Neural Networks and Biological Modeling

Professor Wulfram Gerstner
Laboratory of Computational Neuroscience

Miniproject: Brain damage in the Hopfield Model

Assistant: Friedemann Zenke

Aim and motivations As you all know the mammalian brain is a remarkable computing device. Two very peculiar features of the brain are its error correcting properties and fault tolerance under various complications. Even when your brain is drugged or damaged, under a wide range of conditions it can still perform its tasks (although the results might be less accurate).

The aim of this mini-project is to get a glimpse of these remarkable properties in the framework of the Hopfield model. As you already saw in class the optimal synaptic weights for storing a set of patterns can be achieved using a standard Hebbian learning rule and random patterns. A Hopfield network can be seen as the prototype of holographic memory. Memories are non-local and each memory is distributed over the whole weight matrix. In this exercise we will explore the consequences of unreliable connections and "brain damage' – that is where a substantial part of the connections are removed completely – for memory retrieval.

Description: We consider Hopfield's model of associative memory. It consists of a layer of N fully interconnected units, with binary activities $S_i(t) \in \{-1; 1\}$. The network dynamics are described by

$$S_i(t+1) = \text{sign}\left(\sum_{j=1}^{N} w_{ij} c_{ij} r_{ij}(t) S_j(t)\right) \tag{1}$$

where the synaptic weights w_{ij} are given by summing over all μ patterns ξ^μ

$$w_{ij} = \frac{1}{N} \sum_{\mu=1}^{p} \xi_i^\mu \xi_j^\mu \tag{2}$$

where $\xi_i^\mu = 1$ with probability 0.5 and $\xi_i^\mu = -1$ otherwise.

Note that we extended the standard Hopfield model with the connectivity matrix c_{ij} and the reliability matrix $r_{ij}(t)$. The connectivity matrix is defined as

$$c_{ij} \begin{cases} 1 & : \text{with probability } \gamma \\ 0 & : \text{otherwise} \end{cases}$$

where we furthermore require that $c_{ij} = c_{ji}$ and that $c_{ij}(t = 0) = c_{ij}(t) = c_{ij} \quad \forall t$.

The reliability matrix r is defined similarly

$$r_{ij}(t) \begin{cases} 1 & : \text{with probability } \eta \\ 0 & : \text{otherwise} \end{cases}$$

However r is not necessarily symmetric and it is renewed at each discrete timestep.

Exercise 1: Getting started

Implement a simulation of the above Hopfield network. Use a network of size $N = 200$. To establish a reference simulate the network with $\gamma = \eta = 1$ i.e. as if there were neither a connectivity nor reliability matrix. Determine the maximum capacity of this network. To test the capacity store p random patterns ξ in the weight matrix. Attempt to recall each stored pattern ξ by starting the initial network dynamics S at a distorted pattern ξ' and letting the network relax to a stable state. This stable state corresponds to a pattern recall. You obtain the distorted pattern by randomly flipping 10% of the bits. How can you determine if the network has settled? If you tolerate a maximum of 5% bit errors of the recalled pattern (and you do so for all of the p patterns), how many patterns can you safely store in the network? As result quote the ratio between p_{max} and N

$$\alpha_N = \frac{p_{\mathsf{max}}}{N}$$

Repeat the simulation at least 10 times and give the mean capacity and standard error. Compare your result with the value from the literature.

Measure α_N for $N = 100, 300$ and for a third value of N as large as you can possibly simulate in a reasonable amount of time. What can you say about α?

Hint: If you are experiencing performance issues you might want to check out the function `dot` from the `numpy` package.

Exercise 2: Brain damage

At this point you should have a working associative memory. It is now time to study the effect of synapse loss. Repeat the above measurement for at least 10 different values of γ with $0 < \gamma < 1$ ($\eta = 1$). Ensure the effective weight matrix $w_{ij}c_{ij}$ is always symmetric. Plot your results with error bars and determine the capacity where 10% of all synapses are lost. Furthermore determine the γ_{50} where the maximum storage capacity drops to 50%. Quote the confidence interval for all your results. Is there a critical point where capacity is 0?

Exercise 3: Impaired synaptic reliability

A different scenario is one in which synapses are not lost but their transmission efficiency is impaired. By adjusting η you can control the synaptic transmission probability for each time step. Note that du to the definition of r_{ij} above, synaptic transmission is not necessarily symmetric any more[1]. Again plot the results for at least 10 different values of η ($\gamma = 1$). Show error bars and give explicit values for η_{50} and 90% reliability.

Exercise 4: The grand picture

Extend the above results by simultaneously varying γ and η and plot them in a two-dimensional color plot[2]. Plot at least 100 data points and comment on which mechanism seems to have the worse consequences. Can you explain why intuitively?

[1] Depending on how you defined convergence earlier you might have to reconsider this definition now.
[2] You might want to take a look at the methods `imshow` and `colorbar` from the `pylab` package