VS265 Neural Computation: Problem Set 7

Joshua Abbott
joshua.abbott@berkeley.edu

October 31, 2012

Collaborators: Anwar Nunez-Elizalde and Jessica Hamrick

1 Hopfield network exploration

1.a Setting T with the outer product rule

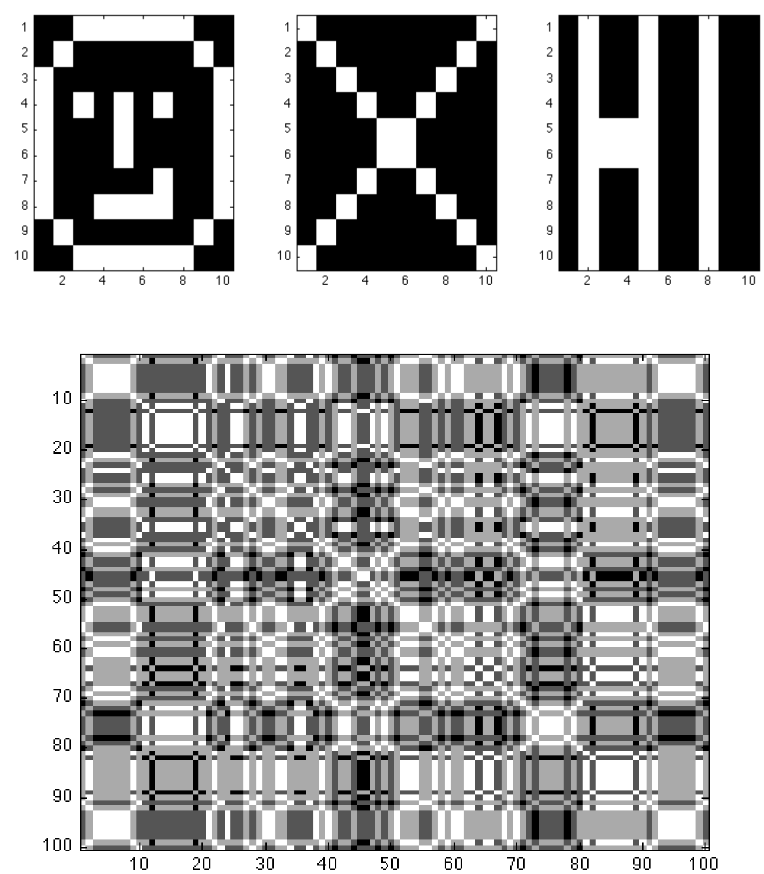

Figure 1: The original three patterns and the corresponding weight matrix T set using the outer product rule

1.b Basins of attraction for three patterns

Initializing the network to each of the stored patterns above results in stable basins of attraction. To explore the Hamming radius for each of the three patterns, I flipped k bits for each image pattern (where k ranges from 0 to 50) and initialized the network to each corrupted image for 100 simulations each (i.e., I corrupted the face pattern by switching k random bits and ran the network, taking note of whether it converged to the uncorrupted pattern, and I did this 100 times for each k to get an idea of how stable each basin was on average).

The results are presented below in Table 1. While the maximum Hamming distances found were a bit lower than the averages (as expected given that there are (50 choose k) possible k-bit flipped versions of each pattern...), all three patterns could generally have up to 38 bits flipped and still be recovered correctly.

Pattern	**Avg. Max Hamming Distance**	**Max Hamming Distance Found**
face	38.53	23
X	38.34	17
hi	38.95	25

Table 1: Hamming radii for three basins of attraction

1.c Capacity of our Hopfield network

Here we created a set of twenty patterns (original three patterns and block letters of "A" through "Q", see Figure 2 below) to explore the capacity of the network.

