(R) GloVe Compartment (1/3) [5 points]

How can we represent the meaning of a word in a way that computers can understand? A popular solution is to represent each word as a vector (a list of numbers). For instance, the word *dog* might be represented as [0.7, 1.9, -4.3, 5.6, 0.0, -0.5]. Computers are great at processing numbers, so vectors are a very computer-friendly way to represent information. How do we decide which numbers should go in the vector for a given word? One popular approach was described by the linguist John Firth in 1957: "You shall know a word by the company it keeps." That is, the vector for a given word is meant to encode which other words tend to occur near this word.

One successful recent approach that builds on this idea is GloVe (Global Vectors for word representation), in which the vectors encode *differences* between words in terms of the other words that occur near them. For example, suppose we were trying to understand what the words *cake* and *cakes* mean. The word *birthday* appears near both of these words, so it is not helpful in distinguishing them. However, if we look at the words *is* and *are*, we would see that *is* appears more often near *cake* than near *cakes*, while *are* appears more often near *cakes* than *cake*. Thus, *is* and *are* can help us distinguish *cake* from *cakes*! Furthermore, we can observe that *candle* and *candles* have a similar relationship to *is* and *are*, so we can conclude that the difference between *cake* and *cakes* is the same as the difference between *candle* and *candles*!

Researchers at NACLO Labs have found some four-dimensional vectors that represent a few English words almost as well as GloVe vectors do! Here are two of the vectors they found:

man: [0.5, 0.9, 0.3, 0.3]

woman: [0.5, 0.9, 0.1, 0.5]

daughter: [0.5, 0.7, 0.2, 0.3]

The researchers had more representations, but they somehow mixed up the vectors and the words that they correspond to! Below are some English words and their corresponding vectors, in no particular order:

1. [0.5, 0.9, 0.2, 0.4]	A. girl
2. [0.5, 0.7, 0.4, 0.1]	B. queen
3. [0.5, 0.9, 0.3, -0.5]	C. prince
4. [0.5, 0.7, 0.2, 1.1]	D. boy
5. [0.5, 0.8, 0.9, 1.3]	E. father
6. [0.5, 0.8, 0.9, 0.5]	F. mother
7. [0.5, 0.8, 0.7, 1.5]	G. person
8. [0.5, 0.8, 0.8, 1.4]	H. king
9. [0.5, 0.9, 0.1, -0.3]	I. ruler
10. [0.5, 0.8, 0.7, 0.7]	J. princess
11. [0.5, 0.7, 0.4, 0.9]	K. son

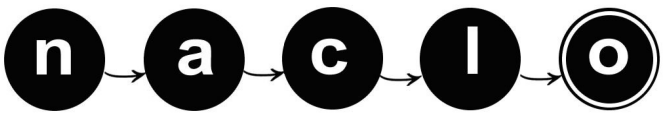

(R) GloVe Compartment (2/3)

R1. Match the vectors (1-11) to the English words they represent. Write the correct letter (A-K) in each box:

1) ☐ 2) ☐ 3) ☐ 4) ☐ 5) ☐ 6) ☐

7) ☐ 8) ☐ 9) ☐ 10) ☐ 11) ☐

Here are a few more word representations, along with their English equivalents, in no particular order. Note that a barometer is a tool for measuring air pressure, while a millibar is a unit of air pressure.

12. [0.3, 0, -0.6, -0.1]	L. clock
13. [0.2, -0.2, -0.3, -0.2]	M. first
14. [0.4, 0, -0.4, -0.4]	N. second
15. [-0.6, 0.6, 0.2, -0.8]	O. one
16. [-0.6, -0.2, -0.4, -0.4]	P. three
17. [0.4, 0.8, -0.4, -0.4]	Q. third
18. [1.6, 0, 1.8, 0.6]	R. two
19. [0, 0, 0, -0.4]	S. barometer
20. [-0.6, -0.4, -0.2, -0.4]	T. half
21. [1.8, 0, 1.6, 0.4]	U. millibar

R2. Match the vectors 12 through 21 to their English equivalents. There are two possible answers; either one will receive full credit. Write the correct letter (L-U) in each box:

12) ☐ 13) ☐ 14) ☐ 15) ☐ 16) ☐

17) ☐ 18) ☐ 19) ☐ 20) ☐ 21) ☐

R3. The word *third* actually has two meanings that are relevant to the problem. The vector that is given for *third* above is the average of the vectors that would represent these two meanings. Suppose English used two different words for these two meanings, rather than a single word. What would you expect the vector to be for each meaning?

Vector: ☐ Definition: ☐

Vector: ☐ Definition: ☐

(R) GloVe Compartment (3/3)

R4. Below are the two vectors found for the words *doctor* and *nurse*. Even though these words are gender-neutral, the method of defining a word based on the words that occur near it also captures general trends and biases that are in the texts which were used to determine what words occur near each other.

a) [0.5, 1.3, 0.3, 1.7]

b) [0.5, 1.3, 0.5, 1.5]

Identify which of these vectors goes with which word, and explain how the vectors encode gender-related properties of the corresponding words: