

Wednesday, March 30 - Day 26

(1)

Topic 11 - Hill Climbing

Next week: video lectures, will email

With Gradient Ascent as our inspiration, we want to think for ways to search for global optima in cases where the search space is:

- (1) discrete
- (2) continuous

Problem Setup:

- * Search space S of possible candidates
- * We may have some constraints
- * Scoring function: $\text{score}(x)$, $x \in S$
(also called "fitness" or "quality")

* A way to generate either:

doesn't
make as
much
sense
in cons.
problems

- all the candidates "near" some given candidate; this is called "the neighborhood", $\text{nbhd}(x)$.
- a random candidate near a given candidate, a "tweak", $\text{tweak}(x)$

"near" is for you to define, it depends on the problem, and the algorithm you're using.

Two running examples:

(1) TSP (n cities)

* discrete size of the search space: $(n-1)!$

* $\text{Score}(T) = \text{cost of the tour}$
(sum of the weights)
- we want to minimize

* $\text{nbhd}(T)$

Suppose $T = C_1 \rightarrow C_2 \rightarrow C_3 \rightarrow \dots \rightarrow C_n \rightarrow C_1$
Define the neighborhood of T to be all ways of picking two cities and swapping them (excluding C_1)

$\text{nbhd}(A \rightarrow B \rightarrow C \rightarrow D \rightarrow A)$

$= \{ A \rightarrow C \rightarrow B \rightarrow D \rightarrow A, \\ A \rightarrow B \rightarrow D \rightarrow C \rightarrow A, \\ A \rightarrow D \rightarrow C \rightarrow B \rightarrow A \}$

* $\text{tweak}(T) =$ a random thing in $\text{nbhd}(T)$

Let's say we have n cities. How big is $\text{nbhd}(T)$?

$$\binom{n-1}{2} = \frac{1}{2}(n-1)(n-2) \quad O(n^2)$$

(2) optimizing a continuous function in two variables $f(x,y)$

* continuous search space $\mathbb{R} \times \mathbb{R}$
 $[0,1] \times [0,1]$

* score of a candidate is the value of the function at that point

* $\text{nbhd}(x,y) =$ all points within some fixed distance δ of (x,y)

* $\text{tweak}(x,y) =$ pick something from $\text{nbhd}(x,y)$

MH #1: Random Search

best = random element of S

while True: (quit whenever you want)

x = random element of S

if $\text{score}(x) > \text{score}(\text{best})$:

best = x

Possible stopping conditions:

- * best score does not improve for N iterations
- * preset # of iterations
- * you get impatient

This is not a good metaheuristic. It doesn't use any old information to guide future choices.

[2 demos] $\left\{ \begin{array}{l} O1: \text{TSP random} \\ O2: \text{CNS func 1 - random} \end{array} \right.$

Gradient Ascent inspires the next one.

MH #2: Steepest Ascent Hill-Climbing (discrete only)

x = random element of S

while True:

$N = \text{nbhd}(x)$

$S =$ element of N with the best score

if $\text{score}(s) > \text{score}(x)$:

$x = s$

else:

quit

What does this do? Climbs right up the hill you start on.

Pros

* Finds a local optimum.

Cons

* Unlikely to find a global optimum unless your search space is very nice.

Demos:

03 - TSP Steep. Asc. swap 2
50 cities

* very slow, especially if the neighborhoods are big, like TSP

04 - TSP Steep. Asc. swap 2
300 cities

What's the slow part?

(1) generating the neighborhood

(2) scoring everything the neighborhood

$$d((x_1, y_1), (x_2, y_2)) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$$

To score with 300 cities, we do 300 distance calculations

(two subtractions, two squarings,
one addition, one square root)

This is slow when we have to do it

$$\binom{299}{2} = 44,551 \text{ times per move.}$$