

Neuroscience 299: Computing with High-Dimensional Vectors
Assignment 2: Overview of different HD Computing/VSA models
Due September 15, 1pm

Writing assignment:

You have two topics to choose from. We ask you to focus on one topic and provide thorough support for your answer, rather than trying to answer both. Both topics have multiple good answers, and we encourage you to prioritize depth over breadth. Though the write-up should be done individually, we also strongly encourage discussing with your peers (for example, over Piazza)!

- a) *The importance of binding.* What is the binding operation useful for? What kinds of modeling, algorithms, or data structures does it enable? Here are some examples that you could discuss: the relation to the “binding problem” in neuroscience, advantages of variable binding in artificial intelligence, how binding operators answer certain challenges to connectionism.
- b) *The merits of different binding implementations.* In most HD/VSA models, the bundling operation is implemented using component-wise addition, possibly followed by a nonlinearity. Implementations of the binding operations, however, can have fairly different properties. Compare and contrast two HD/VSA models with different binding operations, explaining what the differences and the relative pros/cons of each approach.

Programming assignment:


By completing the first assignment, you already have one working HD computing/VSA model, which can be called Multiply-Add-Permute (with integer components). In this assignment, you will implement two other models and run an experiment to compare their performance.

1. *Implement the basic operations used in Holographic Reduced Representations.*

(a) Generate a random vector with n dimensions. Note that the elements of each vector should be independently and identically distributed (i.i.d.) with mean zero and variance $1/n$.

```
#n = dimensionality of vector
def generate_HRR_vector(n):
    return vector
```

(b) Recreate the binding and bundling operations:



$$\tilde{\mathbf{t}} = \tilde{\mathbf{c}} \otimes \tilde{\mathbf{x}}$$

$$t_0 = c_0x_0 + c_2x_1 + c_1x_2$$

$$t_1 = c_1x_0 + c_0x_1 + c_2x_2$$

$$t_2 = c_2x_0 + c_1x_1 + c_0x_2$$

$$t_j = \sum_{k=0}^{n-1} c_k x_{j-k}$$

for $j = 0$ to $n - 1$
(Subscripts are modulo- n)

Figure 4: Circular convolution represented as a compressed outer product for $n = 3$.

Definition of circular convolution (Plate, 1995).

```
def bindHRR(vector1,vector2):
    #your code here
    return bound_vector

def bundleHRR(vector1,vector2, vector3, ..., vectorN):
    #your code here
    return bundled_vector
```

(c) Recreate the calculation of the appropriate similarity metric operations (in this case, cosine similarity):

```
def simHRR(vector1,vector2):
    #your code here
    return similarity_between_vectors
```

2. Implement basic operations used in Fourier Holographic Reduced Representations. Complete the same steps (a), (b), and (c) as in Part 1. Note that for FHRR, the binding operation becomes element-wise multiplication (Hadamard product), vector **elements** are phasors ($x = e^{i\varphi}$, $\varphi \in [-\pi, \pi]$), and the similarity metric is angle distance.

3. Replicate the results reported in the experiment in Figure 4 of the focus paper “A comparison of Vector Symbolic Architectures” for the three HD computing/VSA models you have implemented so far.




Fig. 4. Minimum required number of dimensions to reach 99% accuracy. The solid lines represent linear fitted curves.

Here are the details on the experiment used to create this plot:

For each VSA, we create a database of $N = 1,000$ random elementary vectors from the underlying vector space V . It represents basic entities stored in a so-called item memory. To evaluate the bundle capacity of this VSA, we randomly chose k elementary vectors from this database and create their superposition $B \in V$ using the VSA's bundle operator. Now the question is whether this combined vector B is still similar to the bundled elementary vectors. To answer this question, we query the database with the vector B to obtain the k elementary vectors, which are the most similar to the bundle B (using the VSA's similarity metric). The evaluation criterion is the accuracy of the query result: the ratio of correctly retrieved elementary vectors on the k returned vectors from the database.

Do you get qualitatively the same results as in the focus paper?

Note: you do not have to fit straight lines to your results, just run the experiment several times (e.g., 10) and plot the average dimensionalities.