

CS540 Introduction to Artificial Intelligence

Lecture 20

Young Wu

Based on lecture slides by Jerry Zhu and Yingyu Liang

August 8, 2022

Pirate Game Example

Quiz

- 5 pirates got 100 gold coins. Each pirate takes a turn to propose how to divide the coins, and all pirates who are still alive will vote whether to accept the proposal or reject the proposal, kill the pirate, and continue to the next round. Use strict majority rule for the vote, and use the assumption that if a pirate is indifferent, he or she will vote reject with probability 50 percent.
- How will the first pirate propose?
- $A : (0, 0, 0, 0, 100)$
- $B : (20, 20, 20, 20, 20)$
- $C : (94, 0, 1, 2, 3)$
- $D : (97, 0, 1, 0, 2)$
- $E : (98, 0, 1, 0, 1)$

Pirate Game Example Diagram

Quiz

Summary

Discussion

- Search:
 - 1 Uninformed.
 - 2 Informed.
 - 3 Local Search: Hill Climbing (Valley Finding).
 - 4 Local Search: Simulated Annealing.
 - 5 Local Search: Genetic Algorithm.
 - 6 Adversarial.

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 (nearby) solutions have similar costs.

Local Search Application

Motivation

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

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 similar to gradient descent.

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.

Annealing

Definition

- The annealing process of heated solids.
- Anneal: to subject (glass or metal) to a process of heating and slow cooling to toughen and reduce brittleness.
- Alloys manage to find a near global minimum energy state when heated and then slowly cooled.

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

- The algorithm corresponding to the third idea is called simulated annealing. The Temperature function $T(t)$ should be a decreasing in time t (iteration number).

Temperature

Definition

- T represents temperature which decreases over time. For example, the temperature can change arithmetically or geometrically.

$$T(t+1) = \max\{T(t) - 1, 1\}, T(0) = \text{large}$$

$$T(t+1) = 0.9T(t), T(0) = \text{large}$$

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

Simulated Annealing Example 1

Quiz

- Suppose we are minimizing and $f(s) = 6, f(s') = 5, T = 4$.
What is the probability we move from s to s' in the next step?
What is the probability we move from s' to s in the next step?

Simulated Annealing Example 2

Quiz

- Suppose we are minimizing and $f(s) = 0, f(s') = \log(5), T = 1$. What is the probability we move from s to s' .
- A : 0
- B : $\frac{1}{5}$
- C : $\frac{4}{5}$
- D : 1
- E : I don't understand.

Simulated Annealing Performance

Discussion

- Use hill-climbing first.
- Neighborhood design is the most important.
- In theory, with infinitely slow cooling rate, Simulated Annealing 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)

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

- 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.

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 TSP Example

Discussion

Fitness Example 1

Quiz

- Which one of the following states have the highest reproduction probability? The fitness function is $f(x) = \min \{t \in \{1, 2, 3, 4, 5, 6\} : x_t = 1\}$ with $x_6 = 1$.
- $A : (0, 0, 1, 0, 0)$
- $B : (0, 1, 0, 0, 1)$
- $C : (0, 0, 1, 1, 0)$
- $D : (0, 0, 0, 1, 0)$
- $E : (0, 0, 0, 0, 0)$

Fitness Example 2

Quiz

- Which one of the following states have the highest reproduction probability? The fitness function is $f(x) = \max \{t \in \{0, 1, 2, 3, 4, 5\} : x_t = 1\}$ with $x_0 = 1$.
- $A : (0, 0, 1, 0, 0)$
- $B : (0, 1, 0, 0, 1)$
- $C : (0, 0, 1, 1, 0)$
- $D : (0, 0, 0, 1, 0)$
- $E : (0, 0, 0, 0, 0)$

Variations

Discussion

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

Genetic Algorithm Performance

Discussion

- Use hill-climbing first.
- State design is the most important.
- In theory, cross over is much more efficient than mutation.

Summary

Discussion

- Search:
 - ① Uninformed.
 - ② Informed.
 - ③ Local Search: Hill Climbing (Valley Finding): Start at a random state → Move to the best successor → Repeat.
 - ④ Local Search: Simulated Annealing: Start at a random state → Generate a random successor → Move if better, Move with small probability if worse → Repeat.
 - ⑤ Local Search: Genetic Algorithm: Start with many random states → Cross-over according to fitness → Mutation → Repeat.
 - ⑥ Adversarial.