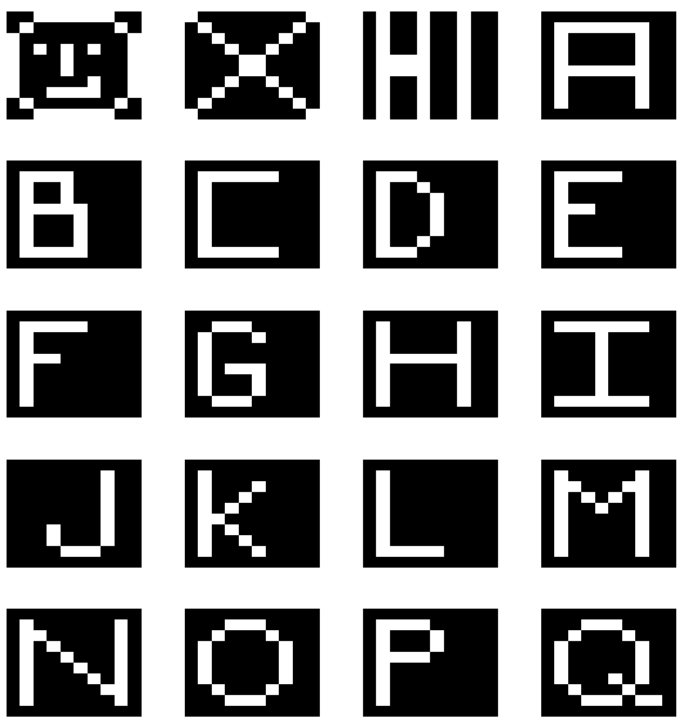

Figure 2: The set of patterns used to explore the capacity of the network.

We find that the network can only handle one additional pattern on average (face, X, hi, and A). However, if we randomly select n patterns from our set of twenty, the number of stable basins varies. This is due to the correlations between patterns. In Figure 3 below we find that the original three patterns have close to zero correlation, while the rest of the patterns (A - Q) are positively correlated for the most part.

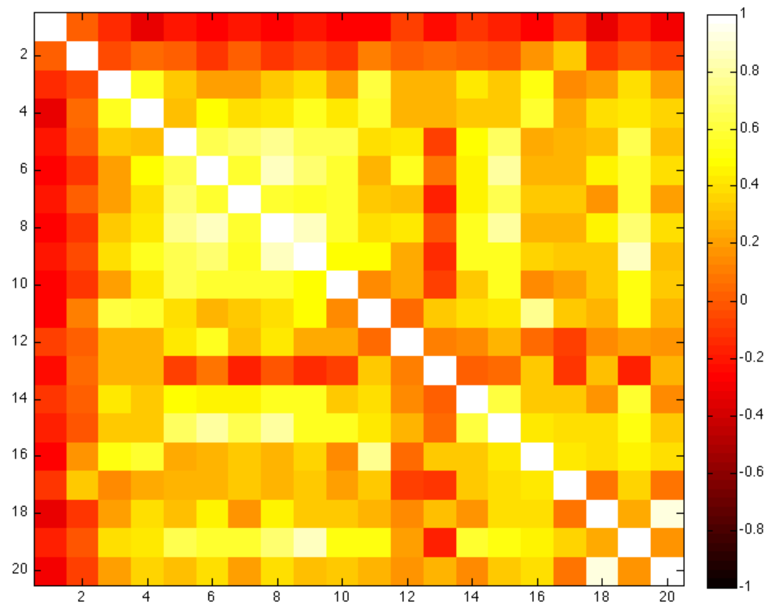

Figure 3: The correlation matrix for the set of patterns used to explore the capacity of the network.

Table 2 below shows the effect Hamming radius around each basin a more patterns are stored (up to 8 patterns for brevity). The entries (i,j) in the table correspond to the average Hamming radius for basin j in a network trying to learn i patterns (averaged over 100 simulations as before). Here, a -1 entry, corresponds to a pattern not being learned at all. Since the average Hamming radius for pattern #'s 3, 4, and 5 are -1 for a network trained on five patterns, this demonstrates the network can only learn one additional pattern from the original three.

	Pattern 1	**Pattern 2**	**Pattern 3**	**Pattern 4**	**Pattern 5**	**Pattern 6**	**Pattern 7**	**Pattern 8**
1 basin	49.20	0	0	0	0	0	0	0
2 basins	41.90	43.20	0	0	0	0	0	0
3 basins	36.40	36.10	38.90	0	0	0	0	0
4 basins	30.30	24.00	12.00	15.90	0	0	0	0
5 basins	24.90	2.50	-1.00	-1.00	-1.00	0	0	0
6 basins	19.80	-1.00	-1.00	-1.00	-1.00	-1.00	0	0
7 basins	19.10	-1.00	-1.00	-1.00	-1.00	-1.00	-1.00	0
8 basins	13.80	-1.00	-1.00	-1.00	-1.00	-1.00	-1.00	-1.00

Table 2: Average Hamming radii for n basins of attraction

2 Exploration of input correlation and capacity

2.a Correlation versus capacity

Sample stimuli and results are displayed below for five different sparsity levels (10%,20%,30%,40%,50%). Using 50% sparsity, the network could successfully store up to 10 basins of attraction (see Table 7). I tried 20 random patterns for 5 different sparsity levels using the Hamming radii finding method outlined above in Problem 1.c. The results in Tables 3-8 display only the first 10 basins/patterns for sake of brevity (and to allow it to fit on the page!)

