
13. Exploring A Hopfield Network (For Week 11)

13.1 Issue

How do Hopfield networks process information, and what kinds of tasks might they be able to accomplish? We attempt to gain some insight into such questions by exploring a Hopfield network that can be trained using four different stimuli.

13.2 Task

We will load an Excel spreadsheet that brings a Hopfield network to life. We will train this network to remember one or more different letters that are 'drawn' on its processing units. We will present stimuli to this network, possibly add different degrees of noise, and then relax the network to explore the kind of responses that it makes.

13.3 Materials

We will use the Hopfield.xlsm spreadsheet that is available from the course website.

13.4 Procedure

Students will download the Hopfield.xlsm Excel spreadsheet. When it is opened, it can be used to learn and recall information. To begin, press the 'Reset Network' button to make sure that the network is in a blank state. Then present a stimulus to the network using one of the 'Present Stimulus' buttons. You will see the stimulus drawn on the network's processors. Teach the network the stimulus by pressing the button 'Train With Hebb Rule'. You can press this button more than once to repeat training if you like.

You now can test learning by presenting some other stimulus to the network. If you like you can add noise to the stimulus. Once the desired stimulus has been drawn on the network's processors, press the 'Relax the Network' button. This will cause the Hopfield network to process the stimulus; it will eventually stop after no changes in processor activity have occurred after 150 epochs. When this happens, what is finally drawn on the processors? How long did it take the network to stabilize?

You can explore the network further by elaborating the procedure described above. For instance, you can present one stimulus, and train the network on it; then you can present a second stimulus, and train the network on it. Now the network knows about two of the possible stimuli. What happens when you present different stimuli to this network, and then relax it? Does it stabilize in more than one different final state? Have you affected processing time in any way?

You can proceed with this approach by training the network on three of the different available stimuli, or on all four. You can train the network on noisy versions of one or more stimuli. After training, you can explore the network's processing to various stimuli.

If you are really daring, you can draw your own stimuli in the four stimulus grids below the network. You have a small 5X5 grid in which to put a pattern – make the pattern by filling the cells with values of 1 or -1. How does the network behave when it learns your stimuli, or noisy versions of them?

14. Mental Rotation With A Hopfield Network (For Week 11)

14.1 Issue

Unlike many of the artificial neural networks that we have studied this term, Hopfield networks generate response latencies for free. In other words, the number of iterations required for a network to stabilize can be taken to be analogous to reaction time. So, it might be possible to use a Hopfield network to model cognitive tasks for which reaction time is the key dependent measure

14.2 Task

We will load an Excel spreadsheet that brings a Hopfield network to life. We will train this network on a stimulus in one orientation, and then present stimuli in various orientations to it. How long does it take the network to rate the stimulus into an upright orientation?

One famous use of reaction time can be found in Shepard and Metzler's study of mental rotation. They had subjects perform a same-different judgment about stimulus pairs; in order to make this judgment subjects had to mentally rotate one of the stimuli into the same orientation as the other.




Fig. 1. Examples of pairs of perspective line drawings presented to the subjects. (A) A "same" pair, which differs by an 80° rotation in the picture plane; (B) a "same" pair, which differs by an 80° rotation in depth; and (C) a "different" pair, which cannot be brought into congruence by any rotation.

Shepard and Metzler measured how long it took their subjects to make their judg-

ments, and plotted this as a function of the angular disparity between the two images. In one of the most famous results in experimental cognitive psychology they found a linear relationship between the independent and the dependent measure, which is illustrated below.

We won't use a Hopfield network to bring their experiment to life. However, we will explore the possibility of using this kind of network to model mental rotation by training it on stimuli in one orientation, and then presenting it stimuli in different orientations. How long does it take for the network to rotate the various stimuli? Is this related in any systematic way to angular disparity? Is the network performing mental rotation? If not, then what are the implications of this kind of simulation?




Fig. 2. Mean reaction times to two perspective line drawings portraying objects of the same three-dimensional shape. Times are plotted as a function of angular difference in portrayed orientation: (A) for pairs differing by a rotation in the picture plane only; and (B) for pairs differing by a rotation in depth. (The centers of the circles indicate the means and, when they extend far enough to show outside these circles, the vertical bars around each circle indicate a conservative estimate of the standard error of that mean based on the distribution of the eight component means contributed by the individual subjects.)

14.3 Materials

We will use the HopfieldRotate.xlsm spreadsheet that is available from the course website.

14.4 Procedure

Students will download the HopfieldRotate.xlsm Excel spreadsheet. When it is opened, it can be used to learn and recall information. To begin, press the 'Reset Network' button to make sure that the network is in a blank state. Then present a stimulus to the network using one of the 'Present Stimulus' buttons. In particular, present the network the upright 'T'. You will see the stimulus drawn on the network's processors. Teach the network the stimulus by pressing the button 'Train With Hebb Rule'. You can press this button more than once to repeat training if you like.

You now can test learning by presenting some other stimulus to the network. Note that there are eight stimuli all together, and each can be interpreted as a 'T' drawn at a particular angle (admittedly some are poorly drawn!). Present a stimulus and relax the network. How long did it take for the network to stabilize? Did it stabilize to the correct answer? Does it take the same amount of time to stabilize to the same stimulus?

Present the various stimuli to the network, and get some sense of the number of iterations required to stabilize for each stimulus. Use Excel to plot your results. What are the similarities and differences between your graph and the graphs from Shepard and Metzler that are provided above? What are the implications of your results?

Arguably there is something needed in this simulation to make it more realistic. What is it, and how in general might you provide it?

You can draw your own eight stimuli in the 5X5 grids below the network by putting your own values of -1 and 1 in each cell. Can you create your own mental rotation example? Does it generate similar results to the ones that you have already observed?