

CSE 473: Artificial Intelligence Autumn 2018

Heuristics & Pattern Databases for Search

Steve Tanimoto

Presented by Emilia Gan

With thanks to Dan Weld, Dan Klein, Richard Korf, Stuart Russell, Andrew Moore, and Luke Zettlemoyer


*With modifications from various sources (as noted on individual slides) E. Gan Fall '18

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



States are Board Positions



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

Heuristic search

5

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 Abbeel ai.berkeley.edu 6

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"

Greedy Best First Algorithm

- Recall: BFS and DFS pick the next node off the frontier based on which was "first in" or "last in".
- Greedy Best First picks the "best" node according to some rule of thumb, called a *heuristic*.

Definition: A *heuristic* is an approximate measure of how close you are to the target.

A heuristic guides you in the right direction.

<https://cs.stanford.edu/people/abisee/gs.pdf>

A*

- Expands the path with the **lowest cost + h** value on the frontier
- The frontier is implemented as a **priority queue** ordered by $f(p) = \text{cost}(p) + h(p)$

Admissibility of a heuristic

Def.: Let $c(n)$ denote the cost of the optimal path from node n to any goal node. A search heuristic $h(n)$ is called **admissible** if $h(n) \leq c(n)$ for all nodes n , i.e. if for all nodes it is an **underestimate** of the cost to any goal.

<https://www.cs.ubc.ca/~mack/CS322/lectures/2-Search6.pdf>

The main drawback of the presented best-first graph search algorithms is their space complexity.


Idea: use the concepts of iterative-deepening DFS

- bounded depth-first search with increasing bounds
- instead of **depth** we bound **f**
 (in this chapter $f(n) := g(n) + h(n.\text{state})$ as in A*)
- **IDA*** (iterative-deepening A*)
- tree search**, unlike the previous best-first search algorithms

https://ai.dmi.unibas.ch/_files/teaching/fs17/ai/slides/ai17.pdf

Iterative Deepening DFS (IDS) in a Nutshell

- Use DFS to look for solutions at depth 1, then 2, then 3, etc
 - For depth D , ignore any paths with longer length
 - Depth-bounded depth-first search



<https://www.cs.ubc.ca/~mack/CS322/lectures/2-Search6.pdf>


(Heuristic) Iterative Deepening: IDA*

- Like Iterative Deepening DFS
 - But the depth bound is measured in terms of the f value
- If you don't find a solution at a given depth
 - Increase the depth bound:
 - to the minimum of the f -values that exceeded the previous bound

<https://www.cs.ubc.ca/~mack/CS322/lectures/2-Search6.pdf>


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 $f\text{-limit} = \text{min-cost of any node pruned}$



IDA*(Iterative Deepening A*) Search

- Perform depth-first search LIMITED to some $f\text{-bound}$.
- If goal found: ok.
- Else: increase $f\text{-bound}$ and restart.
- How to establish the $f\text{-bounds}$?
 - initially: $f(S)$
 - generate all successors
 - record the minimal $f(\text{succ}) > f(S)$
- Continue with minimal $f(\text{succ})$ instead of $f(S)$



<https://www.slideshare.net/hemak15/lecture-17-iterative-deepening-a-star-algorithm>

```

path      current search path (acts like a stack)
node      current node (last node in current path)
g          the cost to reach current node
f          estimated cost of the cheapest path (root..node..goal)
h(node)   estimated cost of the cheapest path (node..goal)
cost(node, succ) step cost function
is_goal(node) goal test
successors(node) node expanding function, expand nodes ordered by g + h(node)
ida_star(root) return either NOT_FOUND or a pair with the best path and its cost

procedure ida_star(root)
  bound := h(root)
  path := [root]
  loop
    t := search(path, 0, bound)
    if t = FOUND then return (path, bound)
    if t = ∞ then return NOT_FOUND
    bound := t
  end loop
end procedure

```

https://en.wikipedia.org/wiki/Iterative_deepening_A*


```

function search(path, g, bound)
  node := path.last
  f := g + h(node)
  if f > bound then return f
  if is_goal(node) then return FOUND
  min := ∞
  for succ in successors(node) do
    if succ not in path then
      path.push(succ)
      t := search(path, g + cost(node, succ), bound)
      if t = FOUND then return FOUND
      if t < min then min := t
      path.pop()
    end if
  end for
  return min
end function

