

ITS336 Lecture 5.


Local Search

5.1. Local Search Algorithms

In many search problems, the path from the initial state to the goal is **not** (4.1) important, but the final configuration of the goal is the target of the problem, e.g., 8-queens problem.

Local search is a group of algorithms that focus on this kind of problems. Since the path is not matter, local search algorithms typically keep only one node while operating.

Apart from searching for goal states, the algorithms can be used to solve **optimization problems**. The optimization problems aim to find the state with the best value of an **objective function**.



5.2. Hill Climbing Search

Hill climbing search algorithm repeatedly moves the current node to the direction that increases (or decreases) the objective value. It ends when the current node reaches a peak where no neighbor has a higher value. (4.1.1)

Code 5.1 Hill Climbing Algorithm implemented in Python

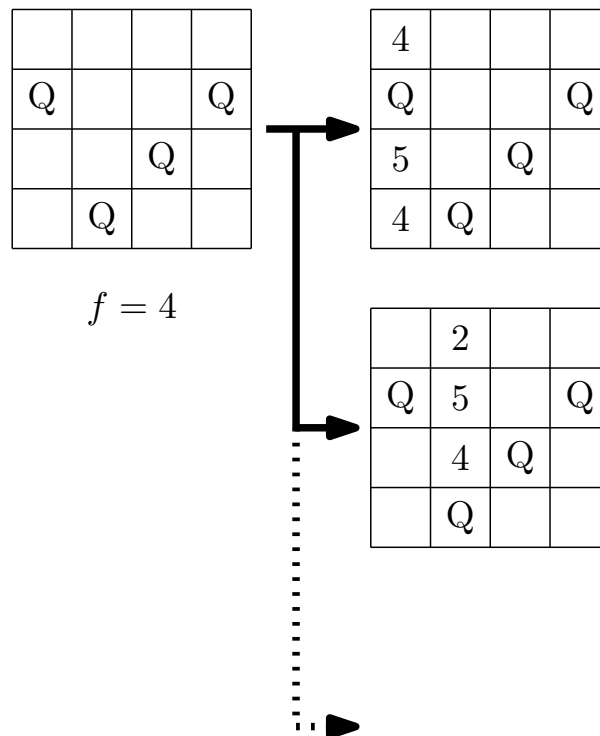
```
1 def hill_climbing(s0, succ, f):
2     # s0 = initial state
3     # succ = successor function
4     # f = evaluation function
5     u = s0
6     while True:
7         s = succ(u)
8         v = max(s, key=f)
9         if f(v) > f(u):
10            u = v
11        else:
12            return u
```

Example 5.1 To solve the 8-queens problem by the Hill Climbing algorithm, we need to define a function that evaluates each state.

We can use the number of attacks among the queens on the board as an evaluation value of a state. Thus, $f(s)$ returns the number of attacks given a state s . The search can start from a randomly created board with one queen placed in each column.

A successor state can be generated by selecting one queen and change its row position.

Here, we search for the state with $f = 0$.



Exercise 5.1 A local search problem represents a state by an ordered pair. The initial state is $(0, 0)$. The successors of a state (x, y) are $(x - 1, y)$, $(x + 1, y)$, $(x, y - 1)$, and $(x, y + 1)$. The objective function is shown in the following figure. Use the hill climbing algorithm to search for the state with maximum objective value.

	0	1	2	3	4	5	6	7	8	9
0	3	4	4	4	5	4	5	5	5	4
1	4	5	5	5	4	4	5	6	5	3
2	5	6	7	6	4	3	6	7	6	2
3	5	6	6	6	4	2	6	6	6	2
4	5	5	5	4	4	1	6	5	5	3
5	3	3	4	3	7	7	7	6	4	4
6	4	3	4	3	7	8	7	8	7	5
7	4	5	5	4	6	7	9	8	7	5
8	4	5	5	4	6	8	8	8	7	6
9	3	3	3	4	6	6	7	7	7	5

5.2.1. Random-restart Hill Climbing

To avoid getting stuck in local maxima, the hill climbing search can be conducted multiple times. In each iteration, it starts from different randomly generated initial states.

Exercise 5.2 Conduct the random-restart hill climbing on the following space.

