

Neural Computation (VS 265), Problem Set 4 - Attractor Networks

Due date: October 29, 3:30pm

Fall 2024

General guidelines:

- We are grading problem sets anonymously. **Include your student ID in the submission, but do not include your name.**
- You may work in small groups of 2-3. Note that you are responsible for writing up and submitting your submission individually.
- You are expected to attach any code you used for this assignment but will be evaluated primarily on the writeup.
- If you are including animations as part of your submission, please attach these files separately in your submission (it should not be necessary to run any code to view your animations).

Part 1: Hopfield Network



- Implement the discrete Hopfield (1982) network model and store the three simple 10×10 image patterns above (available on the course website) in the weight matrix. Next, corrupt one of the images by randomly flipping some percentage of the binary pixels. Feed the corrupted version as input to the network, and run the dynamics until equilibrium. Plot the energy as a function of each iteration to confirm it is descending. Make a figure showing side by side the initial corrupted image, an intermediate image between the network's initial and final state, and the final image once the network has converged. For small amounts of corruption, it should converge to the same image. Try corrupting the image to the point where the network no longer converges to the correct image, and make a similar figure to the previous one (initial, intermediate, and final state of the network).
- The *capacity* of a Hopfield network refers to the maximum number of patterns that can be stored and successfully retrieved from the network. To explore the capacity, create datasets containing different numbers of random patterns and store them in the weight matrix. To evaluate the model's retrieval accuracy, initialize the network state to one of the stored patterns, let it run to equilibrium, and then compute the Hamming distance between the final state and the initial state. Repeat this process multiple times for each dataset and plot the mean percentage of errors as a function of the number of stored patterns. You will need to choose the dimensionality of the network and the number of stored patterns for each dataset so that your plot traces out the regime from error-free retrieval to 50% bit errors.
- The example above (i) used the Hopfield network to store images. However, the Hopfield network can store any type of pattern. What other types of patterns might the brain want to store that a Hopfield network could facilitate?

Part 2: Ring Attractor Network

The dynamics of Kechen Zhang's (1996) ring attractor model are defined as:

$$\begin{aligned}\tau \dot{u} + u &= w * v \\ v &= \sigma(u)\end{aligned}$$

where $u = u(\theta, t)$ and $v = v(\theta, t)$ denote the synaptic input current and firing-rate, respectively, of the neuron coding heading θ at time t within the ring attractor network. $w = w(\theta, t)$ defines the weights between neurons as a function of their relative angular positions within the network, and $*$ denotes convolution. $\sigma(u)$ is the point-wise nonlinearity on each neuron converting synaptic input current to firing-rate. (See the paper for additional details and explanation.) For this part, you will implement the dynamics of this network.

- i. First, assume that $\sigma(x)$ is defined as $a \log^\beta(1 + e^{b(x+c)})$, with parameters $\beta = 0.8, b = 10, c = 0.5$, and $a = 6.34$. The next step is to find a set of weights w such that the network state v stabilizes to a single bump of activity. The solution is given by the following equation:

$$\hat{w}_n = \frac{\hat{u}_n \hat{v}_n}{\lambda + |\hat{v}_n|^2}$$

with v set the desired bump of activity and $u = \sigma^{-1}(v)$ (the inverse function of σ). \hat{w}_n, \hat{v}_n , and \hat{u}_n denote the n -th Fourier coefficients of w, v , and u respectively. λ is a constant that regularizes the solution. A solution for w with $\lambda = 0.01 \max(|\hat{v}_n|^2)$ (see figure 4B of the paper), appropriate for a ring attractor network of 64 neurons, is available on the website, but if you are familiar with Fourier transforms you may want to try solving for your own w using different values of λ .

- ii. Verify that your solution works by initializing the values of u randomly and show that the dynamics converge to a stabilized bump. Save an animation or show multiple timesteps of these updating dynamics for a few different random initializations. (Hint: see Figure 3 of the paper for the type of results you should expect to get.)
- iii. Next, add a shifting mechanism so that the bump position shifts over time. You can accomplish this by adding a small term to w proportional to the derivative of the weights (calculated numerically from the w above). For further details, please visit section 5 of the paper. Again, save an animation or show multiple timesteps of your result.