

CSE 473: Artificial Intelligence  
Autumn 2014

**Heuristics & Pattern  
Databases for Search**

Dan Weld

With many slides from  
Dan Klein, Richard Korf, Stuart Russell, Andrew Moore, & UW Faculty

## Logistics

- PS1 – due Monday 10/13
- Office hours
  - Jeff today 10:30am CSE 021
  - Galen today 1-3pm CSE 218
  - See Website for all

3

## Recap: Search Problem

---

- **States**
  - configurations of the world
- **Successor function:**
  - function from states to lists of (state, action, cost) triples
- **Start state**
- **Goal test**

## N-Queens as Search?

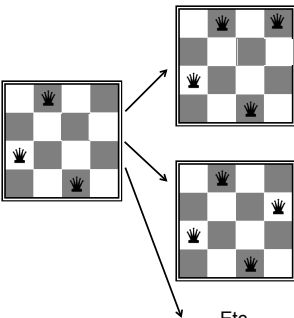


- Given N x N chess board
- Can you place N queens so they don't fight?



Cool picture from Dan Klein & Pieter Abeel ai.berkeley.edu 5

## States are Board Positions



Etc...

6

## Search Methods

- Depth first search (DFS)
- Breadth first search (BFS)
- Iterative deepening depth-first search (IDS)

- Best first search
- Uniform cost search (UCS)
- Greedy search
- A\*
- Iterative Deepening A\* (IDA\*)
- Beam search, hill climbing

Heuristic search

- **Stochastic Search**
- **Constraint Satisfaction**

7

### IDA\* for N-Queens?

- Given  $N \times N$  chess board
- Can you place  $N$  queens so they don't fight?

Cool picture from Dan Klein & Pieter Abeel at berkeley.edu 8

### Best-First Search

- Generalization of breadth-first search
- Fringe = **Priority** queue of nodes to be explored
- Cost function  $f(n)$  applied to each node

Add initial state to priority queue  
 While queue not empty  
     Node = head(queue)  
     If goal?(node) then return node  
     Add children of node to queue

← "expanding the node"  
9

### Which Algorithm?

---

- Uniform cost search (UCS):

10

### Which Algorithm?

---

- A\*, Manhattan Heuristic:

### Which Algorithm?

---

- Best First / Greedy, Manhattan Heuristic:

### Iterative-Deepening A\*

- Like iterative-deepening depth-first, but...
- Depth bound modified to be an **f-limit**
  - Start with  $f\text{-limit} = h(\text{start})$
  - Prune any node if  $f(\text{node}) > f\text{-limit}$
  - Next  $f\text{-limit} = \text{min-cost of any node pruned}$

13

### IDA\* Analysis

- Complete & Optimal (ala A\*)
- Space usage  $\propto$  depth of solution
- Each iteration is DFS - no priority queue!
- # nodes expanded relative to A\*
  - Depends on # unique values of heuristic function
  - In 8 puzzle: few values  $\Rightarrow$  close to # A\* expands
  - In eastern-europe travel: each f value is unique  $\Rightarrow 1+2+\dots+n = O(n^2)$  where  $n$ =nodes A\* expands if  $n$  is too big for main memory,  $n^2$  is too long to wait!
- Generates duplicate nodes in cyclic graphs

14

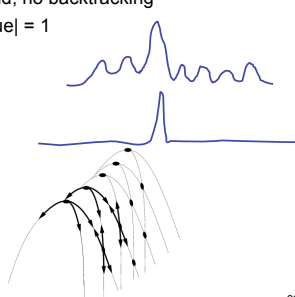
### Beam Search

- Idea
  - Best first
  - But discard all but N best items on priority queue
- Evaluation
  - Complete? No
  - Time Complexity?  $O(b^d)$
  - Space Complexity?  $O(b + N)$

© Daniel S. Weld 19

### Hill Climbing "Gradient ascent"

- Idea
  - Always choose best child; no backtracking
  - Beam search with |queue| = 1
- Problems?
  - Local maxima
  - Plateaus
  - Diagonal ridges



© Daniel S. Weld 20

# Heuristics

It's what makes search actually work

### Admissable Heuristics

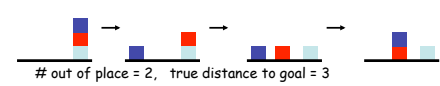
- $f(x) = g(x) + h(x)$
- $g$ : cost so far
- $h$ : underestimate of remaining costs

Where do heuristics come from?

