

CS540 Introduction to Artificial Intelligence

Lecture 17

Young Wu

Based on lecture slides by Jerry Zhu and Yingyu Liang

July 22, 2019

Search Algorithms

- Breadth-First, BiDirectional
 - ~~Depth-First~~, Iterative Deepening
 - Uniform Cost
 - ~~Best First Greedy~~
 - ~~A~~ (or A^*), Iterative Deepening A, ~~Beam~~
- Handwritten notes and annotations:
- Green box around "Breadth-First, BiDirectional" with an arrow pointing to "optimal if $c=1$ ".
 - Green oval around "Iterative Deepening" with an arrow pointing to "optimal if $c=1$ ".
 - Green box around "Uniform Cost" with an arrow pointing to "optimal".
 - Green oval around " A^* " with an arrow pointing to "optimal".
 - Green oval around "Iterative Deepening A" with an arrow pointing to "optimal".
 - Green oval around "IDA*" with an arrow pointing to "optimal".
 - Green arrow from "Iterative Deepening A" to "IDA*".
 - Green arrow from "IDA*" to "optimal".

Iterative Deepening A Star Search


Discussion

- A^* can use a lot of memory.
- Do path checking without expanding any vertex with $g(s) + h(s) > 1$. } $g+h \leq 1$
- Do path checking without expanding any vertex with $g(s) + h(s) > 2$. } $g+h \leq 2$
- ...
- Do path checking without expanding any vertex with $g(s) + h(s) > d$. } $g+h \leq 3$
}

Beam Search

Discussion

- Version 1: Keep a priority queue with fixed size k . Only keep the top k vertices and discard the rest.
- Version 2: Only keep the vertices that are at most ϵ worse than the best vertex in the queue. ϵ is called the beam width.



$\epsilon = 2$ if see $9th > 3 \rightarrow$ do not enqueue

\Rightarrow Small queue during search.

Beam Search Performance

Discussion


- Beam is incomplete.
- Beam is not optimal.

Traveling Salesperson Example

Motivation

goal A \xrightarrow{BCDE} A

tour: A B C D E A



State Space \nearrow too large, State

APCEBA \leftarrow

always a solution
may not be the best

Search vs. Local Search

Motivation

- Some problems do not have an initial state and a goal state.
- Every state is a solution. Some states are better than others, defined by a cost function (sometimes called score function in this setting), $f(s)$.
- The search strategy will go from state to state, but the path between states is not important.
- There are too many states to enumerate, so standard search through the state space methods are too expensive.

local
Search

$$f(s) = g(s) + h(s) \rightarrow \text{Search}$$

Local Search

Motivation

- Local search is about searching through a state space by iteratively improving the cost to find an optimal or near-optimal state.
- The successor states are called the neighbors (sometimes move set).
- The assumption is that similar ^{neighbors} (nearby) solutions have similar costs.

gradient descent

→ cost (w)

Local Search Application

Motivation

- Optimization problems (gradient descent methods are all local search methods)
- Traveling salesman
- Boolean satisfiability (SAT)
- Scheduling

Boolean Satisfiability Example, Part I

Quiz (Graded)

Q2 ABCE

- Assume all variables A, B, C, D, E are set to True. How many of the following clauses are satisfied?

- OR A: $A \vee \neg B \vee C$ ✓
- B: $\neg A \vee C \vee D$ ✓
- C: $B \vee D \vee \neg E$ ✓
- ~~D: $\neg C \vee \neg D \vee \neg E$~~
- E: $\neg A \vee \neg C \vee E$ ✓

SAT
max # of clauses
satisfied
return True.

Boolean Satisfiability Example, Part II

Quiz (Graded)

Q4 CD

- Assume all variables A, B, C, D, E are set to True. Which one of the variables should be changed to False to maximize the number of clauses satisfied?

- $A \vee \neg B \vee C$
- $\neg A \vee C \vee D$
- $B \vee D \vee \neg E$
- $\neg C \vee \neg D \vee \neg E$
- $\neg A \vee \neg C \vee E$

C ✓

D ✓

$A B C \neg D E$

A B C D E

↓

A B \neg C D E

Hill Climbing (Valley Finding)

Description

- Start at a random state.
- Move to the best neighbor state (one of the successors).
- Stop when all neighbors are worse than the current state.
- The idea is similiar to gradient descent.

Hill Climbing

Algorithm

- Input: state space S and cost function f .
- Output: $s^* \in S$ that minimizes $f(s)$.
- Start at a random state s_0 .
- At iteration t , find the neighbor that minimizes f .

$$s_{t+1} = \arg \min_{s \in S'(s_t)} f(s)$$


- Stop when none of the neighbors have a lower cost.

$$\text{stop if } f(s_{t+1}) \leq f(s_t)$$

Hill Climbing Performance

Discussion

- It does not keep a frontier, so no jumping and no backtracking.
- It is simple, greedy, and stops at a local minimum.





- A simple modification is picking random initial states multiple times and finding the best among the local minima.