Figure 4: A sample set of sparse patterns used to explore the capacity of the network. Displayed left to right, the sparsities are at 10, 20, 30, 40, and 50 % of the image.

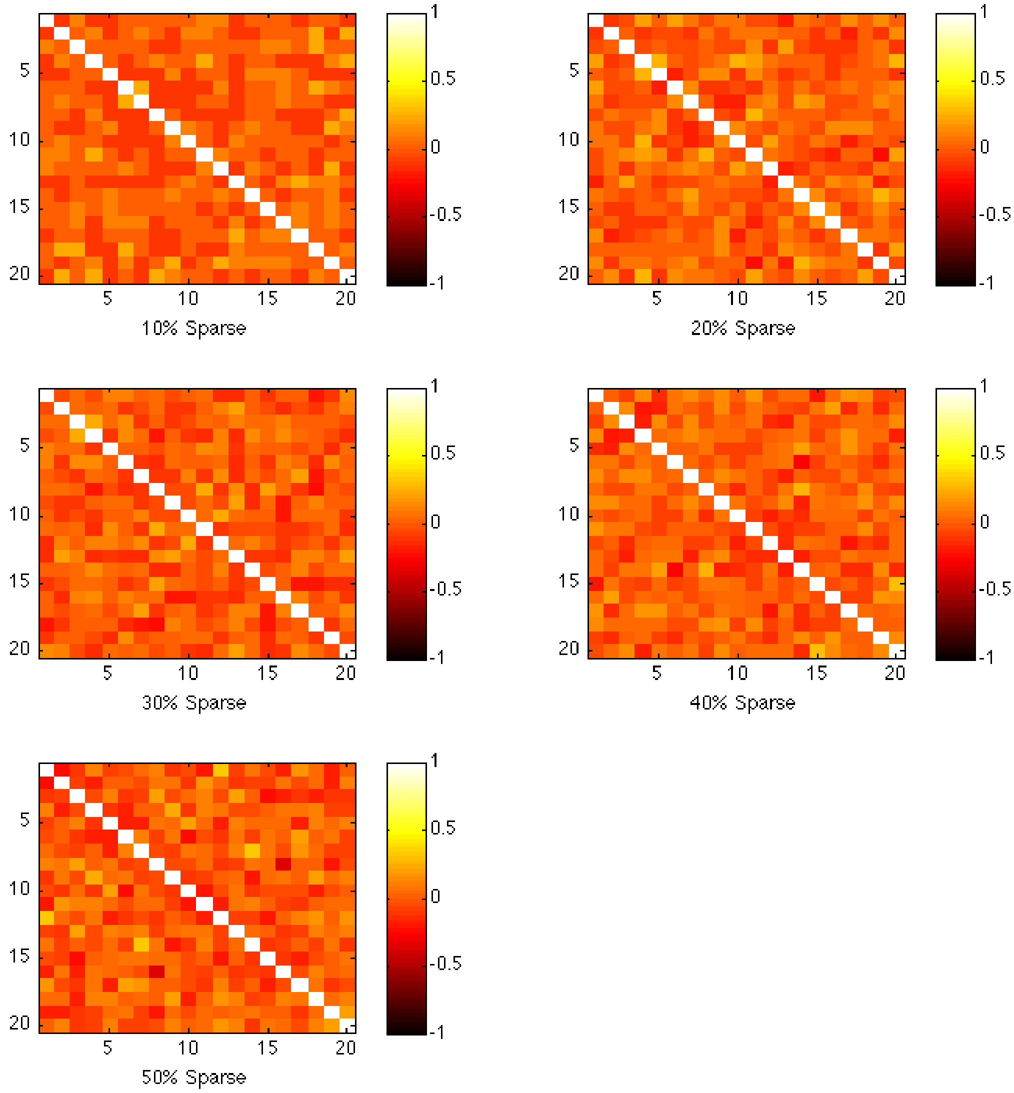

Figure 5: The correlation matrices of the sparse patterns used to explore the capacity of the network.

