Neuroscience 299: Computing with High-Dimensional Vectors
Assignment 6: Analogical Reasoning
Due October 13, 1pm

Reminder: Please do *either* the writing assignment *or* the programming assignment. Expected length for the writing assignment is approximately 250-500 words, but there is no strict minimum or maximum.

Writing assignment:

Bongard problems. Your response should address the following three components:

1) What is a Bongard problem, and why is it an interesting problem for artificial intelligence?
2) Find a state-of-the-art approach to solving Bongard problems and determine the pros and cons of this approach. (This approach need not be based on HD Computing/VSA.)
3) Finally, suggest how HD Computing/VSA methods can help solve this problem. Sketch at a high level which components you would represent with high-dimensional vectors and how you would manipulate these vectors.

You may find the following links helpful:

- Wikipedia article: https://en.wikipedia.org/wiki/Bongard_problem
- List of Bongard problems: https://www.foundalis.com/res/bps/bpidx.htm
- Spencer Kent's thesis includes discussion of vector symbolic visual analogies: https://digitalassets.lib.berkeley.edu/etd/ucb/text/Kent_berkeley_0028E_20225.pdf (especially Sections 3.3-3.5)

Programming assignment:

You will replicate an estimate of analogical similarity from an experiment in Chapter 6 of Tony Plate's thesis:
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.48.5527&rep=rep1&type=pdf

For this implementation, we will use Holographic Reduced Representations (as you may have implemented in Assignment 2). You may re-use your existing implementation of basic operations, but you should double-check that your implementation is correct.

You are encouraged to read Chapter 6 of Plate's thesis for further information and background, but to simplify the assignment we have written out the steps explicitly.

Here are the steps you should complete:

1. Generate base vectors and token vectors (Plate uses n=2048, and angle brackets denote unit normalization):

Base vectors		Token vectors
person	**bite**	**jane** = $\langle$**person** + **id**$_{jane}\rangle$
dog	**flee**	**john** = $\langle$**person** + **id**$_{john}\rangle$
cat	**cause**	**fred** = $\langle$**person** + **id**$_{fred}\rangle$
mouse	**stroke**	**spot** = $\langle$**dog** + **id**$_{spot}\rangle$
	lick	**fido** = $\langle$**dog** + **id**$_{fido}\rangle$
bite$_{agt}$	**bite**$_{obj}$	**rover** = $\langle$**dog** + **id**$_{rover}\rangle$
flee$_{agt}$	**flee**$_{from}$	**felix** = $\langle$**cat** + **id**$_{felix}\rangle$
cause$_{antc}$	**cause**$_{cnsq}$	**mort** = $\langle$**mouse** + **id**$_{mort}\rangle$
stroke$_{agt}$	**stroke**$_{obj}$	
lick$_{agt}$	**lick**$_{obj}$	

Table 6.5: Base vectors and token vectors. All base and **id** vectors a randomly chosen with elements independently distributed as $N(0, 1/n)$.

2. Generate vectors for the following analogical episodes:

P: Spot bit Jane, causing Jane to flee from Spot.

Episodes in long-term memory:	
E1:	Fido bit John, causing John to flee from Fido.
E2:	Fred bit Rover, causing Rover to flee from Fred.
E3:	Felix bit Mort, causing Mort to flee from Felix.
E4:	Mort bit Felix, causing Felix to flee from Mort.
E5:	Rover bit Fred, causing Rover to flee from Fred.
E6:	John fled from Fido, causing Fido to bite John.
E7:	Mort bit Felix, causing Mort to flee from Felix.
E8:	Mort fled from Felix, causing Felix to bite Mort.
E9:	Fido bit John, John fled from Fido.
E10:	Fred stroked Rover, causing Rover to lick Fred.
E11:	Fred stroked Rover, Rover licked Fred.