First Choice Hill Climbing


Discussion

- If there are too many neighbors, randomly generate neighbors until a better neighbor is found.
- This method is called first choice hill climbing.

not for
jump out
of local min

Walk SAT Example

Discussion



- Pick a random unsatisfied clause.

- Select and flip a variable from that clause:

- 1 With probability p , pick a random variable. *avoid local min.*

- 2 With probability $1 - p$, pick the variable that maximizes the number of satisfied clauses.

- Repeat until the solution is found.

- Walk SAT uses the idea of stochastic hill climbing.

$p = 0.1$

$A B C D E = T$

$\neg B \vee \neg C \vee \neg D$

$\frac{1}{3} B, \frac{1}{3} C, \frac{1}{3} D$

check # of clauses

satisfied $B \rightarrow F$

$C \rightarrow F$

$D \rightarrow F$


Simulated Annealing

Description

- Each time, a random neighbor is generated.
- If the neighbor has a lower cost, move to the neighbor.
- If the neighbor has a higher cost, move to the neighbor with a small probability.
- Stop until bored.
- It is a version of Metropolis-Hastings Algorithm.

→ Similar to first choice hill climbing

→ jump out of local min.



Acceptance Probability

Definition

- The probability of moving to a state with a higher cost should be small.

1 Constant: $p = 0.1$


2 Decreases with time: $p = \frac{1}{t}$

3 Decreases with time and as the energy difference increases:

$$p = \exp\left(-\frac{|f(s') - f(s)|}{\text{Temp}}\right)$$

- The algorithm corresponding to the third idea is called simulated annealing.

cost difference



$P_2 > P_1$

Temperature

Definition

$$p = \exp\left(-\frac{|f(s') - f(s)|}{t}\right)$$

t Temp

- ~~The t in the above expression does not have to be time.~~
- It can represent temperature which decreases over time. For example, the temperature can change geometrically.

$$\overline{t} = 0.9t$$

$$T = 0.9, 0.8, 0.7$$

$$T = \frac{1}{7}, \frac{1}{2}, \frac{1}{3}, \dots$$

- High temperature: almost always accept any s' .
- Low temperature: first choice hill climbing.

Simulated Annealing

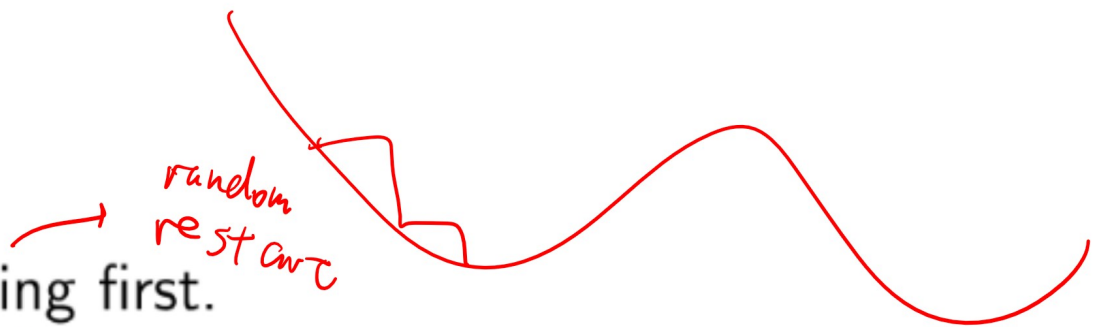
Algorithm

- Input: state space S , temperature function Temp , and cost function f .
- Output: $s^* \in S$ that minimizes $f(s)$.
- Start at a random state s_0 .
- At iteration t , generate a random neighbor s' , and update the state according to the following rule.

$$s_{t+1} = \begin{cases} s' & \text{if } f(s') > f(s_t) \\ s' & \text{with probability } \exp\left(-\frac{|f(s') - f(s_t)|}{\text{Temp}(t)}\right) \\ s_t & \text{otherwise} \end{cases}$$

Simulated Annealing Performance

Discussion

- 
- Use hill-climbing first.
 - Neighborhood design is the most important.
 - In theory, with infinitely slow cooling rate, SA finds global minimum with probability 1.

Genetic Algorithm

Description

- Start with a fixed population of initial states.
- Find the successors by:
 - 1 Cross over.
 - 2 Mutation.

Reproduction Probability

Definition

- Each state in the population has probability of reproduction proportional to the fitness. Fitness is the opposite of the cost: higher cost means lower fitness. Use F to denote the fitness function, for example, $F(s) = \frac{1}{f(s)}$ is a valid fitness function.

$$p_i = \frac{F(s_i)}{\sum_{j=1}^N F(s_j)}, i = 1, 2, \dots, N$$

- A pair of states are selected according to the reproduction probabilities (using CDF inversion).

Cross Over

Definition

- The states need to be encoded by strings.
- Cross over means swapping substrings.
- For example, the children of 10101 and 01010 could be the same as the parents or one of the following variations.

