Last update: September 18, 2012

LOCAL SEARCH ALGORITHMS

CMSC 421: Chapter 4, Sections 1 and 2

CMSC 421: Chapter 4, Sections 1 and 2 1

Iterative improvement algorithms

 \diamondsuit In many optimization problems, the **path** to a goal is irrelevant

- the goal state itself is the solution
- \diamond State space = a set of goal states
 - find one that satisfies constraints (e.g., no two classes at same time)
 - or find **optimal** one (e.g., highest possible value, least possible cost)

\diamond *Iterative improvement* algorithms

- $\diamond~$ keep a single "current" state, try to improve it
- Constant space
- Suitable for both offline and online search

Example: the n-Queens Problem

- \diamondsuit Put *n* queens on an $n \times n$ chessboard
 - No two queens on the same row, column, or diagonal
- \diamondsuit Iterative improvement:
 - Start with one queen in each column
 - move a queen to reduce number of conflicts



♦ Even for very large n (e.g., n = 1 million), this usually finds a solution almost instantly

Example: Traveling Salesperson Problem

- \diamond Given a *complete* graph (edges between all pairs of nodes)
- \diamond Find a least-cost *tour* (simple cycle that visits each city exactly once)
- \diamondsuit Iterative improvement:
 - Start with any tour, perform pairwise exchanges



 \diamondsuit Variants of this approach get within 1% of optimal very quickly with thousands of cities

Outline

- \diamondsuit Hill-climbing
- \diamondsuit Simulated annealing
- \diamond Genetic algorithms (briefly)
- \diamond Local search in continuous spaces (very briefly)

Hill-climbing (or gradient ascent/descent)

 \diamondsuit "Like climbing Everest in thick fog with amnesia"

```
\begin{array}{l} \textbf{function Hill-CLIMBING}(\textit{problem}) \\ \textit{current} \leftarrow \textit{new node containing problem's initial state} \\ \textbf{loop} \\ \textit{next} \leftarrow \textit{a highest-valued neighbor of current} \\ \textit{if VALUE}[\textit{next}] \leq \textit{VALUE}[\textit{current}] \textit{then return STATE}[\textit{current}] \\ \textit{current} \leftarrow \textit{next} \\ \\ \textbf{end} \end{array}
```

- \diamond VALUE[x] is x's objective-function value
 - how good we consider x to be
- \diamondsuit At each step, move to a neighbor of higher value in hopes of getting to a solution having the highest possible value
- \diamondsuit Can easily modify this for problems where we want to minimize rather than maximize

Hill-climbing, continued

 \diamondsuit State space "landscape":



- \diamond Random-restart hill climbing:
 - repeat with randomly chosen starting points
- \Diamond If finitely many local maxima, then $\lim_{\text{restarts}\to\infty} P(\text{complete}) = 1$

Simulated annealing

```
function SIMULATED-ANNEALING(problem, temperatures)

node \leftarrow a new node containing problem's initial state

for i \leftarrow 1 to \infty do

T \leftarrow temperatures[i]

if T = 0 then return STATE[node]

next \leftarrow a randomly selected neighbor of node

\Delta E \leftarrow \text{VALUE}[next] - \text{VALUE}[node] /* difference in desirability */

if \Delta E > 0 then node \leftarrow next

else with probability e^{\Delta E/T}, set node \leftarrow next
```