	Pat 1	Pat 2	Pat 3	Pat 4	Pat 5
1 basin	49.9000	0	0	0	0
2 basins	35.2000	38.1000	0	0	0
3 basins	-1.0000	-1.0000	-1.0000	0	0
4 basins	-1.0000	-1.0000	-1.0000	-1.0000	0
5 basins	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000

Table 3: Average Hamming radii for n basins of attraction at 10% sparsity

	Pat 1	Pat 2	Pat 3	Pat 4	Pat 5
1 basin	49.7000	0	0	0	0
2 basins	42.2000	44.3000	0	0	0
3 basins	24.2000	20.4000	19.2000	0	0
4 basins	-1.0000	-1.0000	-1.0000	-1.0000	0
5 basins	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000

Table 4: Average Hamming radii for n basins of attraction at 20% sparsity

	Pat 1	Pat 2	Pat 3	Pat 4	Pat 5	Pat 6	Pat 7	Pat 8	Pat 9	Pat 10
1 basin	48.1000	0	0	0	0	0	0	0	0	0
2 basins	42.2000	42.7000	0	0	0	0	0	0	0	0
3 basins	35.7000	35.2000	37.8000	0	0	0	0	0	0	0
4 basins	23.5000	26.7000	29.5000	26.1000	0	0	0	0	0	0
5 basins	26.2000	22.1000	25.5000	18.2000	21.5000	0	0	0	0	0
6 basins	12.4000	11.4000	11.1000	13.2000	8.2000	5.7000	0	0	0	0
7 basins	1.0000	5.1000	6.8000	4.3000	-1.0000	2.2000	-1.0000	0	0	0
8 basins	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	4.4000	-1.0000	-1.0000	0	0
9 basins	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	1.3000	-1.0000	-1.0000	-1.0000	0
10 basins	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000

Table 5: Average Hamming radii for n basins of attraction at 30% sparsity

	Pat 1	Pat 2	Pat 3	Pat 4	Pat 5	Pat 6	Pat 7	Pat 8	Pat 9	Pat 10
1 basin	49.9000	0	0	0	0	0	0	0	0	0
2 basins	42.0000	43.8000	0	0	0	0	0	0	0	0
3 basins	39.1000	38.5000	37.9000	0	0	0	0	0	0	0
4 basins	38.3000	37.0000	36.8000	35.7000	0	0	0	0	0	0
5 basins	37.6000	34.1000	36.6000	35.2000	35.8000	0	0	0	0	0
6 basins	33.7000	35.4000	35.4000	33.3000	34.8000	33.3000	0	0	0	0
7 basins	34.0000	34.1000	30.3000	31.7000	34.8000	33.3000	32.3000	0	0	0
8 basins	30.9000	30.9000	29.5000	29.1000	30.0000	33.8000	30.2000	27.9000	0	0
9 basins	7.8000	28.4000	26.0000	26.8000	28.3000	28.8000	27.6000	24.6000	27.1000	0
10 basins	3.8000	27.5000	27.9000	29.4000	26.1000	28.1000	28.0000	19.6000	24.2000	28.9000

Table 6: Average Hamming radii for n basins of attraction at 40% sparsity

2.b Energy landscape

Assuming the results above are correct, I started exploring random input with a network trained on 10 patterns of level 50% sparsity and with 100 random initializations, the network never converged to a basin that was a stored pattern. This should be expected since the stored patterns and the input are both random (and hence mostly decorrelated).

	Pat 1	**Pat 2**	**Pat 3**	**Pat 4**	**Pat 5**	**Pat 6**	**Pat 7**	**Pat 8**	**Pat 9**	**Pat 10**
1 basin	49.3000	0	0	0	0	0	0	0	0	0
2 basins	41.7000	42.9000	0	0	0	0	0	0	0	0
3 basins	36.6000	37.9000	37.1000	0	0	0	0	0	0	0
4 basins	36.5000	35.2000	34.2000	35.7000	0	0	0	0	0	0
5 basins	29.8000	31.8000	30.7000	31.5000	28.3000	0	0	0	0	0
6 basins	26.1000	22.1000	28.3000	32.3000	28.3000	30.7000	0	0	0	0
7 basins	24.4000	31.3000	27.0000	28.9000	22.2000	28.6000	32.7000	0	0	0
8 basins	28.2000	27.2000	26.5000	30.4000	22.2000	25.1000	30.4000	27.7000	0	0
9 basins	30.7000	27.6000	25.1000	28.5000	19.6000	25.8000	27.7000	28.8000	30.7000	0
10 basins	28.2000	27.1000	13.1000	22.2000	2.2000	18.9000	27.0000	24.3000	25.7000	11.8000

Table 7: Average Hamming radii for *n* basins of attraction at 50% sparsity