(11010, 00101), (10010, 01101)
(10110, 01001), (10100, 01011)

1 0 1 0 1
0 1 0 1 0

1 0 1 0 1
0 1 0 1 0

Mutation

Definition

- The states need to be encoded by strings.
- Mutation means randomly updating substrings. Each character is changed with small probability q , called the mutation rate.
- For example, the mutated state from 000 could stay the same or be one of the following.


one of 001, 010, 100, with probability $q(1 - q)^2$

one of 011, 101, 110, with probability $q^2(1 - q)$

and 111, with probability q^3

Cross Over, Modifications

Definition



2-point
Cross over

- The previous cross over method is called 1 point cross over.
- It is also possible to divide the string into N parts. The method is called N point cross over.
- It is also possible to choose each character from one of the parents randomly. The method is called uniform cross over.

10111
01111



1 0 1 0 1 } 1 5% 0 5%
0 1 0 1 0 } 0 5% 1 5% - -

Mutation, Modifications

Definition

- For specific problems, there are ways other than flipping bits to mutate a state.


- 1 Two-swap: ABCDE to EBCDA
- 2 Two-interchange: ABCDE to EDCBA



Genetic Algorithm SAT Example

Discussion

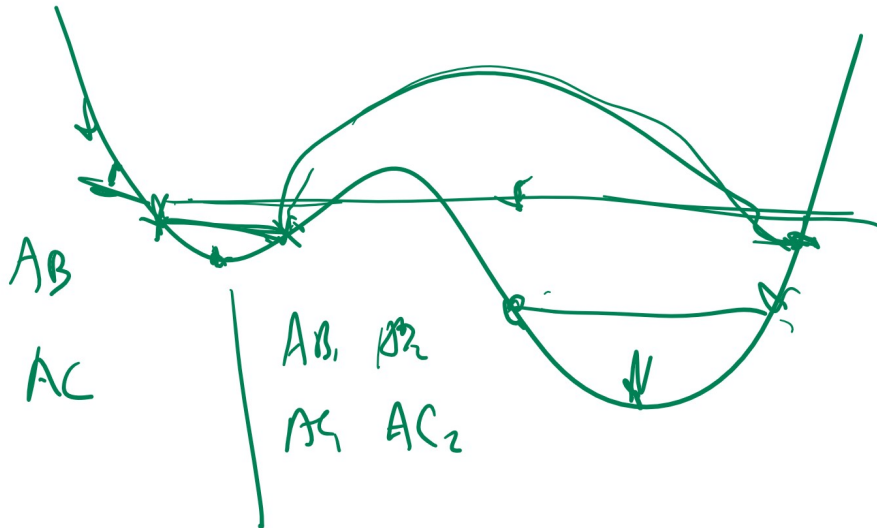
01010



10

000
↓
010


A
B
C
~~0~~



~

Genetic Algorithm TSP Example

Discussion



Fitness Example

Quiz (Graded)

- Fall 1999 Final Q5
- Which ones (multiple) of the following states have the highest reproduction probability?

population of 5

The fitness function is $5A + 3BC - D + 2E$.

max not cost

$$\Rightarrow F(s)$$

(A, B, C, D, E)

- A: (1, 1, 0, 1, 1)
- B: (0, 1, 1, 0, 1)
- C: (1, 1, 0, 0, 0)
- D: (1, 0, 1, 1, 1)
- E: (1, 0, 0, 0, 0)

$$5 + 0 - 1 + 2 = 6$$

$$p = \frac{F(s)}{\sum_{s'} F(s')}$$

$$5 + 0 - 1 + 2 = 6$$

$$\frac{6}{27}$$

$$\frac{5}{27}$$

prob parent 1 $\rightarrow \frac{6}{27}$ prob A $\frac{5}{27}$ prob B ...

Genetic Algorithm, Part I

Algorithm

- Input: state space S represented by strings s and cost function f or fitness function F .
- Output: $s^* \in S$ that minimizes $f(s)$.
- Randomly generate N solutions as the initial population.

$$s_1, s_2, \dots, s_N$$

- Compute the reproduction probability.

$$p_i = \frac{F(s_i)}{\sum_{j=1}^N F(s_j)}, i = 1, 2, \dots, N$$

Genetic Algorithm, Part II

Algorithm

- Randomly pick two states according to p_i , say s_a, s_b .
Randomly select a cross over point c , swap the strings.

$$s'_a = s_a [0...c) s_b [c...m)$$

$$s'_b = s_b [0...c) s_a [c...m)$$

- Randomly mutate each position of each state s_i with a small probability (mutation rate).

$$s'_i [k] = \begin{cases} s_i [k] & \text{with probability } 1 - q \\ \text{random} & \text{with probability } q \end{cases}, k = 1, 2, \dots, m$$

- Repeat with population s' .


Variations

Discussion

$F(s)$	6	5	5	6	5
ranking	4	1	1	4	1
prob	$\frac{4}{11}$	$\frac{1}{11}$	---	---	---

- Parents can survive.
- Use ranking instead of $F(s)$ to compute reproduction probabilities.
- Cross over random bits instead of chunks.

↳ uniform cross over



x, y
 11 2
 2 1

Score = # times
 its flying.