



Class 4 - Greedy Heuristics and Hill Climbing

Guiding Genetic Algorithms with Language Models for Combinatorial Optimization

MISE Research Program - July 2025




Pedro Paredes

Informed Search

Recall: Finding a Path on a Grid

Suppose you have a grid and you want to find a *shortest* path from point A to point B avoiding obstacles





DFS Solution

```
1  def dfs_shortest_path(r, c, goal, path_len, visited):
2      if (r, c) == goal:
3          return path_len
4
5      # explore four directions
6      bs = 100000
7      for dr, dc in ((1,0), (-1,0), (0,1), (0,-1)):
8          nr, nc = r + dr, c + dc
9          inside = 0 <= nr < len(grid) and 0 <= nc < len(grid[0])
10         if inside and grid[nr][nc] == 0 and (nr, nc) not in visited:
11             visited.add((nr, nc))
12             bs = min(bs, dfs_shortest_path(nr, nc, goal, path_len + 1, visited))
13             visited.remove((nr, nc))
14     return bs
```



Informed Search - Dijkstra's Algorithm

The simple DFS solution is pretty inefficient since we repeatedly search the same paths and we search paths that are not very promising



We can make our search more efficient using *informed search* (also known as *Dijkstra's Algorithm*):

- Instead of going depth-first, let's look at more promising paths first, i.e., shorter
- Instead of keeping track of visited cells, we will keep track of the shortest path to get to each of them. If we visit a cell and the current distance is greater than or equal to the previous best, ignore this path.


To prioritize shorter paths first, we will use a data structure known as a *heap*, which lets us find the minimum element in it in $O(\log n)$ time.



```
1  import heapq
2
3  def dijkstra_shortest_path(start, goal):
4      rows, cols = len(grid), len(grid[0])
5      dist = {start: 0}
6      pq = [(0, start[0], start[1])]
7
8      while pq:
9          d, r, c = heapq.heappop(pq)
10         if (r, c) == goal:
11             return d
12         if d != dist[(r, c)]:
13             continue
14
15         for dr, dc in ((1, 0), (-1, 0), (0, 1), (0, -1)):
16             nr, nc = r + dr, c + dc
17             inside = 0 <= nr < rows and 0 <= nc < cols
18             if inside and grid[nr][nc] == 0:
19                 nd = d + 1
20                 if nd < dist.get((nr, nc), float('inf')):
21                     dist[(nr, nc)] = nd
22                     heapq.heappush(pq, (nd, nr, nc))
23     return None
```



Packing a Truck



How many trucks do we need to carry all of the packages?

Truck size



Informed Search for Packing a Truck

We will use a similar approach as before, but we need to encode the problem:

- A state is a tuple $(i, bins)$, which means we have placed the first i packages and we have a list $bins$ with the remaining capacity of each non-empty bin
- A transition is placing one more item, so we can either place it in any bin that has capacity for it, or we open a new bin
- The most promising states are the ones with less bins

Because the state space can be really large, we don't keep track of the best way to get to each state

```
1  import heapq
2
3  def dijkstra_binpack(weights, C):
4      N = len(weights)
5      start = (0, ())
6      pq = [(0, start)]
7
8      while pq:
9          g, (i, bins) = heapq.heappop(pq)
10         if i == N:
11             return g
12
13         w = weights[i]
14
15         # 1) try every existing bin that fits
16         for idx, free in enumerate(bins):
17             if free >= w:
18                 new_bins = list(bins)
19                 new_free = free - w
20                 if new_free == 0:
21                     new_bins.pop(idx) # bin now full
22                 else:
23                     new_bins[idx] = new_free
24                 new_state = (i + 1, tuple(new_bins))
25                 heapq.heappush(pq, (g, new_state))
26
27         # 2) open a fresh bin for this item
28         if w > C:
29             return None
30         new_bins = bins + (C - w,)
31         new_state = (i + 1, new_bins)
32         ng = g + 1
33         heapq.heappush(pq, (ng, new_state))
34
35     return None
```

should not reach here on valid input

Heuristics



Heuristics

A *heuristic* is a non-optimal approximation to a problem

Often we use intuition about a problem to come up with heuristics

We can use heuristics in several ways:

- To pick which states to explore first
- To discard non promising solutions



A* Algorithm

The A* algorithm is an algorithm that combines heuristics with informed search



A* Algorithm for Bin Packing

[See code demo](#)

Hill Climbing



Hill Climbing

Hill climbing is a method to find a solution to a problem by always moving in the direction of a “better” solution

- Start from a random state s
- Move to a neighbor with better score (break ties randomly)
- If we get stuck and didn't find a solution, restart

A neighbor of a state is a state obtained by a local change: e.g., moving one queen in the n -queens example



Hill Climbing for N-Queens

See code demo for N-queens



What's next?

Next class monday

No assignment sheet, study this week's lectures and solve all remaining problems

Class 5: Introduction to Genetic Algorithms