```


https://en.wikipedia.org/wiki/Iterative_deepening_A*

f-limited, f-bound = 100

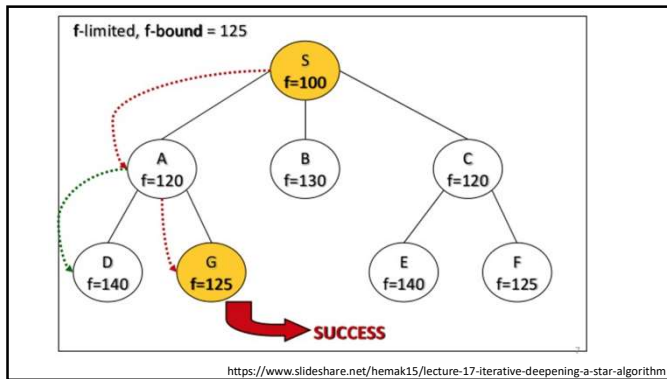


<https://www.slideshare.net/hemak15/lecture-17-iterative-deepening-a-star-algorithm>

f-limited, f-bound = 120



<https://www.slideshare.net/hemak15/lecture-17-iterative-deepening-a-star-algorithm>



IDA* Analysis

- Complete & Optimal (a la 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

- IDA* is a tree search variant of A* based on iterative deepening depth-first search
- main advantage: **low space complexity**
- disadvantage: **repeated work** can be significant
- most useful when there are **few duplicates**

https://ai.dmi.unibas.ch/_files/teaching/fs17/ai/slides/ai17.pdf

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?
 - $O(b + N)$

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


15

HILL-CLIMBING CAN GET STUCK!

Diagonal ridges:

From each local maximum all the *available* actions point downhill, but there is an uphill path!

Zig-zag motion, very long ascent time!



Gradient ascent doesn't have this issue: *all* state vector components are (potentially) changed when moving to a successor state, climbing can follow the direction of the ridge

http://www.cs.cmu.edu/~ariel/proj25/83716/c_slides/78316-2a.pdf

Importance of Heuristics

$h1$ = number of tiles in wrong place

| | | |
|---|---|---|
| 7 | 2 | 3 |
| 4 | 1 | 6 |
| 8 | 5 | |

| D | IDS | $A^*(h1)$ |
|----|---------|-----------|
| 2 | 10 | 6 |
| 4 | 112 | 13 |
| 6 | 680 | 20 |
| 8 | 6384 | 39 |
| 10 | 47127 | 93 |
| 12 | 364404 | 227 |
| 14 | 3473941 | 539 |
| 18 | | 3056 |
| 24 | | 39135 |

23

Importance of Heuristics

$h1$ = number of tiles in wrong place

$h2 = \sum$ distances of tiles from correct loc

| | | |
|---|---|---|
| 7 | 2 | 3 |
| 4 | 1 | 6 |
| 8 | 5 | |

| 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 |

Decrease effective branching factor

24

Need More Power!

Performance of Manhattan Distance Heuristic

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

Need even better heuristics!

© Daniel S. Weld

Adapted from Richard Korf presentation

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

26



cost of the optimal solution of sub-problem
 \leq cost of the optimal solution of complete problem



Open Education Edinburgh
Published on Feb 26, 2014

<https://www.youtube.com/watch?v=HZWV4uOJWk8>

Pattern Databases

- idea: pre-compute and store the solution costs for all possible sub-problems in database
- computing heuristic = DB lookup
- construct DB by searching backwards from the goal state and recording costs
 - very expensive operation, but needs to be computed only once



Open Education Edinburgh
Published on Feb 26, 2014

<https://www.youtube.com/watch?v=HZWV4uOJWk8>

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

27

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

28

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 cubic 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


29

Efficiency

Time for the preprocessing to create a PDB is usually negligible compared to the time to solve one problem-instance with no heuristic.

Memory is the limiting factor.

Many small pattern databases



<https://webdocs.cs.ualberta.ca/~holte/CMPT651/pdb20041031.pdf>

Rubik's Cube

| PDB Size | n | Nodes Generated |
|-------------|---|-----------------|
| 13,305,600 | 9 | 2,654,669 |
| 17,740,800 | 6 | 2,639,960 |
| 26,611,200 | 4 | 3,096,919 |
| 53,222,400 | 2 | 5,329,829 |
| 106,444,800 | 1 | 61,465,541 |

Summary

| State Space | Best n | Ratio |
|-----------------------|--------|-------------|
| (3x3)-puzzle | 10 | 3.85 |
| 9-pancake | 10 | 6.59 |
| (8,4)-Topspin (3 ops) | 9 | 3.76 |
| (8,4)-Topspin (8 ops) | 9 | 20.89 |
| (3x4)-puzzle | 21+ | 185.5 |
| Rubik's Cube | 6 | 23.26 |
| 15-puzzle (additive) | 5 | 2.38 |
| 24-puzzle (additive) | 8 | 1.6 to 25.1 |

RATIO = $\frac{\text{\#nodes generated using one PDB of size } M}{\text{\#nodes generated using } n \text{ PDBs of size } M/n}$

<https://webdocs.cs.ualberta.ca/~holte/CMPT651/pdb20041031.pdf>

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

30

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

© Daniel S. Weld

Adapted from Richard Korf presentation

31

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

© Daniel S. Weld

Adapted from Richard Karf presentation

32