The encoding should follow the following kind of structure (note that this is for ("P: Spot bit Jane…" and you will need to modify it for each of the episodes.):

$$\mathbf{P}_{bite} = \langle \mathbf{bite} + \mathbf{bite}_{agt} \circledast \mathbf{spot} + \mathbf{bite}_{obj} \circledast \mathbf{jane} \rangle$$

$$\mathbf{P}_{flee} = \langle \mathbf{flee} + \mathbf{flee}_{agt} \circledast \mathbf{jane} + \mathbf{flee}_{from} \circledast \mathbf{spot} \rangle$$

$$\mathbf{P}_{objects} = \langle \mathbf{jane} + \mathbf{spot} \rangle$$

$$\mathbf{P} = \langle \mathbf{cause} + \mathbf{P}_{objects} + \mathbf{P}_{bite} + \mathbf{P}_{flee} + \mathbf{cause}_{antc} \circledast \mathbf{P}_{bite} + \mathbf{cause}_{cnsq} \circledast \mathbf{P}_{flee} \rangle$$

Note: Please make your code for this section especially clear (with good commenting/syntax practices) so that it is easy to verify that your encodings are correct.

3. Compute the dot products between the target episode P and base episodes E1-11. Run multiple simulations (i.e., generate different random vectors) to get standard deviations of these results. Compare these results to those in Plate's thesis. Do you get similar results?

P: Spot bit Jane, causing Jane to flee from Spot.

Episodes in long-term memory:		Aspects of similarity OA	FOR	HOR	RFB	HOS	OLI	Type	Dot-products Avg	Sd
E1:	Fido bit John, causing John to flee from Fido.	✓	✓	✓	✓	✓	✓	LS	0.70	0.016
E2:	Fred bit Rover, causing Rover to flee from Fred.	✓	✓	✓	×	✓	✓	AN^{cm}	0.47	0.022
E3:	Felix bit Mort, causing Mort to flee from Felix.	×	✓	✓	×	✓	✓	AN_1	0.39	0.024
E4:	Mort bit Felix, causing Felix to flee from Mort.	×	✓	✓	×	✓	✓	AN_2	0.39	0.024
E5:	Rover bit Fred, causing Rover to flee from Fred.	✓	✓	✓	$\frac{1}{2}$	✓	×	$SS^{\times I}$	0.58	0.019
E6:	John fled from Fido, causing Fido to bite John.	✓	✓	✓	✓	×	✓	$SS^{\times H}$	0.47	0.018
E7:	Mort bit Felix, causing Mort to flee from Felix.	×	✓	✓	×	✓	×	$FA^{\times I}$	0.39	0.024
E8:	Mort fled from Felix, causing Felix to bite Mort.	×	✓	✓	×	×	✓	$FA^{\times H}$	0.28	0.025
E9:	Fido bit John, John fled from Fido.	✓	✓	×	✓	×	✓	SS^{-H}	0.43	0.019
E10:	Fred stroked Rover, causing Rover to lick Fred.	✓	×	✓	×	×	×	OO_1	0.25	0.024
E11:	Fred stroked Rover, Rover licked Fred.	✓	×	×	×	×	×	OO_2	0.12	0.023

Table 6.6: Results from Experiment 1. The averages and standard deviations are for dot-products of the probe with each episode in memory, over 100 runs. The aspects of similarity required for analogy are first-order relation names (FOR), higher-order relation names (HOR), higher-order structure (HOS), and object-level isomorphism (OLI).

The next four steps ask you to elaborate on the following aspects of your experiments. Explain your answers in about a paragraph per question. You are allowed, but not required, to draw inspiration from Plate's thesis, but follow proper citation practices.

4. Comment on the relative similarities (dot product) between P and different episodes. Is the dot product providing reasonable estimates of analogical similarity? (You may find it helpful to draw upon the "Aspects of similarity" table, and Plate's discussion, in your response.) Why does the similarity in the obtained results reflect the types of analogical similarity in the episodes?
5. What are the key design choices for the encoding so that they lead to meaningful results in this task?

6. Find at least one way to improve the representations used in this analogy task and explain why your method would be an improvement.
7. What are some potential issues that could occur when scaling this sentence encoding schema to real world analogical reasoning?