

Heuristics Local Search

CSE 473
University of Washington

Admissible Heuristics

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

Where do heuristics come from?

© D. Weld, D. Fox

2

Relaxed Problems

- Derive admissible heuristic from **exact** cost of a solution to a **relaxed** version of problem

For transportation planning, relax requirement that car has to stay on road \rightarrow Euclidean dist

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

© D. Weld, D. Fox

3

Simplifying Integrals

vertex = formula

goal = closed form formula without integrals

arcs = mathematical transformations

$$\int \dots^{n+}$$

heuristic = number of integrals still in formula

what is being relaxed?

© D. Weld, D. Fox

4

Traveling Salesman Problem

- Problem Space

States = partial path (not nec. connected)
 Operator = add an edge
 Start state = empty path
 Goal = complete path

- Heuristic?




What can be Relaxed?

© D. Weld, D. Fox

5

Heuristics for eight puzzle



start

goal

- What can we relax?

© D. Weld, D. Fox

6

Importance of Heuristics

- h_1 = number of tiles in wrong place
- h_2 = distances of tiles from correct loc



D	IDS	A*(h1)	A*(h2)
2	10	6	6
4	112	13	12
6	680	20	18
8	6384	39	25
10	47127	93	39
12	364404	227	73
14	3473941	539	113
18		3056	363
24		39135	1641

© D. Weld, D. Fox

7

Need More Power!

Performance of Manhattan Distance Heuristic

8 Puzzle	< 1 second
15 Puzzle	1 minute
24 Puzzle	65000 years

Need even better heuristics!

© D. Weld, D. Fox

Adapted from Richard Korf presentation

8

Pattern Databases

[Culberson & Schaeffer 1996]

- Pick any subset of tiles
 - E.g., 3, 7, 11, 12, 13, 14, 15
- Precompute a table
 - Optimal cost of solving just these tiles
 - For all possible configurations
 - 57 Million in this case
 - Use breadth first search back from goal state
 - State = position of just these tiles (& blank)

© D. Weld, D. Fox

Adapted from Richard Korf presentation 10

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?

© D. Weld, D. Fox

Adapted from Richard Korf presentation 11

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 subsets of cubies)
 - Most problems solved optimally in 1 day
 - Compare with **574,000 years** for IDDFS

© D. Weld, D. Fox

Adapted from Richard Korf presentation 12

Drawbacks of Standard Pattern DBs

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

© D. Weld, D. Fox

Adapted from Richard Korf presentation 13

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

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	■

- 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

© D. Weld, D. Fox

Adapted from Richard Korf presentation 14

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: 65000 years



© D. Weld, D. Fox

Adapted from Richard Korf presentation 15

Local Search Algorithms

- In many optimization problems, the path to the goal is irrelevant; the goal state itself is the solution

- State space = set of "complete" configurations

- Find configuration satisfying constraints, e.g., n-queens

- In such cases, we can use local search algorithms that keep a single "current" state, try to improve it

© D. Weld, D. Fox

16

Hill Climbing

"Gradient ascent"

- Idea


Always choose best child; no backtracking
Beam search with |queue| = 1

- Problems?

Local maxima

Plateaus

Diagonal ridges



© D. Weld, D. Fox

17

Stochastic Hill Climbing

- Randomly disobeying heuristic
- Random restarts

© D. Weld, D. Fox

18

Simulated Annealing

- Objective: avoid local minima
- Technique:
 - For the most part use hill climbing
 - When no improvement possible
 - Choose random neighbor
 - Let a be the decrease in quality
 - Move to neighbor with probability $e^{-a/T}$
 - Reduce "temperature" (T) over time
- Pros & cons
 - Optimal?
 - If T decreased slowly enough, *will* reach optimal state
- Widely used
- See also
 - WalkSAT



© D. Weld, D. Fox

19

Local Beam Search

- Idea
 - Best first but only keep N best items on priority queue
- Evaluation
 - Complete?
 - Time Complexity?
 - Space Complexity?

© D. Weld, D. Fox

21

Genetic Algorithms

- Start with random population
 - Representation serialized
 - States are ranked with "fitness function"
- Produce new generation
 - Select random pair(s):
 - probability \sim fitness
 - Randomly choose "crossover point"
 - Offspring mix halves
 - Randomly mutate bits

© D. Weld, D. Fox

22

Genetic algorithms



• Fitness function: number of non-attacking pairs of queens (min = 0, max = $8 \times 7/2 = 28$)


• $24/(24+23+20+11) = 31\%$

• $23/(24+23+20+11) = 29\%$ etc

© D. Weld, D. Fox

23

Genetic algorithms



© D. Weld, D. Fox

24