)$ 
14:         $result \leftarrow$  BACKTRACK( $CSP, assignment$ )
15:        if  $result \neq failure$  then
16:          return  $result$ 
17:         $CSP.REMOVE(inferences)$ 
18:       $assignment.REMOVE(var = value)$ 
```

Yugoslav Arc Consistency



Ordering

- SELECT-UNASSIGNED-VARIABLE(CSP , $assignment$)
 - Choose most-constrained variable¹
- ORDER-DOMAIN-VARIABLES(CSP , var , $assignment$)
 - Least-constraining value
- Why?


1. or MRV: “Minimum Remaining Values”

Restructuring

Tree-structured CSPs:

- *Linear time* solution
- Directional arc consistency: $X_i \rightarrow X_{i+1}$
- Cutsets
- Sub-problems

Cutset Example



(Heuristic) Local Search

- Hill climbing
 - Random restarts
- Simulated annealing
- Fast?
- Complete?
- Optimal?


Continuous Domains

- Linear:

$$\begin{aligned} \max_x \quad & \mathbf{c}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq 0 \end{aligned}$$

- Convex

$$\begin{aligned} \min_x \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & g_i(\mathbf{x}) \leq 0 \\ & h_i(\mathbf{x}) = 0 \end{aligned}$$



Is This Even Relevant in 2025?

- Absolutely yes.
- LLMs are bad at CSPs
- CSPs are common in the real world
 - Scheduling
 - Optimization
 - Dependency solvers

Logic Preview

$$R_{HK} \Rightarrow \neg R_{SI}$$

$$G_{HK} \Rightarrow \neg G_{SI}$$

$$B_{HK} \Rightarrow \neg B_{SI}$$

$$R_{HK} \vee G_{HK} \vee B_{HK}$$

...

Goal: find assignment of variables that satisfies conditions

References

- Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th Edition, 2020.
- Mykal Kochenderfer, Tim Wheeler, and Kyle Wray. *Algorithms for Decision Making*. 1st Edition, 2022.
- Stanford CS231
- Stanford CS228
- UC Berkeley CS188