

Q1. Local Search

(a) Hill Climbing

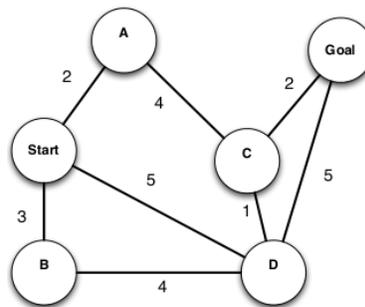
- (i) Hill-climbing is complete. True False
- (ii) Hill-climbing is optimal. True False

(b) Simulated Annealing

- (i) The higher the temperature T is, the more likely the randomly chosen state will be expanded. True False
- (ii) In one round of simulated annealing, the temperature is 2 and the current state S has energy 1. It has 3 successors: A with energy 2; B with energy 1; C with energy $1 - \ln 4$. If we assume the temperature does not change, What's the probability that these states will be chosen to expand after S eventually?
- (iii) On a undirected graph, If T decreases slowly enough, simulated annealing is guaranteed to converge to the optimal state. True False

(c) Local Beam Search

The following state graph is being explored with 2-beam graph search. A state's score is its accumulated distance to the start state and lower scores are considered better. Which of the following statements are true?



- States A and B will be expanded before C and D.
- States A and D will be expanded before B and C.
- States B and D will be expanded before A and C.
- None of above.

(d) Genetic Algorithm

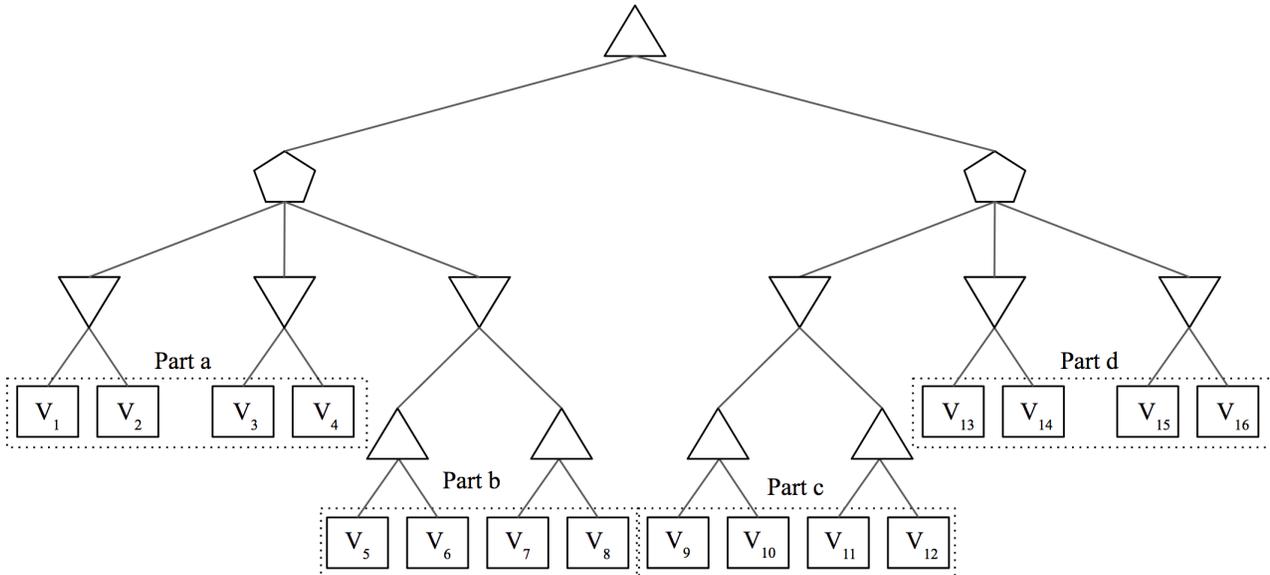
- (i) In genetic algorithm, cross-over combine the genetic information of two parents to generate new offspring. True False
- (ii) In genetic algorithm, mutation involves a probability that some arbitrary bits in a genetic sequence will be flipped from its original state. True False

(e) Gradient Descent

- (i) Gradient descent is optimal. True False
- (ii) For a function $f(x)$ with derivative $f'(x)$, write down the gradient descent update to go from x_t to x_{t+1} . Learning rate is α .

Q2. MedianMiniMax

You're living in utopia! Despite living in utopia, you still believe that you need to maximize your utility in life, other people want to minimize your utility, and the world is a 0 sum game. But because you live in utopia, a benevolent social planner occasionally steps in and chooses an option that is a compromise. Essentially, the social planner (represented as the pentagon) is a median node that chooses the successor with median utility. Your struggle with your fellow citizens can be modelled as follows:



There are some nodes that we are sometimes able to prune. In each part, mark all of the terminal nodes such that **there exists a possible situation** for which the node **can be pruned**. In other words, you must consider **all** possible pruning situations. Assume that evaluation order is **left to right** and all V_i 's are **distinct**.

Note that as long as there exists ANY pruning situation (does not have to be the same situation for every node), you should mark the node as prunable. Also, alpha-beta pruning does not apply here, simply prune a sub-tree when you can reason that its value will not affect your final utility.

- (a) V_1
 V_2
 V_3
 V_4
 None

- (b) V_5
 V_6
 V_7
 V_8
 None

- (c) V_9
 V_{10}
 V_{11}
 V_{12}
 None

- (d) V_{13}
 V_{14}
 V_{15}
 V_{16}
 None