 \diamondsuit Idea: escape local maxima by allowing some "bad" moves

- but gradually decrease their frequency
- \diamondsuit Next two slides: a simple example, in Python

```
import cmath, random
def anneal_example(iterations):
   state = 0; history = [state]
   for i in range(iterations):
     T = 10 * (0.9 ** i)
     next = random.random()
     DeltaE = value(next) - value(state)
     target = cmath.e ** (DeltaE/T):
     if target > random.random():
        state = next
     history.append(round(state,3))
   return history
```

```
def value(state):
    return -abs(0.5-state)
```

function ANNEAL_EXAMPLE(iterations) state $\leftarrow 0$; history \leftarrow list containing state for $i \leftarrow 0$ to iterations -1 do $T \leftarrow 10 \times 0.9^{i}$ next \leftarrow a random number in [0,1) $\Delta E \leftarrow \text{VALUE}(next) - \text{VALUE}(state)$ target $\leftarrow e^{\Delta E/T}$ if target > random no. in [0,1) then state $\leftarrow next$ append state to history return history

function VALUE(state) return -|0.5 - state|

 \diamond A state is a number $s \in [0, 1]$. All states are neighbors

- \bigcirc Desirability = max{-(dist. from 0.3), .025 (dist. from .8)}
- \diamondsuit Start with *state* = 0; iterate for *i* = 0 to *iterations*-1
- \diamond On *i*'th iteration, temperature is 10×0.9^{i}
 - 10.0, 9.00, 8.10, 7.29, 6.56, 5.90, 5.31, 4.78, 4.30, 3.87, \ldots

Simple example, continued



Properties of simulated annealing

 \diamondsuit At fixed "temperature" T, probability of being in any given state x approaches a Boltzman distribution

 $p(x) = \alpha e^{\frac{E(x)}{kT}}$

 \diamondsuit For every state x other than x^* and for small T,

$$p(x^*)/p(x) = e^{\frac{E(x^*)}{kT}}/e^{\frac{E(x)}{kT}} = e^{\frac{E(x^*)-E(x)}{kT}} \gg 1$$

 \diamond It can be shown that if we decrease T slowly enough,

• $\Pr[\text{reach } x^*]$ approaches 1

 \diamondsuit Devised by Metropolis et al., 1953, for physical process modelling

• Widely used in VLSI layout, airline scheduling, etc.

Local beam search

 $\begin{array}{l} \textbf{function BEAM-SEARCH}(\textit{problem}, k) \textbf{ returns} a solution state \\ start with k randomly generated states \\ \textbf{loop} \\ generate all successors of all k states \\ \textbf{if} any of them is a solution \textbf{then} return it \\ \textbf{else} select the k best successors \end{array}$

 \diamondsuit Not the same as k parallel searches

- Searches that find good states will recruit other searches to join them
- Problem: often all k states end up on same local hill
- \diamondsuit Stochastic beam search:
 - Choose k successors randomly, biased towards good ones
 - Close analogy to natural selection

Genetic algorithms

- \diamondsuit Genetic algorithms
 - $\diamond~$ stochastic local beam search
 - $\diamond~$ generate successors from **pairs** of states
 - Each state should be a string of characters
 - Substrings should be meaningful components
- \diamond **Example**: *n*-queens problem
 - i'th character = row where i'th queen is located



Genetic algorithms

```
function GENETIC-ALGORITHM( population, FITNESS)

loop

new-population \leftarrow empty set

for i = 1 to |population|

choose x and y randomly^* from population, using FITNESS

child = REPRODUCE(x,y)

if small random probability then child = MUTATE(child)

add child to new-population

population \leftarrow new-population

if some individual is fit enough, or enough time has elapsed then

return arg max {FITNESS(x)|x \in population}

function REPRODUCE(x, y)
```

```
c \leftarrow \text{random integer in } \texttt{range(len}(x))
```

```
return x[:c] + y[c:]
```

* To choose x, usually use *relative fitness*: for each $x \in population$, $\Pr[choose x] = FITNESS(x)/\sum_z FITNESS(z)$. To choose y, usually use relative fitness or a uniform distribution

