Final Exam 2022

(1.) Hopfield network

Implement a binary Hopfield network (HN) with $I = 400$ spins. Each spin can be set to $+1$ or -1.

The HN energy function is given by

$$E = -\frac{1}{2} \sum_{\substack{i,j \\ i \neq j}} w_{ij} x_i x_j \quad , \text{ where}$$

$x_i = \pm 1$ is the state of spin i, and w_{ij} are weights. Note that $w_{ij} = w_{ji}$, $w_{ii} = 0$ and there are no bias terms.

(a) Define $N = 2$ distinct memories (x_1 & x_2 I-dim vectors) (you can encode a black-and-white 20×20 image if you like) and compute w_{ij} using the Hebbian rule: $w_{ij} = \sum_{n=1}^{2} x_i^{(n)} x_j^{(n)}$.

Start $N = 1000$ runs from randomly chosen spin configurations and find the minima of E by asynchronous updates (i.e., update one spin at a time and go through all spins in order).

What is the number of times exact memories 1 & 2 have been recovered? Plot a histogram of all runs as a function of the number of mismatches (spin flips) with respect to the closest memory:

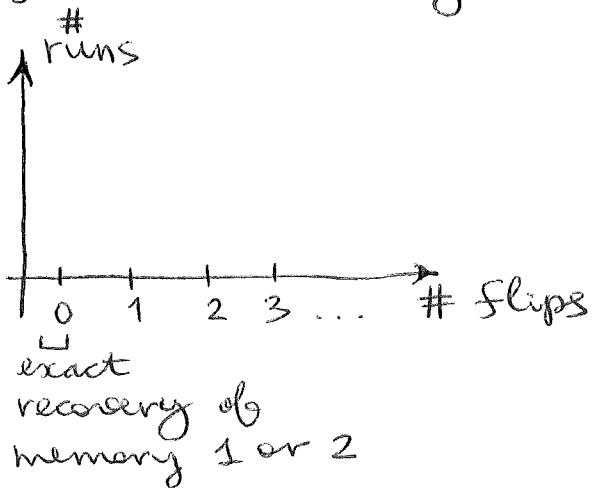

(b) Repeat part (a) with $N=6$ distinct memories. Have you reached the regime in which HN is overloaded?

(c) Now, set 20% of all weights w_{ij} " w_{ji} to $\emptyset$ randomly.

Repeat the analysis of part (a) with 2 memories originally introduced there, starting from the same initial spin configurations for consistency. Has HN been able to recover the memories after suffering the deletion of weights?

(2.) Binary classification

Implement a single-hidden-layer ANN for $K=2$ classification. Download dataset #8 (ionosphere) from the link provided on the website (568). Set 50 randomly chosen observations as your test set and train the ANN on the rest of the data. Try at least 2 activation functions and experiment with M, the number of hidden nodes. Report prediction accuracy on the test set as a function of M, for both activation functions. Note: Run the ANN 10 times for each parameter combination ($\leftarrow \{M, \text{activ. f'n}\}$) and record the best result for the plots mentioned above. What is the globally best prediction accuracy? Were you able to beat the 94% mentioned on the website (dataset)?