© Daniel S. Weld 24

### Relaxed Problems

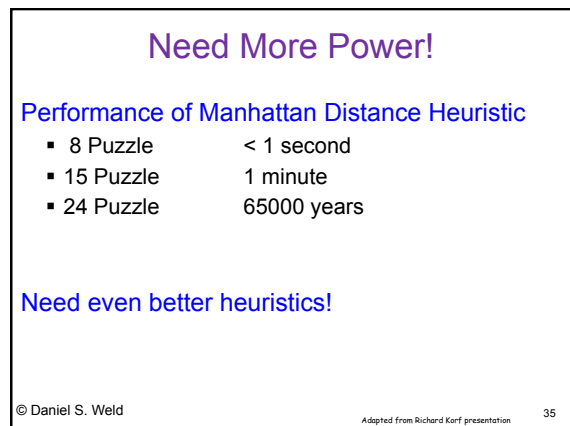

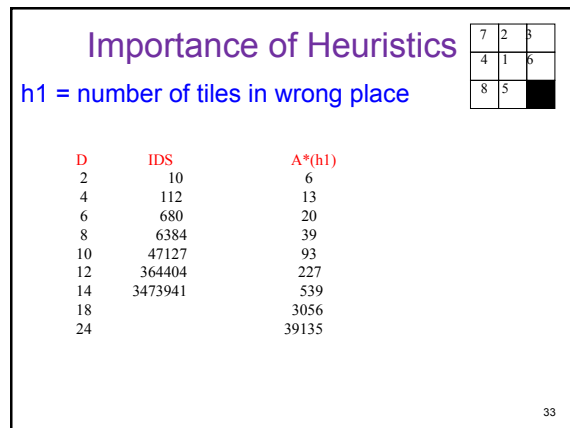

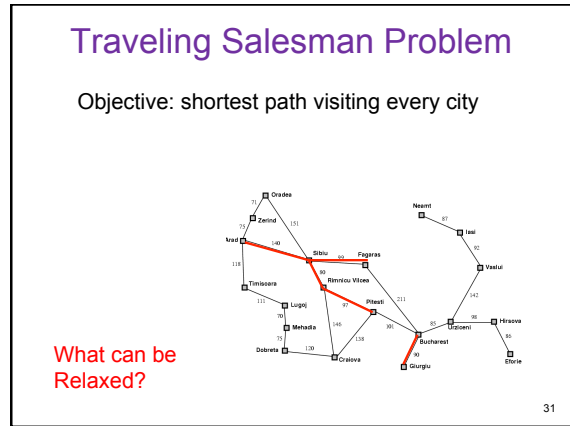

- Derive admissible heuristic from **exact cost of a solution to a relaxed version of problem**
  - For blocks world, distance = # move operations
  - heuristic = number of misplaced blocks
  - What is relaxed problem?



# out of place = 2, true distance to goal = 3

- Cost of optimal soln to relaxed problem  $\leq$  cost of optimal soln for real problem

© Daniel S. Weld 25



## Subgoal Interactions

- **Manhattan distance assumes**
  - Each tile can be moved independently of others
- **Underestimates because**
  - Doesn't consider interactions between tiles

1	2	3
4	6	5
7	8	

© Daniel S. Weld Adapted from Richard Korf presentation 36

## Pattern Databases

[Culberson & Schaeffer 1996]

- **Pick any subset of tiles**
  - E.g., 3, 7, 11, 12, 13, 14, 15
  - (or as drawn)
- **Precompute a table**
  - Optimal cost of solving just these tiles
  - For all possible configurations
    - 57 Million in this case
  - Use A\* or IDA\*
    - State = **position of just these tiles (& blank)**

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	

© Daniel S. Weld Adapted from Richard Korf presentation 37

## Using a Pattern Database

- **As each state is generated**
  - Use position of chosen tiles as index into DB
  - Use lookup value as heuristic,  $h(n)$
- Admissible?

© Daniel S. Weld Adapted from Richard Korf presentation 38

## Combining Multiple Databases

- **Can choose another set of tiles**
  - Precompute multiple tables
- **How combine table values?**
- **E.g. Optimal solutions to Rubik's cube**
  - First found w/ IDA\* using pattern DB heuristics
  - Multiple DBs were used (dif cubie subsets )
  - Most problems solved optimally in 1 day
  - Compare with *574,000 years* for IDDFS

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	

© Daniel S. Weld Adapted from Richard Korf presentation 39

## Drawbacks of Standard Pattern DBs

- **Since we can only take *max***
  - Diminishing returns on additional DBs
- **Would like to be able to *add* values**

© Daniel S. Weld Adapted from Richard Korf presentation 40

## Disjoint Pattern DBs

- **Partition tiles into disjoint sets**
  - For each set, precompute table
    - E.g. 8 tile DB has 519 million entries
    - And 7 tile DB has 58 million
- **During search**
  - Look up heuristic values for each set
  - **Can add values without overestimating!**
- Manhattan distance is a special case of this idea where each set is a single tile

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	

© Daniel S. Weld Adapted from Richard Korf presentation 41

## Performance

- **15 Puzzle: 2000x speedup vs Manhattan dist**
  - IDA\* with the two DBs shown previously solves 15 Puzzles optimally in 30 milliseconds
- **24 Puzzle: 12 million x speedup vs Manhattan**
  - IDA\* can solve random instances in 2 days.
  - Requires 4 DBs as shown
    - Each DB has 128 million entries
  - Without PDBs: 65,000 years