Genetic algorithms



Fitness Selection Pairs Reproduction Mutation







Same simple example as before

- \diamondsuit A state is a number $s \in [0,1].$ All states are neighbors
- Fitness = max{-(dist. from 0.3), .025 (dist. from .8)}

```
import random
stringsize = 4; top = 10**stringsize
def example(times=100, popsize=200, mutation_prob=.001):
    population = [int(random.random() * top) for j in range(popsize)]
    history = []
    for i in range(times):
        fitness = [value(population[j]) for j in range(popsize)]
        fitmax = -float('inf')
        for j in range(popsize):
            if fitness[j] > fitmax:
                fitmax = fitness[j]; jbest = j
        history.append(population[jbest])
        fit_sum = sum(fitness)
        probs = [x/fit_sum for x in fitness]
        cutoff = [sum(probs[:j+1]) for j in range(popsize)]
        children = []
        for j in range(popsize):
            r = random.random()
            for k in range(popsize):
                if r < cutoff[k]: break
            par1 = population[k-1]:
            par2 = population[int(random.random() * popsize)]
            split = int(random.random() * (stringsize+1))
            child = str(par1)[:split] + str(par2)[split:]
            if random.random() < mutation_prob:
                where = int(random.random() * stringsize)
                what = str(int(random.random() * 10))
                child = child[0:where] + what + child[where+1:]
            children.append(int(child))
        population = children
    return history
def value(s):
    return max(-abs(round(0.3*top) - s), .025 - abs(round(<math>0.8*top) - s))
```



- 100 iterations, population size = 200, mutation probability = .001
- Running time is much worse than for simulated annealing

Discussion

- \diamondsuit Genetic algorithms \neq biological evolution
 - for example, real genes encode replication machinery
- Genetic algorithms are widely used by engineers for optimization problems
 circuit layout, job-shop scheduling
 - Not clear whether this is due to performance or intuitive appeal
- \diamondsuit More work is needed to identify conditions under which they perform well

Hill-climbing in continuous state spaces

 \diamond Suppose we want to put three airports in Romania – what locations?

- 6-D state space defined by $(x_1, y_2), (x_2, y_2), (x_3, y_3)$
- Objective function $f(x_1, y_2, x_2, y_2, x_3, y_3)$ measures desirability
 - $\diamond~$ e.g., sum of squared distances from each city to nearest airport



Hill-climbing in continuous state spaces

- \diamondsuit A technique from numerical analysis:
- \diamondsuit Given a surface z = f(x, y) and a point (x, y)
- \diamond A **gradient** is a vector

 $\nabla f(x,y) = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)$

- points in the direction of the steepest slope
- length is proportional to the slope



 \diamond Gradient methods compute ∇f and use it to increase/reduce f

 $\diamond \text{ e.g., } \mathbf{x} \leftarrow \mathbf{x} - \alpha \nabla f(\mathbf{x})$

• If $\nabla f = 0$ then you've reached a local maximum/minimum

Hill-climbing in continuous state spaces

Suppose we want to put three airports in Romania – what locations? $\langle \rangle$ $\nabla f = \left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial y_1}, \frac{\partial f}{\partial x_2}, \frac{\partial f}{\partial y_2}, \frac{\partial f}{\partial x_3}, \frac{\partial f}{\partial y_3}\right)$

Look for $x_1, y_1, x_2, y_2, x_3, y_3$ such that $\nabla f(x_1, y_1, x_2, y_2, x_3, y_3) = 0$ $\langle \rangle$



Continuous state spaces, continued

 \diamond Sometimes can solve for $\nabla f(\mathbf{x}) = 0$ exactly (e.g., with one city)

 \Diamond Newton-Raphson method (1664, 1690)

- solve $\nabla f(\mathbf{x}) = 0$ by iterating $\mathbf{x} \leftarrow \mathbf{x} \mathbf{H}^{-1}(\mathbf{x}) \nabla f(\mathbf{x})$ • \mathbf{H} is an $n \times n$ matrix with $\mathbf{H}_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$
- I used to know this stuff, but that was a **long** time ago
- \diamond *Discretization* methods turn continuous space into discrete space
 - e.g., *empirical gradient* considers $\pm \delta$ change in each coordinate

