4511W, Spring-2020
ASSIGNMENT 3:
**Assigned: 03/03/20 Due: 03/18/20 at 11:55 PM** (submit via Canvas, you may take a picture of handwritten solutions, but you must put them in a pdf) Submit only pdf or txt files

Written/drawn:
Problem 1. (15 points)
Run A* on the following graph starting at "Karlsruhe". Ensure you show both the fringe (with corresponding f-cost values) and explored set.

(Note: the edge between Nurnberg and Munchen is 167 not just 67.)

Node	Heuristic
Augsburg	350
Erfurt	260
Frankfurt	300
Karlsruhe	430
Kassel	200
Mannheim	370
Munchen	320
Nurnberg	160
Stuttgart	0
Wurzburg	90

Problem 2. (15 points)
On the graph below, run the basic hill climbing algorithm starting at node A. (1) Run basic hill-climbing. (2) Then run stochastic hill-climbing (starting with node A) using the randomly generated numbers below. You should attempt to travel to cities in lexicographical order. So if you could go to cities A, F, and Z at probabilities (A,10%), (F,70%) and (Z,20%)... then on a random number 0-0.1 should be A, 0.1-0.8 should be F, and 0.8-1.0 should be Z (the numbers are designed so the edge cases won't be hit).

For this problem we want to minimize the heuristic, so use the method we discussed in class for finding the probability to go to a different node. Specifically, for better nodes you should compute the probability using reciprocals (i.e. using an absolute scale rather than relative). For example if h(node X)=5 and h(node Y)=9 are the two places you want to go from some node, then:
$P(\text{node X}) = \frac{1/5}{1/5+1/9} = 0.643$. Assume 1/0 is infinity (so you always go to there).

Run 1: 0.566 0.753 0.532 0.753 0.598 0.004 0.103 0.544 0.693 0.974

Run 2: 0.014 0.186 0.975 0.512 0.407 0.974 0.098 0.350 0.686 0.493

Run 3: 0.039 0.860 0.330 0.733 0.022 0.110 0.678 0.046 0.101 0.769

Heuristic
A=32
B=30
C=12
D=20
E=4
F=6
G=8
H=1
I=2
J=0

Problem 3. (15 points)
For the hill-climbing done above. (1) What is the probability when doing stochastic hill-climbing to reach the optimal answer? (2) If the problem only has 10 states (as shown above), is random-restart or the stochastic-hill climbing better?

Problem 4. (15 points)
Answer the following questions:

(1) Assume you have a consistent heuristic when running A* search on some graph. What property will the fringe set have as you run the algorithm?

(2) Assume h1 and h2 are admissible. Prove h3 is admissible, where h3 is defined as: (h1+h2)/2

Problem 5. (15 points)
Find an admissible non-trivial heuristic (i.e. not h(n) = 0, or something similarly trivial) for the following cases:

(1) The vacuum world with 5 states (as opposed to the normal 2). All 5 states are initially dirty and arranged in one dimension (i.e. on a line), so the actions = {Suck, Left, Right} remain unchanged. The agent starts in the middle state and wants to minimize the number of actions before all squares are clean. (Assume no additional dirt is added to the problem while the robot is running.)

(2) Tic-tac-toe on an NxN board where you need K in-a-row to win. You want to find how to win in the least number of moves.

Problem 6. (15 points)
Below is a mini-max tree to apply alpha-beta pruning on. When performing the depth-first search, you can choose any branch to go down, without being consistent on picking (e.g. you don't always need to go left branch first). You need to do this two separate ways:
(1) List the sequence of nodes you should explore (and indicating where you can prune) to result in the **most** pruning possible.
(2) List the sequence of nodes you should explore (and indicating where you can prune) to result in the **least** pruning possible.

Programming (python/lisp):
Problem 7. (10 points)
Use the given genetic algorithm in search.py and apply it to the n-queens problem. Analyze this on a number of different configuration to answer the following questions:
(1) Is the genetic algorithm good in general to solve n-queens? Why or why not?
(2) What is the purpose of f_ thresh and how did you handle/use it in your testing?
(3) What is one positive thing about the genetic algorithm in general? One negative aspect?