	0	1	2	3	4	5	6	7	8	9
0	3	4	4	4	5	4	5	5	5	4
1	4	5	5	5	4	4	5	6	5	3
2	5	6	7	6	4	3	6	7	6	2
3	5	6	6	6	4	2	6	6	6	2
4	5	5	5	4	4	1	6	5	5	3
5	3	3	4	3	7	7	7	6	4	4
6	4	3	4	3	7	8	7	8	7	5
7	4	5	5	4	6	7	9	8	7	5
8	4	5	5	4	6	8	8	8	7	6
9	3	3	3	4	6	6	7	7	7	5

Code 5.2 Random Restart Hill Climbing Algorithm

```
1 from hillclimbing import *
2
3
4 def random_restart_hill_climbing(succ, f, rnd, n):
5     # succ = successor function
6     # f     = evaluation function
7     # rnd   = function that randomly generates state
8     # n     = number of iterations
9     best_value = -float("inf")
10    for i in range(n):
11        s0 = rnd()
12        u = hill_climbing(s0, succ, f)
13        if f(u) > best_value:
14            best_state = u
15            best_value = f(u)
16    return best_state
```

5.2.2. Stochastic Hill Climbing Search

The hill climbing search basically goes directly to the closest local maximum with exploring the search space.

To allow the search to explore the space, we use the probability p to judge if a newly generated state should be accepted.

$$p = \frac{1}{1 + e^{\frac{\text{current} - \text{new}}{T}}}$$

where T is the parameter to control how easy the worse state is selected.

Code 5.3 Stochastic Hill Climbing Algorithm

```

1 import random
2 import math
3
4
5 def stochastic_hill_climbing(s0, succ, f, n, t):
6     # succ = successor function
7     # f = objective function
8     # n = number of iterations
9     # t = control parameter
10    best_value = -float("inf") # negative infinity
11    u = s0 # u = current state
12    for i in range(n):
13        s = succ(u) # generate a set of successors
14        v = random.choice(s) # randomly choose one to be the next state
15        p = 1/(1+math.exp((f(u)-f(v))/t)) # calculate probability
16        r = random.uniform(0.0, 1.0) # randomly pick "r"
17        if r < p: # move to "v" if r < p
18            u = v
19        if f(u) > best_value: # keep the best state up to now
20            best_state = u
21            best_value = f(u)
22    return best_state

```

Consider the following values of p based on the different values of T

T	$e^{\frac{-13}{T}}$	p	$e^{\frac{13}{T}}$	p
1	0.00000226	1.000	442413.39	0.00000226
5	0.0498	0.953	13.5	0.0691
10	0.273	0.786	3.67	0.214
20	0.522	0.657	1.92	0.343
50	0.771	0.565	1.30	0.435
10^{10}	1.00	0.500	1.00	0.500

5.3. Simulated Annealing

Simulated annealing is based on an idea from thermodynamics.

- To grow a crystal, we heat the materials to a molten state. Then, cool it down until the crystal structure is frozen in.
- If the cooling is done too quickly, the materials become brittle.


The parameter T in the stochastic hill-climbing can be compared to the temperature. We start with T in a high value to allow random walks, and we decrease T when the search is going on. This allows the exploration of the search space at the beginning, then it goes to the optimal state in the end.

Code 5.4 Simulated Annealing Algorithm


```
1 import random
2 import math
3
4
5 def simulated_annealing(s0, succ, f, t0, cooling):
6     # succ = successor function
7     # f     = evaluation function
8     # rnd   = function that randomly generates state
9     # n     = number of iterations
10    best_value = -float("inf")
11    u = s0
12    t = t0
13    i = 0
14    while t > 0:
15        s = succ(u)
16        v = random.choice(s)
17        if f(u) > f(v):
18            u = v
19        else:
20            p = 1/(1+math.exp((f(u)-f(v))/t))
21            r = random.uniform(0.0, 1.0)
22            if r < p:
23                u = v
24        i += 1
25        t = cooling(t, i)
26        if f(u) > best_value:
27            best_state = u
28            best_value = f(u)
29    return best_state
```

5.3.1. Cooling Schedule

$$T_i = T_0 - i \frac{T_0 - T_N}{N}$$



$$T_i = T_0 \left(\frac{T_N}{T_0} \right)^{\frac{i}{N}}$$



5.4. Local Beam Search

Instead of keeping track of only 1 node, **local beam search** algorithm keeps track of k states while conducting the search. It starts from randomly generated k states. (4.1.3)

The algorithm selects the k best neighbors. It may select all of the neighbors if there are less than k states, and randomly select in the case of ties. Then, these k states are used in the next iteration.

Local beam search with k states is different from conducting k random restarts hill climbing search.

