CSC 665: Artificial Intelligence

Homework 0

By turning in this assignment, I agree to abide by SFSU's academic integrity code and declare that all of my solutions are my own work.

Welcome to your first CSC 665 assignment! The purpose of this homework is to give you a sense of the background knowledge we'll be drawing on throughout the course. If you can handle the problems here, then you should be able to succeed in the course. If not, then you may need to do some additional work to refresh yourself on the prerequisites.

This homework is worth 2% of extra credit, and is thus technically optional. However, you are strongly encouraged to complete the problems in this assignment as preparation for the remainder of the course.

1 About you

a. (1 point) What are your pronouns?

b. (1 point) Please list any relevant courses you have taken in probability, linear algebra, calculus, discrete math, or algorithms. Include both the course name and number.

c. (1 point) In an effort to create an active class environment, I sometimes call on students by name to invite them to share their thoughts or answer a question. If the prospect of being called on fills you with dread or anxiety, however, then I will refrain from bringing you into the discussion against your will. Please indicate whether you are comfortable being called on. (Note that your grade will not be negatively affected if you answer no.)

2 Optimization

In this class, we will cast a lot of AI tasks as optimization problems. In an optimization problem, we are asked to find the value of a parameter that maximizes or minimizes a function. Concretely, given a function f, we might like to find a value θ^* such that $f(\theta^*) \leq f(\theta)$ for any θ in the domain of f.

a. (1 point) Let a, b, and c be positive real numbers. Consider the quadratic function

$$f(\theta) = a\theta^2 + b\theta + c.$$

Note that θ here is a real number. What value of θ minimizes $f(\theta)$?

b. (1 point) Let $x_1, \ldots, x_n$ be real numbers. Consider the quadratic function

$$g(\theta) = \sum_{i=1}^{n} (\theta - x_i)^2.$$

What value of θ minimizes $g(\theta)$?

c. (1 point) Let $x_1, \ldots, x_n$ again be real numbers, and let $w_1, \ldots, w_n$ be positive real numbers that we can interpret as representing the importance of each of the x_i's. Consider the weighted quadratic function

$$h(\theta) = \sum_{i=1}^{n} w_i(\theta - x_i)^2.$$

What value of θ minimizes $h(\theta)$?

d. (1 point) What issue could arise in the minimization of h if some of the w_i's are negative?

3 Probability

Later in the course we will study probabilistic graphical models, which offer a framework for representing complex probability distributions that can be used to model the real world. We will review some of the relevant concepts from probability theory in class, but as a preview, these are the types of problems that you might be asked to solve.

a. (1 point) Consider a standard 52-card deck of cards with 13 card values (Ace, King, Queen, Jack, and 2–10) in each of the four suits (clubs, diamonds, hearts, spades). If a card is drawn at random, what is the probability that it is a spade or a two?

b. (1 point) Two factories — Factory A and Factory B — design batteries to be used in mobile phones. Factory A produces 60% of all batteries, and Factory B produces the other 40%. 2% of Factory A's batteries have defects, and 4% of Factory B's batteries have defects. What is the probability that a battery is both made by Factory A and defective?

c. (1 point) Consider the following (made up) facts about COVID incidence and testing:

- In the absence of any special information, the probability that a person has COVID is 1%.
- If a person has COVID, the probability that a test will correctly read positive is 80%.
- If a person does not have COVID, the probability that a test will incorrectly produce a false positive is 10%.

Suppose you take a COVID test and it reads positive. Given the facts above, what is the probability that you have COVID?

d. (2 points) Suppose you repeatedly roll a fair six-sided die until you roll a 1 (and then you stop). Every time you roll a 3, you win a points, and every time you roll a 6, you lose b points. You do not win or lose any points if you roll a 2, 4, or 5. What is the expected number of points (as a function of a and b) you will have when you stop?

Hint: You will find it helpful to define a recurrence. If you define V as the expected number of points you get from playing the game, what happens if you roll a 3? You win a points and then get to play again. What about the other cases? Can you write this as a recurrence?

4 Counting

a. (1 point) Suppose we have an $n \times n$ grid of points in which we'd like to place an arbitrary axis-aligned rectangle (i.e., the sides of the rectangle are parallel to the axes of the grid). Each corner of the rectangle must be one of the points in the grid, but otherwise there are no constraints on the location or size of the rectangle. In particular, it is possible for all four corners of the rectangle to be the same point (resulting in a rectangle with 0 area). How many ways are there to place such a rectangle on the grid? Give your answer using big-O notation.

b. (1 point) Now suppose we wish to place three such rectangles on the grid. Again, there are no constraints on the locations of the rectangles. In particular, it is possible for all three rectangles to be on top of each other. How many possible ways are there to place three rectangles on the grid? Again, give your answer using big-O notation.

5 Programming in Python

In this problem, you will implement three short functions and one medium function. Detailed instructions can be found in the source code file.

Do not import any outside libraries (e.g. numpy). Only the Python standard library and the libraries imported in the starter code are allowed.

a. (1 point) Implement `find_alphabetically_first_word`.

b. (1 point) Implement `euclidean_distance`.

c. (1 point) Implement `find_non_singleton_words`.

d. (3 points) Implement `mutate_sentences`.

Submission

Submission is done on Canvas. You should submit two files: one containing your solutions to the written problems, the other containing your implementations of the coding problems.

- Submit your written solutions in a single PDF file. Make sure to clearly indicate the number and letter of the problem corresponding to each solution. It is okay to hand-write your solutions and then scan them into a PDF, but *only if your handwriting is legible.*
- Submit your coding solutions in the provided `.py` file. Do not modify the name of this file after downloading it from the course website.