Final Exam (2020)

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 (you can encode a black-and-white 20x20 image if you like) $\underbrace{x_1 \,\&\, x_2}_{I\text{-dim vectors}}$ 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:

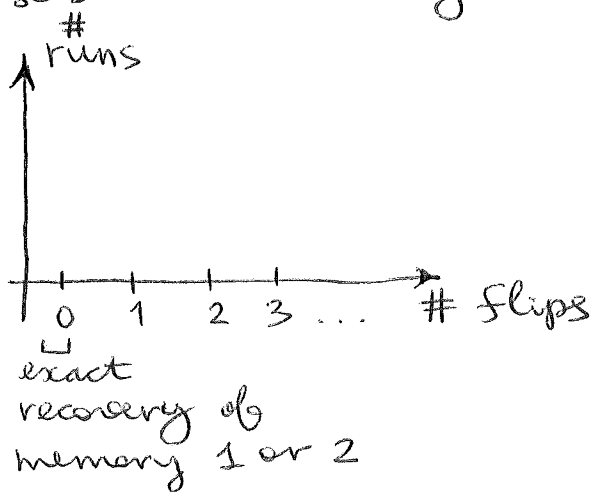

(b) Repeat part (a) with N=6 distinct memories. Have you reached the regime in which HN is overloaded?

(c) Now, set 20% of all weights $w_{ij} = w_{ji}$ to ∅ 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: identifying phases of the 2D Ising model

(a) Read about the 2D Ising model on a 40x40 square lattice dataset in Mehta et al., 2018 (pp. 30-32), and download the data from the link provided on the course website.

(b) Use logistic regression to classify the Ising model samples in the dataset into ordered ($T/J \lesssim 2.26$) and disordered ($T/J \gtrsim 2.26$). [J defined in Eq. 178)]
Clearly explain your approach, including regularization terms and optimizers used.
Divide the data into training & test sets in 3 separate categories:

- strongly ordered ($T/J < 2.0$)
- near-critical ($2.0 \leq T/J \leq 2.5$)
- strongly disordered ($T/J > 2.5$)

(i) Combine training data from all 3 categories, fit the model, and report classification accuracy separately for training and test sets in each category as a function of the regularization parameter λ.

(cf. Fig. 21, p. 32).

(ii) Repeat the analysis in (i) but train the model only on the ~~[illegible]~~ training sets in strongly ordered & strongly disordered categories (i.e., make the entire near-critical dataset a test set).

Compare the results of (i) and (ii).

Note: this is a binary classification task, all datasets are assigned a '0' label if $T > T_c \simeq 2.26\,J$, and a '1' label if $T < T_c$, where T_c is the Onsager critical temperature.