Exercise 5.3 Conduct the local beam search when the beam size is 3, and the initial states are $(0, 0)$; $(1, 8)$; $(7, 4)$

	0	1	2	3	4	5	6	7	8	9
0	3	4	4	4	5	4	5	5	5	4
1	4	5	5	5	4	4	5	6	5	3
2	5	6	7	6	4	3	6	7	6	2
3	5	6	6	6	4	2	6	6	6	2
4	5	5	5	4	4	1	6	5	5	3
5	3	3	4	3	7	7	7	6	4	4
6	4	3	4	3	7	8	7	8	7	5
7	4	5	5	4	6	7	9	8	7	5
8	4	5	5	4	6	8	8	8	7	6
9	3	3	3	4	6	6	7	7	7	5

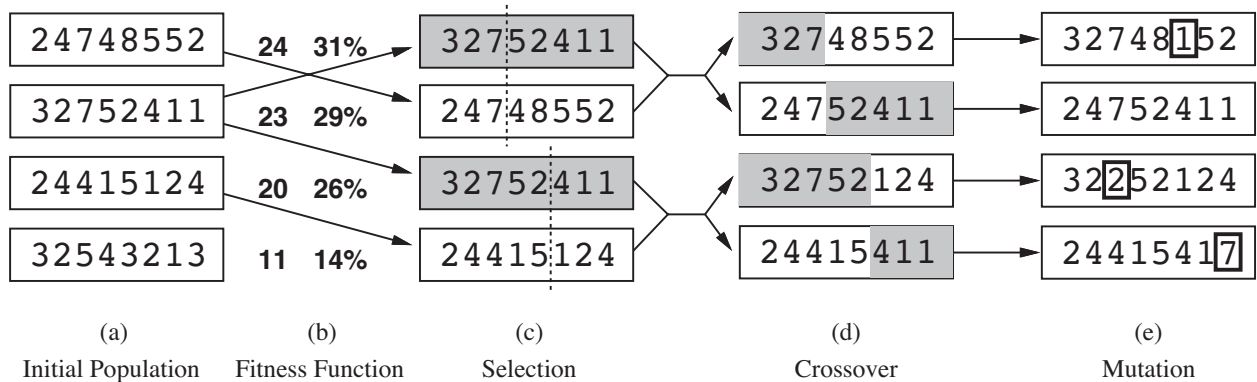

Stochastic beam search chooses k neighbors at random with higher evaluation states are more likely to be selected. Here, we define the probability p as

$$p = e^{-f(s)/T}$$

5.5. Genetic Algorithm

Genetic Algorithm is a variant of stochastic beam search. Different from other local search algorithms, GA generates successor states from a pair of states. (4.1.4)

Each state is denoted by a string over a finite set of alphabets, e.g. 0/1. Successor states are generated by two basic operations: crossover, and mutation. This is different from the other local search algorithms.



Code 5.5 Genetic Algorithm

```
1 import random, math
2
3 def genetic_algorithm(population, f, n):
4     # population = initial set of individuals
5     # f = evaluation function
6     # n = number of iterations
7     new_population = []
8     for i in range(n):
9         for j in range(len(population)):
10            v = evaluate(population, f)
11            x = select(population, v)
12            y = select(population, v)
13            child = cross_over(x, y)
14            r = random.uniform(0.0, 1.0)
15            if r < 0.05:
16                child = mutate(child)
17            new_population.append(child)
18        population = new_population
19    return max(population, key=f)
20
21 def evaluate(population, f):
22    return [f(p) for p in population]
23
24 def select(population, val):
25    m = sum(val)
26    r = random.uniform(0, m)
27    c = 0
28    for p, v in zip(population, val):
29        c += v
30        if c > r:
31            return p
32
33 def cross_over(x, y):
34    c = random.randint(0, len(x))
35    return x[0:c] + y[c+1:len(x)]
36
37 def mutate(x):
38    c = random.randint(0, len(x)-1)
39    x[c] = random.randint(0, 1)
40    return x
```

Exercise 5.4 To find the maximum value of the following function

$$f(x, y) = (1 - x^2)e^{-x^2-y} + x^3 - y^3$$

If x and y are set to vary between -3 and 3 . How can we represent an individual?

References

Russell, S. and Norvig, P. (2010). *Artificial Intelligence: A Modern Approach* (3rd edition). Pearson/Prentice Hall.

Michalewicz, F. and Fogel, D. B. (1998). *How to Solve It: Modern Heuristics*. Springer.